Essays in the Economics of Subjective Well-Being

Glenn Goldsmith

Balliol College, University of Oxford

A thesis submitted in fulfillment of the requirements for the degree of Doctor of Philosophy in Economics

Trinity Term, 2009

Approx 99,000 words

Essays in the Economics of Subjective Well-Being

Glenn Goldsmith, Balliol College

Doctor of Philosophy in Economics

Trinity Term, 2009

This thesis explores three major issues in the burgeoning empirical literature on the determinants of subjective well-being (SWB).

While economic theory assumes that it is current consumption that matters to SWB, empirical work has focused almost exclusively on the effect of income. In Part I, we use household panel data from Russia and Britain to show that neither the standard theoretical account, nor the standard empirical practice may be adequate. Consumption, income, and wealth each contribute separately to SWB, in particular via perceptions of status and anticipation of the future; and omitting consumption from SWB equations significantly understates the importance of money to SWB.

Distinguishing between consumption and income is also important to identifying reference effects. In Part II, we confirm earlier findings that others' income has a positive (informational) effect on SWB in Russia, but show that others' consumption has an offsetting negative (comparison) effect. The net effect depends on how we define individuals' reference groups. We develop a novel econometric model that lets us estimate these reference groups from the data. Contrary to previous results, we conclude that comparison dominates information.

Most SWB analyses focus on the average effects of money, relationships, and other outcomes across a given population; yet there may be significant differences in what is important to different people. In Part III, we employ parametric and semi-parametric random coefficient models to show that there are large differences in the determinants of individual SWB in Britain, and (in contrast to previous work) that such differences cannot simply be attributed to differences in individuals' reporting functions. While individual differences correlate with (some) observable demographic variables, they do not generally correlate with individuals' perceptions about what is important to them. The results of SWB research may therefore be a useful source of information.

Acknowledgements

This thesis could not have been completed without the help and support of a number of people. While not wishing to implicate them in its remaining failings, I would especially like to thank my late MPhil supervisor, Dr. Gavin Cameron, and my DPhil supervisor, Professor John Knight — without whose encouragement and advice it would doubtless have even more.

The RLMS (Phase II) data used in this thesis were funded by the USAID and NIH (R01-HD38700), Higher School of Economics and Pension Fund of Russia, and provided by the Carolina Population Center and Russian Institute of Sociology. The BHPS data were originally collected by the ESRC Research Centre on Micro-social Change at the University of Essex (now incorporated within the Institute for Social and Economic Research), and made available through the ESRC Data Archive. Neither the original collectors of the data nor its distributors bear any responsibility for the analyses or interpretations presented here.

Last, but by no means least, a special thanks Aruna Sathanapally, for her patience, understanding, and support — and for keeping me smiling throughout.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

Individuals the world over spend significant amounts time and energy striving to improve their circumstances — presumably because they think that this will improve their lives in some sense. For their part, economists, and the governments they advise, have typically been content to accept such revealed preferences as indicative (indeed, in some cases definitive) of the true value of various outcomes, economic and otherwise — after all, if individuals think that something will make them happy, who are economists to argue? In recent years, however, there has been increasing interest — amongst both economists and other social scientists — in using self-reports of subjective well-being (SWB) to determine what makes people happier, or more satisfied with their lives (see e.g. Oswald, 1997; Kahneman, Diener, and Schwarz, 1999; Van Praag and Ferrer-i Carbonell, 2004; Layard, 2005a; Frey, 2008).

This thesis explores three major issues in the burgeoning empirical literature on the determinants of subjective well-being (SWB). This introduction provides a short overview of these in Section 1.1, before moving on, in Section 1.2 to discuss two foundational issues underlying SWB research more generally: the nature of SWB and our interest in it; and the means by which it can be measured. We suggest in particular that while it seems clear that self-reported SWB measures have the potential to significantly advance our understanding of human well-being, exactly how useful they are will inevitably depend on both our underlying normative commitments, and the specific uses to which the measures are put.

1.1 Overview

The first issue this thesis seeks to address relates to age-old question of whether money can buy happiness. While economic theory assumes that it is current consumption that matters to SWB, empirical work has focused almost exclusively on the effect of income. In Part I, we use household panel data from Russia and Britain to show that neither the standard theoretical account, nor the standard empirical practice may be adequate. Consumption, income, and wealth each contribute separately to SWB, in particular via perceptions of status and anticipation of the future; and omitting consumption from SWB equations significantly understates the importance of money to SWB.

Distinguishing between consumption and income is also important to the second focus of the thesis: identifying reference effects. In Part II, we confirm earlier findings that others' income has a positive (informational) effect on SWB in Russia, but show that others' consumption has an offsetting negative (comparison) effect. The net effect depends on how we define individuals' reference groups. We develop a novel econometric model that lets us estimate these reference groups from the data. Contrary to previous results, we conclude that comparison dominates information.

The subject of the final two chapters is individual differences. Most SWB analyses focus on the average effects of money, relationships, and other outcomes across a given population; yet there may be significant differences in what is important to different people. In Part III, we employ parametric and semi-parametric random coefficient models to show that there are large differences in the determinants of individual SWB in Britain, and (in contrast to previous work) that such differences cannot simply be attributed to differences in individuals' reporting functions. While individual differences correlate with (some) observable demographic variables, they do not generally correlate with individuals' perceptions about what is important to them. The results of SWB research may therefore be a useful source of information to them.

1.2 Justifying SWB Research

The rise of interest in SWB within the dismal science has not been without its critics (see e.g. Bertrand and Mullainathan, 2001; Johns and Ormerod, 2007; Wilkinson, 2007). SWB research — especially in the survey-based form typically practised by economists — is based on two potentially contentious assumptions: (1) that we care about individuals' subjective well-being; and (2) that we can adequately measure it. A great deal of ink has been spilled addressing various aspects of these assumptions, and we shall not engage in an extended defence them here. Nonetheless, sections 1.2.1 and 1.2.2 present a brief discussion of each in turn. The two major themes that we wish to emphasise relate firstly to the range of potential conceptions of well-being in which we may be interested, and secondly to the importance of both our underlying normative commitments, and the specific purposes of our analysis, in determining the degree to which we are able to accurately measure well-being. It seems clear that SWB measures can usefully add to our knowledge, and make a valuable contribution to both policy debates and individual decision-making. Nonetheless, such measures do have certain weaknesses, which require us to be careful in our uses of data and our interpretation of results.

1.2.1 Conceptions of SWB

This is not the place for anything resembling a rigorous defence of the notion that the pursuit of individual well-being (somehow defined) should be the primary goal of either individuals, or of public policy. Indeed, no such defence is required to justify an interest in SWB research. What is required is simply that subjective well-being be amongst the things we care about — a minimum condition which seems relatively innocuous. As Sen (1985) has observed, "[i]t would, of course, be altogether amazing if moral goodness had nothing to do with well-being... the question is not whether well-being is an intrinsically important variable for moral analysis, but whether it is uniquely so." Most people would agree that, while they might be willing to sacrifice their well-being in the name of other goals, other things equal, they would like to lead happy, satisfying lives. The point of SWB research is to provide information that may assist them (and the governments who serve them) in doing so.[1]

The major challenge to the importance of SWB derives less from concerns about the importance of well-being *per se*, and more from doubts

[1]Even if we did not care about happiness as an end in itself, we might still care about it instrumentally. For example, Clark, Georgellis, and Sanfey (1998) find that job satisfaction negatively predicts quit rates. To the extent that this is true, we may also care about the determinants of job satisfaction for their behavioural consequences. However, this use of happiness research is somewhat tangential to the point of this thesis.

about whether the most appropriate notion of well-being is *subjective* (as opposed to objective). The attraction of subjective theories seems to be due to two related notions: on the one hand a sense that whatever is good for a person be required somehow to *resonate* with them; and on the other, the complementary idea that whatever resonates with a person should (at least presumptively) contribute to their well-being. There are no doubt those who fail to feel the pull of these intuitions, but they, or something like them appear to be fairly generally accepted (in particular by economists), and to lie behind most attempts in the philosophical literature to flesh out various theories of well-being in greater detail (as depending, say, on pleasure and pain, or on the fulfillment of individual desires, to take Parfit's (1984) two most prominent examples).

Most cases in which people have felt it necessary to depart from subjective views are based on specific cases when individuals' subjective judgments seem to go astray — either because of false beliefs,[2] flawed reasoning, or, perhaps most importantly, individuals adjusting their standards downwards when they find themselves in undesirable situations. Sen (e.g. 1987) in particular has been a vocal critic of happiness-based theories of well-being on this final

[2]A particularly relevant example for SWB research is the finding that individuals' SWB appears to depend in part on perceptions of mobility and expectations about the future (see e.g. Senik, 2004; Clark and Senik, 2008). If such perceptions are misguided, then one might legitimately question the extent to which they should be considered "good for" individuals, if at all.

ground — arguing, quite reasonably, that respecting such adjustments in standards can serve to perpetuate unjust disadvantage.[3]

Nonetheless, it does not seem necessary to completely abandon subjective conceptions of well-being in order to accommodate such concerns. Indeed, there is a well-established tradition in the philosophical literature that attempts to deal with such mistakes by shifting focus from individuals' *actual* (and therefore potentially mistaken) subjective attitudes, towards a set of *idealised* ones. While such idealisations can vary in their details, the basic idea is straightforward: to define well-being in terms of the responses of individuals who are fully-informed, rational, and/or whose evaluative standards are not artificially constrained by their situation.[4] Such idealised subjective theories offer the promise of correcting obvious mistakes, while remaining true to individuals' own underlying values.[5]

Of course, even within the space of subjective theories, there still remains a wide variety of possible notions of SWB one might adopt. Such notions can be distinguished along at least three dimensions.[6] Probably the main cleav-

[3]On the other hand, there does seem to be an important sense in which those who emotionally adapt to impoverished circumstances are better off than those who do not. Attempting to delineate exactly what sort of adjustment in standards should be deemed problematic and what should not is not a straightforward task. It nonetheless seems that an adequate theory of well-being should attempt to accommodate both of these concerns.

[4]For recent examples of such views, see e.g. Brandt (1979), Railton (1986), Griffin (1986), Sumner (1996), and Brülde (2007b), but the basic idea stretches back to Mill (1861) and Sidgwick (1913).

[5]Indeed, there is a sense in which correcting such mistakes can be seen as seeking to uphold, rather than undermine, individuals' true values.

[6]The potential distinctions one could draw here are legion. See e.g. Wilkinson (2007),

age between different subjective theories is the type of subjective attitude they consider relevant to well-being judgments. A general, if not entirely satisfactory breakdown can be given in terms of what Brülde (2007b) calls "affective" and "cognitive" theories. While *affective* notions of SWB focus on individuals' feelings or emotional states,[7] *cognitive* conceptions of SWB — as the name suggests — instead hold that it is an individual's cognitive evaluation of her situation that determines her SWB.[8,9]

Theories can also differ in what aspects of an individual's situation they focus on. While there is a general consensus that the attitudes in question must somehow pertain to the subject's *own* situation,[10] it is difficult

Brülde (2007a,b), and Haybron (2007) for further possibilities.

[7]Examples of affective theories include: the classical Benthamite, or *narrow hedonist* view which takes pleasure and pain as the determinants of utility; the *preference-hedonist* view, which focuses on a potentially wider range of *desired* or *preferred* mental states, and which was proposed by Parfit (1984, p.493) to accommodate the argument that "the pleasures of satisfying an intense thirst or lust, listening to music, solving an intellectual problem, reading a tragedy, and knowing one's child is happy... do not contain any distinctive common quality... [except] their relation to our desires." (In fact, while most authors seem to have gone along with this argument (see e.g. Griffin, 1986, p.8), modern advances in neuroimaging may throw doubt on whether it is in fact true.); and *thymic* or *mood-state* views of the sort developed by Haybron (2005), and which focus on one's propensity to experience particular affective states over the medium-term.

[8]Depending on how one chooses to draw distinctions, such views can be taken to include, *inter alia*: *desire-fulfillment* theories, which hold that one's life goes well to the extent that one's desires are fulfilled, and is probably the most prominent example of this type of view; and *value-correspondence* views, which focus on the degree to which one's situation aligns with one's "deeper" values (as opposed to potentially fleeting, or unconsidered desires: see e.g. Griffin (1986), whose "informed desire" account of well-being seems to appeal to such a deeper sense of "value").

[9]As Brülde (2007b) notes, one can also imagine *hybrid* theories, that combine two or more of the individual theories noted above. Indeed, the sense in which "SWB" is used in the psychological literature typically combines affective and cognitive components in this way.

[10]This does not rule out the potential importance of others' outcomes to SWB via either altruism or jealousy, but does suggest that if such outcomes matter to an individuals'

to describe precisely the way in which this should be the case, and different approaches to the issue are possible. Some prefer SWB notions whereby only things one actually experiences matter (e.g. Sumner, 1996), while others reject this "experience requirement" (e.g. Parfit, 1984; Griffin, 1986).[11] Conceptions can also differ in whether they focus on attitudes towards one's life *as a whole* or only on parts of it, such as attitudes held *at a particular time*, or in relation to particular *domains* of one's life (e.g. one's health or finances). Finally, as we have already suggested, different notions of SWB may differ in the extent to which they respect individuals' attitudes, even when these appear to be based on some form of mistake.

Though we shall not defend it here, our own preferred notion of SWB for normative purposes is one that is cognitive, global, rejects the experience requirement, and idealises away from mistakes of fact, reasoning, and artificial constraints on individuals' evaluative standards. In broad terms, we are interested in the extent to which an individual's life measures up to her considered values. As we discuss in more detail in Section 1.2.2, this (or any other set of normative commitments) will have implications for the conclusions we can draw from different measures of SWB — although, as we

SWB, it must be because they directly affect that individual in some way (for example, by influencing their emotional state).

[11] Parfit, for example, asks us to suppose that one of his strongest desires is to be a successful parent. If his children have wretched lives because of mistakes he makes in raising them, then — Parfit argues — his life goes badly for him, even if he does not realise this (say, because his failure as a parent only becomes apparent after his death).

also note, even for a given set of normative commitments, different measures of SWB may also be more or less useful for different purposes.

1.2.2 Measuring SWB

The more contentious question than whether (some notion of) SWB is important is whether we can measure it with sufficient accuracy — in our case via survey self-reports — to be useful. Once again, there is a great deal of existing literature on this topic, which we do not propose to review in any great detail here. What we want to emphasise instead is the fact that what we require of different measures of SWB depends heavily on the purpose to which they are to be put. In this section, we note firstly how different measures of SWB will have particular strengths and weaknesses depending on our underlying normative commitments. We also explain our preference for the the life-satisfaction measures which we use as our dependent variables in the remainder of the thesis. We then go on to briefly discuss how the different ways in which we might utilise these measures will require us to make different assumptions about their comparability. In particular, our use of panel data requires us to assume that our measures are inter-temporally comparable. What other assumptions we will require will vary depending on our purpose.

Alternative Measures of SWB

Just as there are a variety of notions of SWB in which we might be interested, there are also a variety of alternative SWB measures, each of which will approximate different well-being notions to varying degrees. In particular, such measures may: (1) take different aspects of individuals' lives as their objects; (2) focus on different evaluative attitudes; and (3) be idealised to a greater or lesser extent. The closer a measure matches our preferred SWB notion along each of these dimensions, the more useful it generally will be. Nonetheless, practical concerns may sometimes make slightly less well-matched measures more useful. For example, time-specific measures of SWB will often be more useful for regression purposes than global evaluations of an individuals' entire life-course, even if the latter better matches one's preferred SWB concept.

Survey-based studies, and in particular those conducted by economists, have often (and partly of necessity) been content to play fairly fast and loose with alternative measures — variously identifying SWB with either self-reported satisfaction with one's life, job, or other economic circumstances, answers to the question "how happy are you?", or measures of mental health (such as the General Health Questionnaire). In this regard, researchers have no doubt been comforted by the frequent assertion that these measures tend to track each other quite closely: to the extent that this is true, equivocation over the appropriate well-being concept, while not immaterial, is unlikely to be crucial. At the same time, however, both evidence and introspection suggest that various well-being concepts can in some cases diverge markedly, and an appreciation of these divergences may be important both to the conduct, and to the interpretation of SWB research.

Given the normative commitments set out in Section 1.2.1 above, our preferred SWB measure would ideally be one that focuses on individuals' assessments of how well (i) their entire lives measure up against (ii) their considered values and (iii) does so in a way that avoids as much as possible potential sources of bias or error (such as artificially constrained aspirations). Of the measures available in the Russian and British data that we employ in the remainder of the thesis, those that most closely approximate these criteria are the global life-satisfaction questions.[12] These measures nonetheless depart from our ideal in three main ways.

First, although we might ideally prefer to focus on individuals' entire lifespans, this is neither feasible from a surveying perspective, nor particularly useful from an econometric one. Because we rely on inter-temporal variation in our measures to identify the effects of various factors on SWB, we must necessarily focus on measures of current well-being.[13] This means that we

[12] These questions, and the data in general, are described in more detail in Chapters 2 and 3 respectively.

[13] And indeed, the measures we employ contain specific time-referents asking individuals to focus on their lives "at present" and their "current situation".

are unable to garner much information about how individuals' assess different time-patterns of SWB from a holistic perspective.

Second, our measures focus on the degree to which individuals are "satisfied" with their lives. While this framing of the question seems to appropriately target a cognitive evaluation, Haybron (2007) has argued that such satisfaction judgments may actually combine cognitive evaluation ("thinking things satisfactory") with a potentially less cognitive notion of "actually *being* satisfied". He argues that this second aspect means that satisfaction reports may not precisely track the sort of cognitive evaluation in which we are most interested. Unfortunately, while the multi-item Satisfaction With Life Scale (SWLS) of Diener et al. (1985) appears to address this potential shortcoming, we do not have access to it here.[14] On the other hand, the correlations between the pure satisfaction component of the SWLS (which corresponds roughly to the single-item questions in the RLMS and BHPS) and its other components seem strong enough to allay severe doubts about

[14] The SWLS contains 5 items, and asks respondents to indicate the degree to which they agree with the following statements:

(1) In most ways my life is close to ideal
(2) The conditions of my life are excellent
(3) I am satisfied with my life
(4) So far I have gotten the important things I want in life
(5) If I could live my life over I would change almost nothing

The first two of these in particular seem to be phrased closer to what we would ideally like (the last two items have different temporal scope). Items (1) and (3) have the highest factor loadings for the single largest factor (0.84 and 0.83 respectively) and also have the highest item-total correlations (both 0.75).

the measures we use here.

Third, the measures to which we have access are not idealised in any sense, suggesting that we may need to be awake to the possibility that they could be compromised in particular circumstances. As we discuss in the next section, for example, it is possible that life-satisfaction judgments could be affected by changes in individuals' reporting scales.[15]

The usefulness of results derived from SWB measures other than global life-satisfaction may need to be interpreted carefully if they do not match our preferred normative focus. Because most alternative measures employed in the literature — including measures of affective states, as well as satisfaction in various of life's sub-domains — are likely to contribute to a global notion of life-satisfaction, they should be somewhat informative about its determinants. Nonetheless, they will not necessarily be especially useful in capturing the relative importance of different outcomes where the effects may differ across different sub-domains.

To take an example, discovering that income is important to financial satisfaction is unlikely to be especially informative about its value in the context of one's life as a whole.[16] Similarly, it seems to us that the re-

[15] Again, the first item of the SWLS scale noted in Footnote 14 seems to better address this. By focusing on the extent to which an individual's life is "close to *ideal*" [emphasis mine] the measure may get around potential problems with constrained aspirations noted in Section 1.2.1 below. Correlations between this measure and the pure satisfaction item suggest that this may not be a huge problem, but it may still we worth bearing in mind.

[16] We might think that it would at least identify the sign of the effect, if not its relative

cent finding by Kahneman et al. (2006) that, while the rich report higher life-satisfaction, they "are barely happier than others in moment-to-moment experience", may need to be interpreted with some caution. Kahneman et al. (2006) interpret the divergence as being due to "focusing illusion" that causes life-satisfaction to systematically overstate the importance of money. However, while their results clearly suggest that life-satisfaction is a biased measure of net affect, this does not directly address the issue of whether the reported life-satisfaction judgments are mistaken. Mistake is clearly one possible explanation of this result; but another is simply that individuals value more than just the aggregation of their momentary mood assessments.[17]

Conversely, if one's preferred well-being concept differs from ours, then the life-satisfaction measures we employ will have a distinct set of target-mismatch problems. On the one hand, the general level of noise in the measures is likely to increase. On the other, there are particular biases that may be either introduced or alleviated if we adopt alternative views of the appropriate SWB concept. (For example, as we just noted, life-satisfaction judgments will probably overstate the importance of money to net affect.)

importance. However, if higher income also reduced satisfaction in other areas, then even this could be ambiguous.

[17]One way to test the focusing illusion hypothesis would would be to see whether individuals do in fact alter their cognitive evaluations or their behaviour when presented with the sort of findings outlined by Kahneman et al. (2006). Evidence of such changes would support the idea that their initial assessments were mistaken.

The Comparability of SWB Measures

As Ferrer-i-Carbonell and Frijters (2004) note, the purpose to which we wish put our SWB measures will also affect the assumptions we will need to make about their comparability. Abstracting from issues of measurement error, there are at least four assumptions we might wish to make.[18]

(C1) *Monotonicity*: $u_{it} = f_{it}(w_{it})$, where the reporting function, $f_{it}(\cdot)$ is a positive, (weakly) monotonic transformation, invariant only for the same individual i at a single point in time t (implying that $u_{it} > u'_{it} \Rightarrow w_{it} > w'_{it}$).

(C2) *Inter-temporal Ordinal Comparability*: $u_{it} = f_i(w_{it})$, where the reporting function, $f_i(\cdot)$ is a positive, (weakly) monotonic transformation, invariant over t for any given i (implying that $u_{it} > u_{is} \Rightarrow w_{it} > w_{is}$).

(C3) *Interpersonal Ordinal Comparability*: $u_{it} = f_t(w_{it})$, where the reporting function, $f_t(\cdot)$ is a positive, (weakly) monotonic transformation, invariant over i for any given t (implying that $u_{it} > u_{jt} \Rightarrow w_{it} > w_{jt}$).

(C4) *Cardinal Comparability*: $f_{it}(w_{it}) - f_{it}(w'_{it}) = g_{it}(u_{it}, u'_{it})$, where $g_{it}(\cdot)$ is known up to a multiplicative constant (usually assumed to take the form $g_{it}(u_{it}, u'_{it}) = (u_{it} - u'_{it})$).

[18] The discussion here is based on Ferrer-i-Carbonell and Frijters (2004), who distinguish between three assumptions. Their assumption (A1) is similar to our (C2); their (A2) is a conjunction of our (C2) and (C3); and their (A3) of (C2), (C3) and (C4). As far as we can tell, they do not clearly distinguish between (C1) and (C2).

Regardless of our purpose, we will require measures that satisfy (C1), but which of the other assumptions are necessary may vary. If our purpose were to obtain descriptive measures of two populations' relative levels of well-being (say for the purpose of prioritising aid to the worst-off) then we would need to assume at the very least (C3), and would be greatly assisted by (C4). On the other hand, if our purpose is simply to qualitatively investigate the correlates and determinants of subjective well-being, somewhat different concerns arise: different estimators will be available to us depending on which of these assumptions we can sustain. These are set out in Table 1.1 below.

Table 1.1: Assumptions Required for Alternative Estimators

Estimators			**Assumption(s)**
(1)	Ordinal Cross-Section	(Ordered) Logit/Probit	(C3)
(2)	Ordinal Panel Data	(a) Fixed Effects Ordered Logit (Ferrer-i-Carbonell and Frijters, 2004)	(C2)
		(b) Pooled (Ordered) Logit/Probit, Conditional Logit, Das and Van Soest (1999)	(C2) & (C3)
(3)	Linear Cross-Section	OLS	(C3) & (C4)
(4)	Linear Panel Data	Pooled OLS, FE, RE, Mundlak (1987), GMM	(C2), (C3) & (C4)

We discuss the issue of the choice of estimator in more detail in Chapter 2. However, it is apparent that our use of panel data constrains us to estimators of types (2)(a), (2)(b) or (4), and thus, at the very minimum, we need to assume inter-temporal ordinal comparability, (C2). On the other hand, we can, if we wish, often avoid the need to rely on interpersonal (C3) or cardinal (C4) comparability.

Ferrer-i-Carbonell and Frijters (2004) provide a good discussion of the evidence for these various assumptions, and indeed provide some further evidence of their own — showing that the results of regressions imposing interpersonal (C3) or cardinal (C4) comparability appear very similar to those imposing only inter-temporal comparability, (C2). We do not wish to traverse this well-travelled ground again. However, a couple of points seem worth drawing out.

The first is that Ferrer-i-Carbonell and Frijters (2004) actually appear to conflate assumptions (C1) and (C2), and the evidence they cite in support of the latter is for the most part only evidence in favour of the former. As far as we are aware, there seems to be little evidence directly attempting to validate SWB measures over time. While some support for the idea of inter-temporal comparability can be found in the fact that the results of SWB regressions relying on it tend to be both fairly consistent and intuitive, it would be useful to have something stronger to rely on.

One particular concern around inter-temporal comparability is raised by Wilkinson (2007), who suggests that findings that we adapt quite quickly to higher levels of income might instead be explained by re-norming of the scales that individuals use to report SWB.[19] While this does suggest some caution in interpreting findings of adaptation, it bears noting that scale re-norming will result in *common* variation in observed adaptation to different outcomes, and will therefore have difficulty explaining observed differences in adaptation between different variables (such as those found by Di Tella, Haisken-De New, and MacCulloch, 2007, who show that individuals appear to adapt differently to income and status).

A related point has also been made recently by Oswald (2008), who argues that the existing evidence for the concavity of SWB in income could instead be due to concavity in the function that maps "true" SWB into reported SWB.[20] As with the possibility of scale re-norming, this will affect estimates of all coefficients simultaneously. Thus, a finding that the marginal SWB of income declines more quickly than, say, leisure may still be reliable, even if we cannot identify the rates of decline in absolute terms. Nonetheless, Oswald (2008) seems correct to suggest that the form of the reporting function may deserve more attention than it has received — clearly the conclusions that we

[19]This concept is similar, but not identical to the idea of a *satisfaction*, or *aspiration* treadmill, suggested by e.g. Kahneman (1999). But whereas the satisfaction treadmill suggests that individuals may re-norm their satisfaction scales relative to their underlying emotional states, the scale re-norming hypothesis set out here applies more generally to the re-norming of any reporting scale relative an underlying "true" well-being measure.

[20]Though he also presents some evidence that the reporting function for subjectively perceived height is linear, as he notes, it is not entirely clear whether this will generalise.

can draw from SWB research will depend quite crucially on the assumptions we can justifiably make about it.

1.2.3 Conclusion

The preceding discussion has provided a brief justification of the life-satisfaction measures that we employ to track SWB in the remainder of the thesis, both in relative terms as against alternative self-report measures, and more generally. It seems clear that such measures have the potential to contribute significantly to our knowledge about the sources of human well-being. However, as we have suggested, different measures will have different strengths and weaknesses, depending on one's underlying normative commitments. Moreover, the precise conclusions we can draw from any given analysis will often hinge on both these normative commitments, and what further assumptions we are willing to make about our measures.

Part I

The Effect of Consumption, Income, & Wealth on SWB

Chapter 2

Russia

The age-old question of whether money "buys happiness" has been a major focus of SWB research. Yet, while economic theory typically assumes that it is current consumption that matters to current SWB, empirical work has focused almost exclusively on the effect of income. In Part I of this thesis, we use household panel data from Russia and Britain to suggest that both the standard theoretical account and the standard empirical practice may be inadequate.

A few existing papers have examined the SWB impact of non-income measures of material well-being: Headey and Wooden (2004) and Brown, Taylor, and Wheatley Price (2005) focus on net wealth and short term debt respectively, while Headey, Muffels, and Wooden (2008) examine the impact of various combinations of consumption, wealth and income. Part of our contribution here lies in extending this sort of analysis to an as yet unstudied context (Russia), and in attempting to improve on the methodologies of these previous studies (in particular by controlling for unobservable heterogeneity). However, our major focus is more substantive: none of the existing work has really attempted to address the underlying reasons why income and/or wealth may matter to SWB, and the results obtained generally seem to have been interpreted as supporting — rather than undermining — the standard consumption-based model (see e.g. Clark, Frijters, and Shields, 2008). In contrast, we attempt to provide a fuller account of both *how* and *why* income and wealth matter to SWB. While our results support previous studies in suggesting that income coefficients may substantially understate the importance of money to SWB (by 40-80% in Russia, and by more than 600% in Britain), they also suggest a number of further conclusions. First, the consumption-based model seems inadequate to explain the impact on SWB of income and wealth: income and wealth do not seem to simply proxy for consumption (as e.g. Clark, Frijters, and Shields, 2008, suggest). Instead, they appear to matter to SWB independently — in particular by contributing to status, self-esteem, and/or a sense of security about the future.

The existence of these indirect channels to SWB has a number of interesting implications. One such implication is to lend support to a perspective on inter-temporal decision-making that focuses on the importance of anticipatory emotions as both a source of current utility and a motive for e.g. savings and investment behaviour. Prior to the emergence of the the discounted utility model, such a view was championed by W.S. Jevons (1888) and especially H.S. Jevons (1905) — and has recently been resurrected in a number of behavioural models (see e.g. Caplin and Leahy, 2001; Köszegi, 2005).

Another is that different types of income and wealth may have different effects on SWB. In Russia, for example, we find that earned income has a greater effect than unearned income, a finding that we attribute to the greater sense of status or self-esteem it may bring. In Britain, meanwhile, we find that different types of wealth have markedly different effects on SWB: wealth that is more closely linked to the consumption of specific items in the near future has a greater impact than less-specific, or longer-term wealth holdings.

Finally, these results also demonstrate the importance of context to determining the impact of income or wealth on SWB — as evidenced by the large differences between the importance of income in Russia and Britain, as well as some systematic (albeit smaller) differences within the British sample itself. As we discuss later, possible reasons for this divergence include both a greater focus on future economic security, and/or a greater focus on income as a source of status, in Russia.

In this chapter, we focus on the Russian case, before moving on in Chapter 3 to look at Britain. Section 2.1 begins by describing the apparent disconnect between the consumption focus of traditional economic theory and the income focus of standard practice in SWB research, and suggests a number of reasons to suspect that neither may fully capture the effects of money on SWB. Section 2.2 discusses the existing work on this topic in more detail, and explains how we hope to build on it. Section 2.3 describes the Russian context and the RLMS data that we employ, while Section 2.4 sets out in broad terms the basis of our empirical strategy. Section 2.5 presents evidence that consumption, income and wealth each matter independently to SWB in Russia. Section 2.6 shows that status and security are both important in explaining these effects, and that, consistent with the status view, earned income appears to have a greater impact on SWB than unearned income. Section 2.7 concludes.

2.1 Consumption, Income, Wealth, & SWB

According to standard economic theory, it is not money *per se* that brings us satisfaction, but the flow of goods and services — the food, clothing, shelter, cars, TVs and so on — that it allows us to purchase. This is reflected in the standard model of inter-temporal decision-making, which assumes that individuals maximise, at each point in time, the (discounted) sum (or integral) of period utilities over their lifetime T: $U_i(t) = \sum_{\tau=t}^{T} u(c_{i\tau})$. In this formulation, period utility is assumed to be a function only of contemporaneous consumption, c_{it}, and the purpose of both income and wealth is simply to enable such consumption, either now or in the future.[1]

Unfortunately, good quality consumption data can be difficult to come by, so the equations typically estimated in the SWB literature tend to focus on income instead. Under the standard view outlined above, we would expect such equations, which are essentially employing noisy proxies for consumption, to underestimate the importance of money (see e.g. Weinzierl, 2006; Clark, Frijters, and Shields, 2008, for expressions of this view).[2] On the other hand, SWB research is also notable for its focus on sources of well-being beyond the merely material, and these provide reasons to suspect that income and wealth may in fact have effects on SWB that are independent of consumption. As we discuss below, income and wealth could affect SWB

[1]While some models do assume income as the argument in the utility function, it is typically as a proxy for consumption, under the simplifying assumption that the two are equal when individuals do not save. Similarly, lifetime wealth may be used as the argument in a lifetime utility function, as, under certain assumptions (e.g. lack of credit constraints), it will determine the sum or integral of a maximising agent's instantaneous utilities.

[2]Much of the theoretical work on which the happiness literature is based appears to uphold this traditional focus on consumption. The likes of Veblen (1899) and Frank (1984, 1999) place great emphasis on the phenomenon of "conspicuous consumption" (though Frank (2003) speaks instead in terms of relative *income*), while Hirsch's (1976) focus on "positional goods" also suggests a somewhat consumption-centric approach. Even Duesenberry (1949), who is so often cited as the father of the "relative *income*" hypothesis, was clear in his formalisation that it was relative *consumption* that featured in the utility function (see e.g. p.32).

via status, self-esteem, or anticipatory emotions — indeed, they could even do so directly. There actually seems little reason in theory to assume that consumption is the only argument in the SWB function.

Status

How we think we are perceived by others is undoubtedly a major influence on SWB. And while we need not go as far as Fourcade and Healy (2007) — who claim that "in modern society money is central to the evaluation of the moral worth of individuals" — it nonetheless seems reasonable to suppose that money should influence this sense of status.[3]

On the other hand, there is little reason to restrict our attention specifically to consumption in this regard. While it may be true that certain types of consumption are more visible than income or wealth, and may therefore provide a focal point for "demonstration effects" and status concerns (see e.g. Veblen, 1899) it seems implausible to suggest that others' income and wealth are entirely invisible. Not only are certain types of wealth (such as housing) visible, individuals will often be aware of the incomes of others in their industries, workplaces, or social circles,[4] and for this reason both income and

[3]While talk of status typically connotes a concern with relative standing, our usage here need not do so. If the respect accorded us by others depends on our absolute, and not our relative position, then "status" in our sense will inherit this absolute quality.

[4]Certainly, much of the job satisfaction literature assumes as much (see e.g. Brown et al., 2006; Clark and Oswald, 1996), as does Senik's (2004; 2008) recent work providing empirical support for Hirschman and Rothschild's (1973) "tunnel effect" (the thesis that

wealth seem capable of contributing to status.

Self-Esteem

In contrast to status, which reflects others' perceptions of us, self-esteem reflects our perceptions of ourselves.[5] Evidence suggests a strong link between self-esteem and SWB (for a review, see Bauermeister et al., 2003). In the data considered by Diener and Diener (1995) and Furnham and Cheng (2000), self-esteem emerges as the strongest predictor of life-satisfaction (though Diener and Diener, 1995, report that the correlation is weaker in more "collectivist" countries). A number of other studies confirm high correlations between the two concepts (e.g. DeNeve and Cooper, 1998; Shackleford, 2001; Lyubomirksy, Tkach, and DiMatteo, 2006), and there also appear to be strong associations between low self-esteem and depression and anxiety (again see e.g. Bauermeister et al., 2003). While it is not entirely clear how much of this reflects a causal effect from self-esteem to life satisfaction, there is at least a strong intuitive basis to suspect causation.

Similarly, there are good reasons to think that income and wealth should exert an influence on self-esteem, even independently of status concerns. To the extent that income is seen as an earned reward for one's work and achieve-

increases in others' income may improve SWB by acting as a signal of future prospects).

[5]Of course, the two will be related to the extent that we internalise others' perceptions of us (Mead, 1934; Cooley, 1902).

ments, it is likely to enhance one's self-perception. In similar vein, if income represents the flow of one's current achievements, wealth may in a sense represent the stock of one's past achievements, suggesting that both may contribute positively to one's self-image. This conjecture is consistent with what evidence exists on the link between socio-economic status (SES) and self-esteem, and in particular, the finding that the connection between the two appears to rise with age (see e.g. Rosenberg and Pearlin, 1978; Gecas and Seff, 1990; Twenge and Campbell, 2002): one way to explain such results is that the observed weak relationship at younger ages (and particularly amongst children) is due to the fact that SES is less a function of one's own achievements when young, and that the perceived connection between SES and achievement rises over time.[6]

Anticipation

The third reason income and wealth may matter to SWB is through anticipation, in particular via emotions such as hope and anxiety. The link between such emotions and SWB should be fairly obvious (indeed, anxiety is typically considered a major form of subjective *ill*-being). Recognising this, a body of work within behavioural economics has recently begun to incorporate such

[6] Other explanations are of course possible; in particular, Twenge and Campbell (2002) argue that variation in social reference groups can explain much of the variation in the link between SES and self-esteem, even if other factors play a role.

anticipatory concerns into decision models, arguing that such factors can help to explain a number of otherwise puzzling phenomena: from why individuals sometimes choose to experience negative events sooner and positive events later (Loewenstein, 1987; Elster and Loewenstein, 1992); to the risk-free rate and equity premium puzzles (Caplin and Leahy, 2001).[7]

These models typically define anticipation as a function of the distribution of future consumption possibilities, which enters the utility function independently of current consumption: $u_{it} = u(c_{it}) + v(\Phi_t(c_{t+1}, \ldots, c_T))$. Nonetheless, it is easy to see how income and (especially) wealth are likely to affect such possibilities. On the one hand, wealth confers security about future consumption in the face of expected or unexpected shocks; on the other, to the extent that there is persistence in income or wage processes, current income may confer security, by acting as a signal of future income prospects (and thus, future consumption prospects). (There is an obvious parallel here to the "tunnel effect" proposed by Hirschman and Rothschild (1973), except that here, it is own income rather than others' that is presumed to provide information about the future.)

In contrast to the traditional discounted utility model (which models inter-temporal decision-making as a detached weighing of future and current utility), such anticipatory utility models recall the older Jevonian (1888;

[7]See also Köszegi (2005); Köszegi and Rabin (2006, forthcoming).

1905) perspective, which held that the primary motivation for savings, investment, and other future-oriented activities was the utility they brought in the present.

Direct Importance

It is also possible that either income or wealth could be directly important to SWB, independently of the channels just discussed. It may be that we simply like getting or having money: for example, Knutson et al. (2003) find evidence that merely receiving money activates areas of the brain associated with positive emotion, while Schwarz (1990) reports that finding a 10 cent piece on a photocopier increases reports of life-satisfaction, presumably due a temporary elevation in mood. It is unclear to what extent we should expect such effects to generalise into longer term mood, or broader cognitive judgments (and indeed, the experiment reported by Schwarz (1990) is usually interpreted as evidence against the reliability of life-satisfaction measures, which we would prefer to exhibit somewhat more stability). Nonetheless, if a high income, or the accumulation of wealth is a consciously chosen goal for any individual, then they could presumably gain satisfaction from such achievements even if this did not contribute independently to their status, self-esteem, security, or mood.

Variation

As shown in Figure 2.1, each of the (indirect) channels described above involves a two-stage mapping from the primary outcome(s) with which we are concerned (here income or wealth) to SWB. The first stage is from the primary outcome to the relevant intermediate outcome (status, self-esteem, or security). The second is from these intermediate outcomes to overall SWB. Both these mappings will likely be somewhat idiosyncratic: for example, income may be more important to some individuals' self-esteem than others; and some individuals may value either status or security more highly.[8] However, we may also expect some systematic element to this variation: in particular, different *types* of income and wealth may be more or less important in different *contexts*.

Figure 2.1: Channels via which Economic Outcomes May Affect SWB

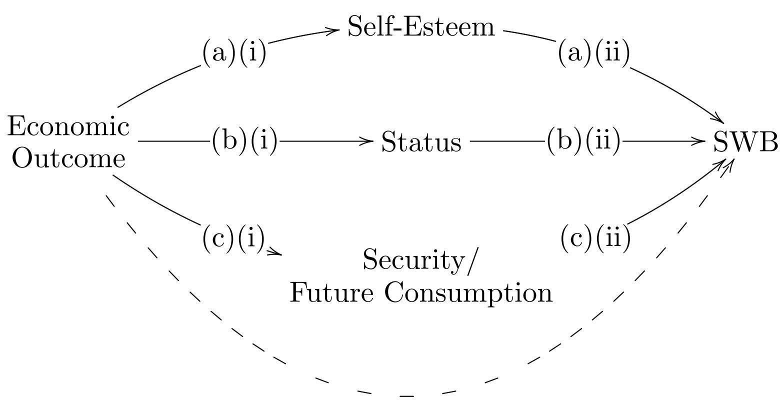

[8]We take up the issue of individual-level variation in more detail in Part III.

To take an example from the existing literature, Senik (2004, 2008) hypothesises (and finds supporting evidence) that the positive effect of (suitably similar) others' income will be more important to SWB in highly volatile economic environments. There are two potential reasons for this increased importance: on the one hand, in a volatile environment, the average income of similar others may provide a stronger signal of one's own future prospects, strengthening the link (c)(i); on the other, volatile environments may also lead individuals to care more about their future economic security, thereby strengthening link (c)(ii). Clearly, this latter effect could also affect the impact of one's own income or wealth on SWB. Similarly, there could be e.g. cultural differences in the importance of money to status, or in the importance of status concerns to SWB (perhaps in part as a result of advertising or other media).

Different types of consumption, income, and wealth could also affect SWB to varying degrees. For example, some have suggested that earned income may yield greater satisfaction than income received as, say, a benefit or subsidy (see e.g. Gecas and Seff, 1990).[9] This could be the case either because *others* view earned income more favourably (i.e. via the link to status in (b)(i) of Figure 2.1), or alternatively because higher earned income is asso-

[9]We discuss some other possibilities (related to anticipatory consumption) in greater detail in Section 3.3.

ciated with a greater *personal* sense of achievement or self-worth (i.e. via (a)(i) in Figure 2.1).[10]

2.2 Existing Literature

While it is not clear that they have recognised the apparent disjunction between the happiness literature and economic theory, Headey and Wooden (2004) and Headey, Muffels, and Wooden (2008) have recently argued that it may be important to include broader measures of material well-being in SWB equations, and, in particular, that focusing on income alone is likely to underestimate the importance of money to well-being. Headey and Wooden (2004) use cross-sectional household survey data from Australia to argue that wealth is more important to well-being than income. Meanwhile, Headey, Muffels, and Wooden (2008) use a broader range of survey data from Australia, Britain, Germany, the Netherlands, and Hungary to estimate SWB regressions that include consumption and/or wealth in addition to income, again finding evidence that using income alone is likely to understate the importance of pecuniary factors to SWB. In a similar vein, Brown, Taylor, and Wheatley Price (2005) employ BHPS data to examine the effect of sav-

[10]In this latter case, the extent to which work-based achievement is "central" to individuals' self-concept may also be important. As Gecas and Seff (1990) argue, those in boring, monotonous work are unlikely to invest a great deal of their identity in workplace achievement, and may therefore have a lower return to earned income than those whose work allows them greater scope for autonomy or "self-expression".

ings and debt on mental health, finding that both exert significant effects in addition to income.

Unfortunately, in part due to data limitations, the results in Headey and Wooden (2004) and Brown, Taylor, and Wheatley Price (2005), as well as most of those in Headey, Muffels, and Wooden (2008) are based on simple or pooled cross-sectional regressions. As we shall see, such results are likely to be contaminated by omitted personality traits, and may significantly overstate the importance of pecuniary factors to SWB (see Ferrer-i-Carbonell and Frijters, 2004, discussed in greater detail in Section 2.4.2 below). While Headey, Muffels, and Wooden (2008) report some results from fixed effects regressions, these use satisfaction with one's standard of living as their dependent variable, and therefore seem likely to overstate the importance of money to global life-satisfaction.[11] These regressions also produce somewhat more mixed results. In Britain (where comprehensive wealth measures were not available in the panel) consumption was much more strongly related to financial satisfaction than income; while in the Netherlands, (where consumption measures were unavailable) wealth was much more strongly correlated with financial satisfaction than income. However, in Hungary (the only country for which all three measures were available in the panel), although wealth

[11]Their reason for omitting fixed effects life-satisfaction results is somewhat unclear, given that global life-satisfaction seemed to be their preferred dependent variable in their OLS specifications.

was again more strongly correlated with SWB than income, consumption attracted a significant negative coefficient, which is somewhat puzzling.[12,13]

This chapter and the next attempt to build on these results in a number of ways. They provide what is to our knowledge the first evidence that consumption, income, and wealth each significantly affect global judgments of life-satisfaction, even once we control for unobservable fixed effects. In line with previous work, these results confirm that focusing on income alone significantly understates the importance of money to SWB. But in contrast to previous results (which have been interpreted by, e.g. Clark, Frijters, and Shields, 2008, as lending support to the traditional consumption-based specification of utility), we also provide evidence that the effects of income and wealth are independent of consumption, and do not merely proxy for it.

In the course of this analysis, Chapter 3 also seeks to resolve a conflict between the two distinct approaches taken to the importance of wealth in the existing literature. While Headey and Wooden (2004) and Headey, Muffels, and Wooden (2008) focus on overall net wealth as the main driver of SWB,

[12]In fact, we suspect that part of the reason for this final result may be the narrow focus on *financial* satisfaction, rather than the broader concept of life-satisfaction. If an individual's satisfaction with her financial situation is assessed primarily as the degree to which her finances allow her to meet her consumption goals, the pattern of correlations would in fact be entirely unsurprising, even if consumption contributes positively to one's overall satisfaction with life.

[13]Headey and Wooden (2004) and Headey, Muffels, and Wooden (2008) make a few other specification choices that seem odd to us, including their decision to treat income and consumption linearly in their SWB equations, in contrast to the more standard practice of assuming (or at least allowing for) diminishing marginal returns (see Layard, Mayraz, and Nickell, 2008, for a more extensive analysis of the functional form of the SWB function).

Brown, Taylor, and Wheatley Price (2005) instead employ disaggregated measures of savings and debt. Comparing the predictive ability of these, as well as a variety of other models, we argue that the disaggregated approach better reflects the influence of wealth variables (and that it does so even when thinking about one's wealth holdings in a disaggregated way is not in individuals' emotional interests). Indeed, focusing on net wealth in the BHPS would have incorrectly led us to conclude that wealth was unimportant to SWB.

In further contrast to the traditional consumption-based model, we find evidence that status, self-esteem, and anticipatory emotions may all play a role in explaining the impact of pecuniary variables on SWB — with the latter in particular lending support to recent behavioural models that attempt to incorporate the effects of anticipation into current utility (e.g. Caplin and Leahy, 2001; Köszegi, 2005).

The existence of these indirect channels also suggests ways in which contextual factors may affect the importance of income and wealth to SWB, and our analysis provides a number of examples of this point. In Russia, we find that earned income has a greater effect on SWB than does income that is unearned (which we attribute to the former having a greater impact on status and self-esteem); while in Britain we find evidence that wealth that is more closely linked to consumption in the near future has a greater effect on

SWB.

The potential importance of context is further highlighted by both the significant difference in the effect of income between Russia (where it is nearly as important as as consumption) and Britain (where its impact is close to an order of magnitude lower than consumption); and by the fact that the importance of income relative to consumption rises substantially in Britain in response to the 2000 stock market crash. These changes suggest that the relative importance of income is greater in situations of high volatility — again lending support to the idea that expectations of and security about the future have important effects on SWB.

2.3 Background

2.3.1 The Russian Context

The post-transition period in Russia has been a time of significant economic and social upheaval (more comprehensive accounts of which can be found elsewhere: see e.g. Klugman and Braithwaite, 1998; Graham, Eggers, and Sukhtankar, 2004). Both the radical economic reforms that began in 1992,[14] and the financial crisis that struck the country in 1998 have had drastic

[14] See e.g. European Bank for Reconstruction and Development (2004) for a detailed account of these reforms

effects on the Russian economy over the period, which, as Table 2.1 shows, has been marked by large swings in national income, and high and variable unemployment and inflation.

Table 2.1: Economic Indicators for Russia, 1994-2004

Year	GDP/capita growth	CPI Inflation	Unemployment
1994	-12.46	307.63	8.10
1995	-4.02	197.47	9.70
1996	-3.34	47.74	9.90
1997	1.70	14.77	11.80
1998	-5.04	27.67	13.30
1999	6.83	85.74	13.40
2000	10.00	20.78	9.80
2001	5.35	21.46	8.90
2002	5.18	15.79	8.60
2003	7.82	13.68	7.90
2004	7.76	10.88	7.90

All figures in (annual) %. Source: World Development Indicators (November 2007)

While clearly unfortunate from the perspective of those concerned, as Senik (2004) has noted, the extent of these changes in the Russian economy does have one small silver lining: providing analysts with high degree of real and plausibly exogenous variation in economic outcomes. Provided we can adequately control for the high levels of inflation, the Russian context seems to afford us "a natural experiment characterized by an unusually high variance in absolute and relative incomes," (Senik, 2002, p. 2100) with which to assess the impact of consumption, income, and wealth on SWB.

2.3.2 Data

We employ Rounds 5-13 from Phase II of the Russian Longitudinal Monitoring Survey (RLMS), a longitudinal data set tracking between 3700 and 4800 Russian households each year (with the exception of 1997) over the period 1994-2004.[15] While not designed as a panel in the strict sense of tracking the same individuals over time, the same dwelling units were sampled in each round. This effectively yields a panel structure, but with significant attrition over the course of the sample. Of those who are present in Round 5, only 45% are still in the sample in Round 13 (and of those, 27% were not observed in at least one of the intervening rounds). We discuss the potential effects this may have on our estimation (and the steps we take to deal with it) in more detail in Section 2.5.3 below.

In each round, the survey asked the following question of each adult household member: "To what extent are you satisfied with your life in general at the present time?" with five possible answers ranging from "not at all satisfied" to "fully satisfied".[16] As can be seen in Figure 2.2, the average response changes significantly over the course of the sample, in a way that

[15] An earlier phase was also conducted, covering the period 1990-1994. However, the quality of the data from this early phase is generally thought to be poor, and the current Phase was begun in 1994, with a different sample of households. Records from this second phase cannot be linked with those from the first.

[16] The present-focus of the life-satisfaction supports our interpretation of the responses as a measure of *current* life satisfaction, and contrasts with questions asked in other contexts that specify alternative time frames for assessment (e.g. "your life until now").

appears broadly to track the country's economic fortunes.[17]

The survey also contains detailed information on household income, individual income, and household expenditure. Ravallion and Lokshin (2001) have discussed elsewhere the high quality of this income and consumption data. The income measures we use as provided. However, for consumption, we focus on a version of the RLMS "Household Expenditure" variable, modified so as to more accurately reflect a true consumption measure. Following Deaton and Zaidi (2002), we construct this "Household Consumption" measure by (a) excluding savings and asset expenditure; (b) excluding spending on consumer durables, and substituting an estimate of their rental value; and (c) running hedonic housing regressions to obtain imputed values for housing rental. Further details regarding the procedures used to generate these variables can be found in Appendix A.2. Figure 2.2 displays the change over time in the mean of (logged) household income and consumption, and illustrates clearly the scale of the economic turmoil in Russia over the early part of the period. Both incomes and consumption decline sharply until 1996 (recall that no survey was conducted in 1997), though the fall in consumption is not as steep, suggesting that households are dis-saving over the period. From 1998, income recovers quite strongly, finishing the period slightly above its 1994 level; however consumption remains almost constant following the

[17]For more on this point, see e.g. Frijters et al. (2006).

crash, presumably as households pay off debt and/or rebuild their wealth stocks.

Figure 2.2: Consumption, Income, and SWB in Russia

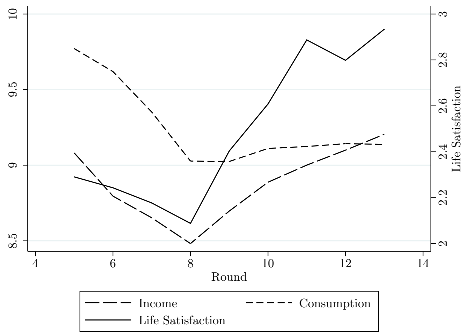

Figure 2.3: Household Debt in Russia

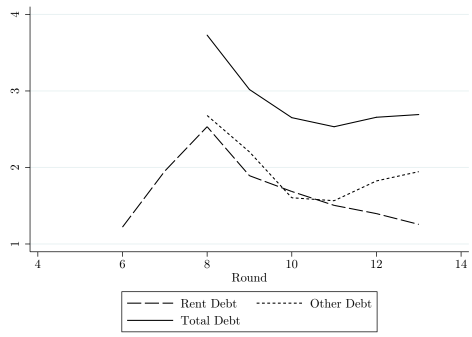

In contrast to the detailed data on consumption and income, the RLMS does not include a comprehensive measure of wealth. In place of such a measure, we use a combination of the following wealth proxies:[18]

(1) the imputed value of the durable stock used to construct the consumption measure;

(2) rent arrears owed by the household (not collected in round 5);

[18] We also experimented with using dummies for home ownership/privatisation, and home value (whether privatised or not). However, in all cases, the coefficients obtained were small and insignificant. This stands in marked contrast to the results reported in Headey and Wooden (2004), who find that the housing stock is one of the most important elements of wealth in explaining well-being. However (despite a potential upward bias in their estimates due to an inability to control for fixed effects) we would not wish to read too much into this. While part of the difference in estimates may reflect real differences in the consequences of home ownership (e.g. in terms of access to credit, or security of tenure) or different tastes for ownership, our estimates are likely to be seriously affected by the survey design, which results in a relative lack of true variation in home values and ownership. The issue of the subjective valuation of home ownership/value seems important in the context of claims that housing privatisation may have either ameliorated or exacerbated inequality in Russia and other transition countries (see e.g. Yemtsov, 2007). However, it is not clear that the RLMS is capable of helping us resolve it.

(3) total household debt (not collected in rounds 5-7);

(4) a savings variable that estimates the time that the household's savings would last at their current level of consumption (collected only in rounds 11 and 12).

While the lack of a comprehensive net wealth measure might initially seem problematic, as we show in Chapter 3, individuals actually appear to base SWB judgments on exactly these sorts of disaggregated factors, rather than their entire net wealth position. These measures therefore seem perfectly adequate for our purpose. Figure 2.3 displays the change in household debt over the sample. As suggested above, the patterns are essentially the inverse of those for consumption and income, with households falling into increasing rent debt until 1996/1997, and then reducing debt thereafter.[19]

As noted above, obtaining *real* income, consumption and wealth values is important given high levels of inflation in Russia across the period. Consequently, all the monetary variables are calculated as monthly values, and deflated to average 1995 roubles, using Goskomstat monthly CPI data.[20]

We do not adjust these measures for household size, preferring instead to include it as a separate regressor. This allows us both to avoid the somewhat

[19] Although access to credit in Russia over the period is generally thought to be limited, much of the rental stock was owned by local governments, and so households may have been more easily able to run up rent debt than otherwise.

[20] As Frijters et al. (2006) note, this seems the lesser evil compared to the alternative of using regional-specific CPI deflators, for which only annual data was available.

vexed issue of the appropriate equivalence scale to use, as well as allowing more flexibility in terms of how we allow the pecuniary variables to enter into the utility function. While the rationale for using per capita consumption (appropriately adjusted for economies of scale) is relatively clear under the standard assumption that the consumption of goods and services is the only argument in the utility function, the case is less straightforward once status and security effects enter the mix.[21]

Additional variables that we use as controls include: age (in months); subjectively evaluated health status (with answers ranging from "very good" to "very bad" on a 5-point scale); marital status (never married, married, divorced, or widow(er)); and employment status.

2.4 Estimating SWB Equations

The empirical strategies we pursue in the remainder of this thesis will vary depending on the precise questions we are attempting to answer. Nonetheless, they will generally share a basic structure, which we describe in general terms in this section.

[21] In fact, while the family size effect is typically significant, and with the expected negative sign, its inclusion changes the point estimates on the pecuniary variables very little.

2.4.1 The SWB Function

The basic equation estimated in most of the literature can be derived from a general utility function of the form $u_{it} = U(u_c(C_{it}), u_z(z_{it}))$.[22] Here, $U(\cdot)$ is a general utility function defined over sub-utility functions in consumption, $u_c(C_{it})$ and a vector of additional variables $u_z(z_{it})$, including demographic and other factors, such as leisure and health, that we expect to influence SWB. This is typically operationalised by assuming: additive separability (both between $u_c(C_{it})$ and $u_z(z_{it})$, and within $u_z(z_{it})$ between the elements of z_{it}); and a log-linear form for $u_c(C_{it})$.[23] Substituting $u_c(x) = \beta_c \ln(x)$ and $u_z(y) = y'\gamma$ into the original equation yields the empirical specification:

$$u_{it} = c_{it}\beta_c + z_{it}'\gamma + \varepsilon_{it} \tag{2.1}$$

where $c_{it} = \ln(C_{it})$.

Generalising this to allow for separate effects of consumption, income, and wealth, we now specify $u_{it} = U(u_c(C_{it}), u_y(Y_{it}), u_a(A_{it}), u_z(z_{it}))$ and again assume additive separability between the sub-utility functions. However, because income and wealth may be zero, and wealth may be negative, we

[22]We omit from this for now any specification of reference effects, which I will address in more detail in Chapter 4.

[23]Layard, Mayraz, and Nickell (2008) present evidence that the elasticity of marginal utility with respect to income is actually slightly higher than that implied by the logarithmic functional form (i.e. 1). They nonetheless suggest that logarithm seems a reasonable approximation.

have to modify the sub-utility functions slightly. For the income sub-utility function, we follow Tinbergen (1991) in specifying $u_y(x) = \beta_y \ln(x+1)$ (which sets $u_y(0) = 0$).[24] To allow negative values for wealth, we simply reflect this below zero, such that $u_a(0) = 0$ and $u_a(-x) = -u_a(x)$:[25]

$$u_a(x) = \begin{cases} \beta_a \ln(x+1) & \text{if } x \geq 0 \\ -\beta_a \ln(1-x) & \text{if } x < 0 \end{cases}$$

Substituting these into (2.5) then yields:

$$u_{it} = m'_{it}\beta + z'_{it}\gamma + \varepsilon_{it} \tag{2.2}$$

where m_{it} is now a vector, $(c_{it}, y_{it}, a_{it})'$ of the logs of monetary variables defined as $c_{it} = \ln(C_{it})$, $y_{it} = \ln(Y_{it}+1)$, and:

$$a_{it} = \begin{cases} \ln(A_{it}+1) & \text{if } A_{it} \geq 0 \\ -\ln(1-A_{it}) & \text{if } A_{it} < 0 \end{cases}$$

and β is the stacked vector $(\beta_c, \beta_y, \beta_a)'$.

[24] Technically, this means that the function no longer invariant to the scale of Y_{it}. However, provided the distribution of Y_{it} begins sufficiently above 1, this has little practical impact.

[25] The assumption that positive and negative wealth have equal effects can be relaxed (and indeed will be in Chapter 3) by allowing the β_as to differ.

2.4.2 The Choice of Estimator

Two major econometric issues face researchers attempting to estimate an equation like (2.2). The first is that responses to SWB survey questions are typically categorical, rather than continuous. The second is the fairly consistent finding that the lion's share of individual SWB is explained by (relatively) fixed personality traits, rather than observable life events (see e.g. Lykken and Tellegen, 1996; Argyle, 1999; Diener and Lucas, 1999; Lyubomirsky, Sheldon, and Schkade, 2005). To the extent that these (typically unobserved) traits are correlated with the explanatory variables in which we are interested, taking account of unobserved individual heterogeneity is likely to be important.

Until recently, however, there has been no accepted method of dealing with fixed effects in the context of an ordered response model. Researchers had thus been left with a choice between:

(1) using standard ordered response models (i.e. ordered probit or ordered logit) without controlling for fixed effects;

(2) treating the categorical responses as continuous, and using standard linear panel data estimators (typically fixed effects or Mundlak (1987) decompositions, as the Hausman tests tend to systematically reject the random effects assumptions); or

(3) collapsing the categorical responses into a single binary variable (indicating whether the response is greater or less than some cutoff point, λ) and using Chamberlain's (1980) conditional (fixed effects) logit estimator (see e.g. Hamermesh, 2001; Winkelmann and Winkelmann, 1998).

While the latter has the advantage of both preserving the categorical nature of the responses, and controlling for fixed effects, it suffers from the major drawback that it loses a very large proportion of data, which, as Ferrer-i-Carbonell and Frijters (2004) note, runs the risk that "measurement errors may well become a large source of residual variation."

To address this problem, Ferrer-i-Carbonell and Frijters (2004) develop a fixed-effects ordered logit estimator and compare its performance to the standard OLS, ordered response, and linear panel data estimators, as well as an alternative ordinal fixed effects estimator proposed by Das and Van Soest (1999). Their paper sets out the derivation of the estimator in detail; however, the approach is essentially to use a variant of the standard (binary) conditional logit estimator proposed by Chamberlain (1980), but with *individual-specific* cut-off points, λ_i, selected to minimise the information lost by collapsing the full scale of ordinal responses into only two categories. Their main findings are that, while failure to control for fixed effects seriously biases estimates, treating survey responses as either continuous or categorical appears to produce very similar results.

Because they yield estimates that have an obvious interpretation in terms of answers to the life-satisfaction question, we will focus on the results obtained from standard linear panel-data regressions, which treat individuals' categorical responses as a continuous measure of u_{it}:

$$u_{it} = m'_{it}\beta + z'_{it}\theta + \eta_i + \varepsilon_{it} \tag{2.3}$$

where, as above, m_{it} is a vector of the logs of one or more monetary variables, z_{it} is now a vector of time-varying control variables, and η_i is an unobserved time-invariant fixed effect that is allowed to correlate with the observed m_{it} and z_{it}. In line with Ferrer-i-Carbonell and Frijters' (2004) findings, and most previous work, Hausman tests conclusively reject the random effects assumptions, and so fixed effects estimates are reported throughout.

As an additional robustness check, we also employ a variant of the Ferrer-i-Carbonell and Frijters estimator. However, because calculating the optimal individual-specific λ_i cut-offs is difficult, we use an approximation (suggested by Ada Ferrer-i-Carbonell) that seems to perform well in practice. Taking λ_i as the mean of individual i's reported life satisfaction across the sample, $\frac{1}{T_i}\sum_{t\in T_i} u_{it}$, and setting $u^*_{it} = \mathbb{1}[u_{it} > \lambda_i]$ we then estimate the latent variable form of (2.3) using a standard conditional logit with u^*_{it} as the dependent variable. The results are typically very similar (in terms of the relative effects of the independent variables), and are not reported here.

2.5 The Effect of Consumption, Income & Wealth on SWB

This section begins our empirical investigation of the effects of consumption, income, and wealth on SWB in Russia, by attempting to answer two questions: (1) whether consumption, income, and wealth have separate, independent effects on SWB; and (2) whether the common practice of estimating SWB equations using income as the independent variable of interest underestimates the importance of money to SWB. The evidence we present suggests that both queries should be answered "yes". Not only do income-based estimates appear to understate the importance money to SWB (consistent with suggestions in Headey and Wooden, 2004; Headey, Muffels, and Wooden, 2008), but consumption, income, and wealth each appear to exert independent effects on current SWB. Section 5.2.2 sets out the basic results. Section 2.5.2 and Section 2.5.3 then consider a number of issues in their interpretation.

2.5.1 Basic Results

As outlined in Section 2.4, our basic empirical strategy involves running regressions of the form set out in (2.3). As our dependent variable throughout, we use individual answers to the 5-point current life-satisfaction question described in Section 2.3.2 above, while the key variables of interest, m_{it}, are one or more the logged household income, consumption, and wealth measures also detailed there.

In all cases, we include in the regression a standard set of controls, z_{it}, consisting of age, age squared, health, marital status (dummies for each of married, divorced and widow(er), with never married as the omitted category) an employment dummy, and family size.[26] We also include a full complement of round dummies to allow for common shocks.

To address the key questions set out above, we compare the results from four alternative models:

(1) the standard income regression employed in most of the SWB literature in which $m_{it} = y_{it}$;

(2) the consumption-based model of standard economic theory in which $m_{it} = c_{it}$;

[26]Cf. Ferrer-i-Carbonell and Frijters (2004), who include almost the same variables, but collapse the marital status dummies into a single "Steady Partner" indicator, and use only the number of children, rather than the family size. (The latter seems preferable in Russia, where extended family units are likely to be more common than in Germany.)

(3) an alternative model that allows both consumption and income to affect current SWB and sets $m_{it} = (y_{it}, c_{it})'$; and

(4) the most general model in which $m_{it} = (y_{it}, c_{it}, a'_{it})'$ and consumption, income and wealth are all allowed to affect SWB independently (because we have four distinct wealth measures, a_{it} may itself be a vector).

The results of regressions based on models (1)-(3) are shown in the corresponding columns of Table 2.2.

Table 2.2: The Effect of Income and Consumption on SWB in Russia

	(1)		(2)		(3)	
	β	IME	β	IME	β	IME
(Log) Household Income	0.116***				0.093***	
	(19.640)				(15.499)	
(Log) Household Consumption			0.152***		0.113***	
			(17.347)		(12.562)	
Age	-0.037***		-0.039***		-0.039***	
	(-3.750)		(-3.859)		(-3.963)	
Age^2	0.000***		0.000***		0.000***	
	(5.917)		(6.468)		(6.497)	
Health	0.158***	3.9	0.160***	2.9	0.158***	2.2
	(19.342)		(19.663)		(19.459)	
Married	0.104***	2.0	0.111***	1.8	0.112***	1.5
	(3.357)		(3.579)		(3.616)	
Divorced	-0.127***	.40	-0.121***	.52	-0.114***	.68
	(-3.491)		(-3.327)		(-3.128)	
Widow(er)	-0.151***	.33	-0.151***	.43	-0.136***	.61
	(-3.867)		(-3.877)		(-3.511)	
Employed	0.136***	3.2	0.150***	2.7	0.134***	1.9
	(9.997)		(11.013)		(9.850)	
Family Size	-0.022***	.83	-0.022***	.87	-0.034***	.85
	(-3.850)		(-3.820)		(-5.979)	
Constant	1.740***		1.292***		0.902**	
	(4.967)		(3.599)		(2.540)	
Observations	78409		78409		78409	
Groups	19235		19235		19235	
R^2	0.144		0.150		0.158	

* $p<0.10$, ** $p<0.05$, *** $p<0.01$. t-statistics (robust to individual-level clustering) in parentheses. Round dummies included but not reported.

The effects of the non-pecuniary variables in the basic income regression in column (1) (and indeed, also in the other columns) are consistent with other findings in the literature. In order to put these into appropriate perspective, we can express the effects in terms of an Income Multiplier Equivalent (IME) — that is, the amount by which income would have to be multiplied to generate an equivalent increase (or decrease) in well-being.[27] Subjective health status has a strong and significant effect on SWB, with a one step increase being equal to a 3.9-fold increase in income. Being married or cohabiting has an IME of around 2, while those who become divorced (or widowed) are even less happy than those who never married, becoming as unhappy as those with approximately 40% (33%) of the latter's income.[28] Becoming unemployed has an effect of similar magnitude (1/3.2=.31). One divergence from previous results is the effect age. While the coefficients on age and age squared have the usual sign, they suggest a minimum at around 70, well above the common finding of around 40-50. A possible explanation for this is that the elderly in Russia may have been left most insecure by the

[27]The formula for the IME is given by:

$$\text{IME}_x = \exp\left(\frac{\beta_x}{\beta_y}\right)$$

where β_y is the coefficient on log income, and β_x is the coefficient on the variable in question. The change in u due to a change in y is given by $\Delta u_y = \beta_y \ln(\text{IME})$, implying that $\text{IME} = \exp(\Delta u_y/\beta_y)$. Setting $\Delta u_y = \Delta u_x = \beta_x$ yields the above expression. Cf. the more common method of quoting absolute monetary equivalents set out in e.g. Clark and Oswald (2002).

[28] These are calculated assuming that the loss of a spouse induces a reduction in family size of 1. Thus, e.g. $\text{IME}_{divorce} = \exp\{(\beta_{divorce} - \beta_{familysize})/\beta_y\}$

volatile economic situation and least able to adjust psychologically to the fall of communism. Nonetheless, as Ferrer-i-Carbonell and Frijters (2004) note, we probably should not place too much weight on these estimates. Age is likely to co-vary with a number of non-observables, and estimated age effects therefore tend to be quite sensitive to the controls included.[29]

Turning to the other columns of Table 2.2, we see in column (2) that including the consumption measure alone attracts a larger coefficient than income, consistent with the idea that the latter may be merely proxying for the former. However, contrary to what we would expect if income were merely a proxy for consumption, when we include the two together in Column (3), income retains a sizeable and significant effect, suggesting that it may in fact have an independent influence on well-being.[30]

The potential weaknesses of the consumption-only model (2) are further highlighted in Table 2.3, which shows the results for model (4), where we introduce our wealth proxies. Because the rounds for which the proxies are available vary, we report a number of alternative baseline regressions for comparison, each using the same sample as the corresponding wealth regression. In all cases, the wealth variables also appear to affect SWB

[29]For further investigations into the age profile of SWB, see e.g. Easterlin (2004) and Blanchflower and Oswald (2008).

[30]We also experimented with using the RLMS household expenditure measure instead of consumption. The coefficients were qualitatively similar to those for consumption, but, as we would expect, slightly smaller and less precisely measured.

Table 2.3: The Effect of Income, Consumption, and Wealth on SWB in Russia

	Durables		RentDebt		OverallDebt		Savings		Combined	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
(Log) Household Income	0.093***	0.094***	0.097***	0.094***	0.091***	0.086***	0.110***	0.107***	0.097***	0.094***
	(15.499)	(15.575)	(14.770)	(14.327)	(11.071)	(10.564)	(5.570)	(5.431)	(14.770)	(14.400)
(Log) Household Consumption	0.113***	0.098***	0.124***	0.123***	0.109***	0.112***	-0.009	-0.013	0.124***	0.110***
	(12.562)	(10.571)	(12.777)	(12.715)	(9.330)	(9.595)	(-0.322)	(-0.453)	(12.777)	(11.103)
(Log) Durable Stock		0.018***								0.016***
		(5.785)								(4.834)
(Log) Rent Arrears				-0.014***						-0.014***
				(-8.490)						(-8.471)
(Log) Debt						-0.014***				
						(-9.303)				
Savings								0.082***		
								(3.514)		
Standard Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	78409	78409	68500	68500	53445	53445	18993	18993	68500	68500
Groups	19235	19235	17869	17869	15610	15610	11381	11381	17869	17869
R^2	0.158	0.159	0.167	0.170	0.166	0.171	0.093	0.095	0.167	0.171

* $p<0.10$, ** $p<0.05$, *** $p<0.01$.
t-statistics (robust to clustering within panelids) in parentheses.

independently of income and consumption, with each entering significantly with the expected sign (positive for wealth, negative for debt), both when included alone (Columns 1 to 8), and with the durable stock and rent debt included jointly (Column 10).[31]

For the most part, the inclusion of the wealth proxies induces only small decreases in the coefficients on income and consumption, and roughly in accordance with what we would expect. The effect of income falls slightly when we include debt measures, presumably because higher (lower) income may be used to pay off (lead to increases in) debt; and the effect of consumption falls somewhat when we include the durable stock, as the two are correlated by construction. The only apparent complication is that the consumption coefficient becomes slightly and insignificantly negative in columns (7) and (8), where the availability of the subjective savings variable reduces us to only two rounds of data (11 and 12). As this occurs even in column (7), before the introduction of the savings variable itself, it cannot be due to omitted variable bias (and in any event savings and consumption should be negatively correlated). Instead, it seems that the reduction is due in large part to the loss of real variation in consumption in these two rounds of the sample. Figure 2.4 displays the mean absolute changes in log consumption

[31] We exclude overall debt and savings in the final "full" specification, due to the limited number of rounds in which they are observed.

Figure 2.4: Annual Changes in Logged Consumption and Income by Round

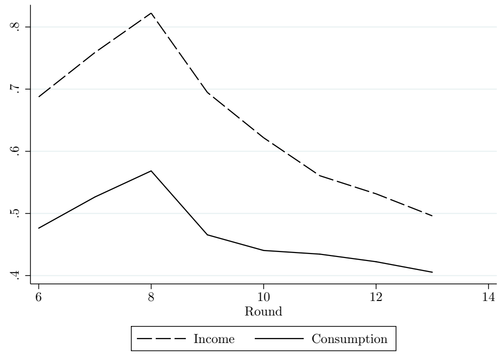

and income over the course of the sample. While there appears to be little consumption smoothing over the early rounds of the RLMS (a proposition defended in more detail by Stillman, 2001), as the effects of the financial crisis have receded, consumption has become much more stable, increasing the chance that measurement error will come to dominate our fixed effects regressions.

Taking these estimates at face value suggests two broad conclusions: first, that consumption, income and wealth *all* affect SWB; and second, that the standard income regression of model (1) understates the importance of money to SWB. The precise degree of this understatement (in any given period, or over a lifetime) will depend on how period income is allocated between consumption now and in the future. Nonetheless, to get a rough idea of the total effect, we assume that all income is consumed in the period it is received, and that the total effect is therefore $y_{it}(\hat{\beta}_y + \hat{\beta}_c) = 0.206y_{it}$, which is around 80% larger than the measured effect of income from column 1, $0.116y_{it}$. Relative to other variables, the correction leads to decreases in the IMEs for health from 3.9 to around 2.2, and for marriage from 2.2 to 1.5, while the effect of unemployment is now equivalent to a 48% (=1-1/1.9) fall in income (compared to 69% previously).

2.5.2 Are the Effects of Income and Wealth Really Independent?

Nonetheless, there are a number of potential reasons why one might doubt these conclusions, and in particular the claim that the effects of income and wealth are independent of the effects of consumption. This section deals with three of these: measurement error in consumption, reverse causation, and omitted variables bias. It argues that in fact, none of them challenges our basic point.

Measurement Error in Consumption

One potential reason to doubt the conclusion that income, consumption and wealth have independent effects on well-being is that the other variables may still merely be picking up measurement error in consumption. However, three factors suggest that this is unlikely to be driving the result. First, this can only be a real concern in the case of income, and the durable stock measure (the latter of which is correlated with our consumption variable by construction, and typically causes the coefficient on consumption to drop by around 0.1). Once we control for fixed effects, the other wealth variables should be (and indeed are) negatively correlated with consumption, and their coefficients are unlikely to be due to picking up unmeasured consumption.

Second, in relation to income, the size of its coefficients and the accuracy with which they are measured suggests that they are unlikely to merely be picking up the effects of consumption. In column (3) of Table 2.2, not only are the coefficients on income and consumption of similar magnitude, but the standard errors on income are *smaller*. This is exactly the opposite of what we should see if both were merely acting as noisy measures of true consumption, as we would expect our consumption measure to better track true consumption than our income measure does.

Finally, if measured income and wealth were truly orthogonal to SWB except for their effect on consumption, then they would be valid instruments for consumption. We could then estimate the true effect of consumption on SWB by instrumenting consumption with measures of income and wealth. To examine the plausibility of the orthogonality assumption, we pursue exactly this instrumenting strategy. The results of using a variety of instruments for consumption are shown in Table 2.4 (we omit the control coefficients and focus only on the consumption coefficient). Column 1 displays the baseline, uninstrumented regression for comparison.[32] Column (2) shows the results of instrumenting solely with income. Here, because the model is just-identified, we cannot test the validity of the orthogonality assumption formally. However, the consumption coefficient of 0.588 seems very large. A rough idea of the plausibility of this coefficient can be obtained by examining a basic linear regression using only income and consumption, under the classical errors in variables assumptions. In this case, it is well-known that $\sigma_u^2/(\sigma_{c^*}^2 + \sigma_u^2) = (\beta_{c^*}/\hat{\beta}_c^{FE} - 1)/(\beta_{c^*}/\hat{\beta}_c^{FE})$, which would imply, rather implausibly, that measurement error constitutes somewhere in the vicinity 75% of the variation in the measured deviations from mean consumption. This suggests that income is in fact correlated with SWB independently of consumption.

Table 2.4: Are Income and/or Wealth Valid Instruments for Consumption?

	(1)	(2)	(3)	(4)	(5)	(6)
Consumption	0.152***	0.588***	0.453***	0.648***	0.480***	0.333***
	(19.496)	(20.810)	(22.038)	(19.936)	(20.522)	(10.651)
Instruments	None	Income	Income Durables	Income Debt	Income Durables Debt	Durables Debt
Standard Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	79250	73769	73769	64234	64234	64964
Groups	19322	14595	14595	13603	13603	13687
Sargan-Hansen Test			51.19	63.66	125.24	79.22
p-value			0.000	0.000	0.000	0.000

* p<0.10, ** p<0.05, *** p<0.01.
t-statistics in parentheses.

Including the wealth proxies as additional instruments allows us to test

[32] This corresponds to column 2 of Table 2.2.

the validity of the instruments more formally. The Sargan-Hansen test of over-identifying restrictions (Hansen, 1982; Sargan, 1958) conclusively rejects the null hypothesis that the instruments are jointly valid for all combinations of at least two instruments. This result — that there is no combination of two instruments such that both can be excluded from the SWB regression — implies that, *at most* one of the additional variables could be actually be orthogonal to SWB. To the extent that we are convinced by the arguments above that neither debt measures, nor income are simply picking up measurement error in consumption, this suggests that the orthogonal variable, if there is one at all, is most likely to be the durable stock. Even if this were the case, however, because it is only one measure of wealth, its exclusion would still not alter the point that both wealth and income appear to have separate effects on well-being.

Of course, the fact that our measures of income, debt, and (perhaps) durables are not orthogonal to SWB (conditional on consumption and other covariates) does not *necessarily* imply that they have a causal effect on SWB. The failure of orthogonality could arise instead because of correlation in the measurement errors, which would also serve to invalidate the variables as instruments. (The common CPI measures used to deflate both could plausibly contribute to such correlation.[33]) Nonetheless, these results still serve to increase our confidence that income, consumption and wealth enter separately into the utility function.[34]

Can Happiness Buy Money?

Another potential reason to doubt the claim that income matters to SWB is the possibility of reverse causation. There is some evidence in the literature that happiness may in fact "buy" money (see e.g. Diener et al., 2002; Lyubomirsky, King, and Diener, 2005), and the difficulty of finding valid instruments to test this directly necessarily leaves us with a degree of uncertainty. However, even if reverse causation is a factor here, it seems unlikely to significantly affect our conclusions.

---

[33] To see this, note the construction of our estimates as:

$$y_{it} = \ln(Y^*_{it}\Upsilon_{yit}P_{it}) \tag{2.4}$$
$$c_{it} = \ln(C^*_{it}\Upsilon_{cit}P_{it}) \tag{2.5}$$

where *s indicate true values, Υs multiplicative errors, and P_{it} a common price index. These imply:

$$y_{it} = y^*_{it} + (\upsilon_{yit} + p_{it}) \tag{2.6}$$
$$c_{it} = c^*_{it} + (\upsilon_{cit} + p_{it}) \tag{2.7}$$

(where lowercase letters indicate logs) and will thus generate correlation in the errors due to the common p_{it} term, even if $\text{Corr}(\upsilon_{yit}, \upsilon_{cit}) = 0$.

[34] Of course, even were we not willing to accept that the effects of income and wealth are independent of consumption, we are still led to the conclusion that the basic income regression understates the impact of money on well-being — probably by an even larger amount than was suggested previously. The precise amount depends on which instrumented value we use, and the extent to which any increase in the instrumented value is driven by correlated measurement error. Instrumenting only with income suggests an increase of around 400%, but correlation in the errors would suggest this is too high.

As we noted at the outset, the large swings in the Russian economy as a whole over the sample period suggest that the majority of the variation in income is likely to be exogenous, even in the absence of instruments. In any event, much of the argument for reverse causation is in fact based on the idea that relatively stable personality traits that correlate with happiness tend also to breed success (Diener et al., 2002). Such personality traits or dispositions are absorbed within our fixed effects, and will not affect our estimates.

There are also plausible channels via which temporary happiness may lead to later success: as Lyubomirsky, King, and Diener (2005) note "[p]ositive emotions produce the tendency to approach rather than to avoid and to prepare the individual to seek out and undertake new goals". However, the lags with which such effects result in increased income should minimise the impact on our estimates.[35,36] All things considered, while there may be some

[35] If there were persistence in SWB (even after conditioning on observed covariates and unobserved fixed effects) then this would undermine this argument to some extent, because changes in both current income and current SWB could be driven by past changes in SWB. We attempted to test for such persistence using difference and system GMM estimation, but were unable to find valid instruments. The OLS estimate of the term on the lagged dependent variable (which is likely to be upward biased) was 0.24, while the FE estimate (likely to be downward biased) is -0.07. In the absence of valid GMM instruments, we can say little more than that the true degree of persistence is likely to lie somewhere between these (see e.g. Nickell, 1981) — but this range at least permits the possibility that there is little or none.

[36] In any event, the evidence that lagged SWB is correlated with future success supports at least two alternative explanations. One, not tested here, is that it is not current happiness that breeds future success, but current success that breeds both, and that, once we have controlled for current income, health and so forth, happiness may no longer predict future outcomes. The other is that individuals' current satisfaction is affected by their expectations about the future (a contention for which we present evidence later).

residual reverse causation occurring, our assessment is that its effect is likely to be small.[37]

Omitted Variables

A final possibility is that that we may merely be picking up the effect of unobservables correlated with our pecuniary variables. This seems most likely in relation to income, which could be correlated with job characteristics that independently influence well-being (such as having a higher status, or more satisfying job). We test for this possibility by including in our regressions a number of more direct measures of job characteristics: (1) whether the individual holds a managerial position; (2) a measure of job income risk (this is the product of two subjective risk assessments: the likelihood of losing one's job and the likelihood of being unable to find another job once lost, both on a 5 point scale); and (3) a subjective measure of professional skill on a 9-point scale (available only in rounds 7-9), which we expect to correlate with status.

The results of regressions including these variables (run only on the em-

[37]Other studies (e.g. Senik, 2004b), have used lagged values of regressors to test for reverse causation. However, it is difficult to know what to draw from such evidence, particularly given the large swings in consumption and income over the course of the RLMS sample. While lagged income and consumption may indeed be orthogonal to happiness-induced success in period t, at the same time, they will both be orthogonal to entirely valid parts of of the remaining consumption and income measures, and may also capture other effects that we are not interested in (such as adaptation).

ployed sub-sample of the population) are reported in for model (3) in Table 2.5.[38] While these job characteristics all contribute positively and significantly to well-being, their effect seems to be almost entirely orthogonal to income. The included job characteristics affect the point estimates on income (or consumption) barely at all, increasing further our confidence that the estimated income effect is causative, rather than merely correlative.

Table 2.5: Can Job Characteristics Account for the Income Effect? (Income and Consumption Specifications)

	(1)	(2)	(3)	(4)
(Log) Household Income	0.117***	0.116***	0.108***	0.103***
	(12.708)	(12.664)	(5.725)	(5.548)
(Log) Household Consumption	0.128***	0.123***	0.122***	0.115***
	(9.285)	(8.915)	(3.836)	(3.604)
Manager		0.065***		0.068
		(3.162)		(1.464)
Job Risk		-0.012***		-0.011***
		(-12.372)		(-5.132)
Professionalism				0.021**
				(2.379)
Observations	37352	37352	10570	10570
Groups	11258	11258	5931	5931
R^2	0.161	0.188	0.087	0.106

* p<0.10, ** p<0.05, *** p<0.01. t-statistics (robust to clustering within panelids) in parentheses.

[38] Because our primary concern here is the robustness of the income coefficient, we do not report the results of including job characteristics with wealth variables, but these results are basically identical.

Conclusion

Taken together, the discussion in this section suggests that, despite initial concerns to the contrary, the apparent effects of income and wealth found in Table 2.2 and Table 2.3 actually do reflect *independent* effects on SWB. We have presented evidence that neither the income nor the debt effects appear to be due to these variables picking up measurement error in consumption, and that the correlation with income is due neither to reverse causation nor to omitted variables. To the extent that there is doubt about any of the coefficients, it seems to apply primarily to the durable stock measure, which is correlated with consumption by construction. However, even if this correlation were spurious, because the estimates for the other wealth measures seem sound, it would not affect our underlying point that consumption, income, and wealth appear to exert independent effects on SWB.

2.5.3 Attrition

As we noted in Section 2.3.2, the RLMS suffers from significant attrition. While there is little reason to *expect* such attrition to bias our results in any particular way, it nonetheless seems wise to confirm this empirically, given the sheer scale of movements into and out of the sample. For each of our main regressions we therefore conduct tests designed to detect bias arising from non-random attrition. These tests suggest that attrition is indeed non-random, but that the effects of re-entry and replenishment are sufficient to maintain the representativeness of the sample. We also confirm our earlier results using inverse-probability weighted (IPW) estimates that attempt to correct for attrition more directly.

Detecting Attrition Bias

Verbeek and Nijman (1992) and Wooldridge (1995) suggest basic variable addition tests to detect non-random attrition. These consist in adding to our earlier regressions one of two alternative variables: either an indicator $1 - s_{(i,t-1)}$ that an individual failed to respond in the previous round, or an indicator $1 - s_{(i,t+1)}$ that they failed to respond in the subsequent round.[39] Intuitively, the coefficients on these variables test whether those who attrite or enter are different in any relevant respect from the remaining sample, and thus whether the representativeness of the sample is compromised by such movements into and out of it. The tests we apply are extensions of this strategy, modified to account for re-entry and sample replenishment in the RLMS.

As Verbeek and Nijman (1992) note, a sufficient condition for the consis-

[39] Verbeek and Nijman (1992) use response indicators rather than *non*-response indicators, but the principle is identical.

tency of the fixed effects estimator on the unbalanced panel is:

$$\mathrm{E}(\tilde{\varepsilon}_{it}|s_i) = 0, \qquad t = 1, 2, \ldots, T. \tag{2.8}$$

(where $s_i = (s_{i1}, s_{i2}, \ldots, s_{iT})$ is a vector of selection indicators, and $\tilde{\varepsilon}_{it} = \varepsilon_{it} - \frac{1}{T_i}\sum_{\tau=1}^{T} s_{i\tau}\varepsilon_{i\tau}$ are the transformed fixed effects errors). The variable addition tests just described amount to tests of the restrictions that $\mathrm{E}(\tilde{\varepsilon}_{it}|s_{i,t-1}) = 0$ and $\mathrm{E}(\tilde{\varepsilon}_{it}|s_{i,t+1}) = 0$, which in this context would imply $\mathrm{E}(\tilde{\varepsilon}_{it}|s_{i,t-1}, s_{i,t+1}) = 0$.[40]

The problem with these tests is that, while insignificant coefficients on each of the added variables would indeed suggest a lack of attrition bias, the converse is not true: even if exit and entry are *individually* correlated with the errors, their effects may counteract each other. Indeed, there are good reasons to think that this may be the case here: on the one hand, the later re-entry of those who attrite only temporarily will help maintain the sample's representativeness; one the other, the sample design means that attriters are replaced by those who move into their previous residence (and who may therefore be to similar to them in relevant respects).[41] To allow for this possibility,

[40]Ideally, of course, we would test alternative potential violations of (2.8). However attempting to include more leads and lags quickly diminishes our sample to the point where such tests have little power. Our tests therefore rely on the implicit assumption that $\mathrm{E}(\tilde{\varepsilon}_{it}|s_i) = \mathrm{E}(\tilde{\varepsilon}_{it}|s_{i,t-1}, s_{i,t+1})$.

[41]Consistent with these conjectures, Heeringa's (1997) analysis of the effect of attrition on the early sample, found that it did not significantly affect representativeness; but to our knowledge no equivalent analysis has been performed for the later rounds.

we therefore weaken the requirement that the coefficients on each added variable be insignificant, and instead focus on whether their *sum* is significantly different from zero (which tests the restriction that $\mathrm{E}(\tilde{\varepsilon}_{it}|s_{i,t-1}, s_{i,t+1}) = 0$ directly).

We conduct four alternative versions of this test for each of our main regressions above, each with a slightly different definition of the non-response indicator. We first distinguish between two alternative selection indicators: s_{it}, which requires only that individual i be in the *RLMS* sample in period t, and r_{it} which requires that the individual be in our *regression* sample in period t (i.e. both in the RLMS sample, and having valid responses to all the questions we use). We then define for each selection indicator, q_{it} (for $q = r, s$) two alternative sets of non-selection variables.

(1) $1 - q_{i,t-1}$ and $1 - q_{i,t+1}$, which indicate non-selection in either the previous or the subsequent period (respectively); and

(2) $\mathbb{1}(q_{i,t-n} = 0, \forall\, n > 0)$ and $\mathbb{1}(q_{i,t+n} = 0, \forall\, n > 0)$ which indicate (respectively) whether the individual is observed in for the first time in t (i.e. they are a new entrant) or for the last time in t (i.e. they afterwards attrite for good).

The results of these tests are displayed below in Table 2.6 for the consumption-income model (3), and Table 2.7 for the wealth models (4). In all cases, the coefficients on the future non-selection indicators are significant and negative, suggesting that attrition is indeed non-random. On the other hand, the coefficients on almost all of the past non-selection indicators are positive and significant, confirming our suspicion that entry at least partially counteracts exit. Conducting F-tests of the constraint that the sum of the coefficients is equal to zero shows that in no case can we reject the hypothesis that the effects of attrition and entry cancel.[42]

Table 2.6: Variable Addition Tests for Attrition Bias (Income & Consumption Specifications)

	Selection ($q = s$)		Non-Response ($q = r$)	
	$n = 1$	$n > 0$	$n = 1$	$n > 0$
$\mathbb{1}(q_{i,t+n} = 0)$	-0.043***	-0.052***	-0.027*	-0.043***
	(-2.734)	(-2.936)	(-1.958)	(-2.588)
$\mathbb{1}(q_{i,t-n} = 0)$	0.035**	0.023	0.039***	0.039**
	(2.280)	(1.306)	(3.060)	(2.563)
Observations	60725	60725	60725	60725
Groups	16912	16912	16912	16912
$F(1; 16911)$	0.12	1.23	0.36	0.03
p	0.733	0.268	0.551	0.870

* p<0.10, ** p<0.05, *** p<0.01.
t-statistics (robust to clustering within panelids) in parentheses.
Independent variables include income, consumption and standard controls

Correcting for Attrition Bias

To confirm that attrition is is not a serious concern, we also run a set of inverse-probability-weighted regressions that attempt to correct for any at-

[42]Similar (unreported) tests on the regressions in Table 2.5 suggest that these results too, are robust.

Table 2.7: Variable Addition Tests for Attrition Bias (Wealth Specifications)

	Selection ($q = s$)		Non-Response ($q = r$)	
	$n = 1$	$n > 0$	$n = 1$	$n > 0$
$\mathbb{1}(q_{i,t+n} = 0)$	-0.049***	-0.066***	-0.039***	-0.060***
	(-2.743)	(-3.240)	(-2.590)	(-3.224)
$\mathbb{1}(q_{i,t-n} = 0)$	0.039**	0.019	0.034**	0.029*
	(2.332)	(0.945)	(2.505)	(1.776)
Observations	51348	51348	51348	51348
Groups	15719	15719	15719	15719
$F(1; 15718)$	0.15	2.52	0.05	1.44
p	0.699	0.113	0.824	0.231

* p<0.10, ** p<0.05, *** p<0.01.
t-statistics (robust to clustering within panelids) in parentheses.
Independent variables include income, consumption, durable stock overall debt, and standard controls

trition bias.[43] The method is described in detail in Appendix A.3, but essentially involves estimating the predicted probability of a given individual remaining in the sample at time t, and weighting each observation by the inverse of this prediction. (The intuition behind this is that if an individual has, say, a 0.1 probability of remaining in the sample at t, then in the absence of attrition, there would be 10 (i.e. 0.1^{-1}) relevantly similar individuals still in the sample, and we can correct for their absence by increasing the weight of the remaining individual to count for the missing nine.) As the Appendix notes, the key to this strategy is finding sufficiently strong predictor variables for attrition, and the procedure suffers from the drawback that we must restrict ourselves to using only an "artificial monotone" sample

[43] The IPW approach has been employed in a number of studies analysing the RLMS (see e.g. Senik, 2004; Mu, 2003)

(i.e. the sample that does not allow reentry if an individual is ever absent).[44] Here, this results in losing around 2/3 of our data in most cases (and even more for some of the wealth regressions.) Nonetheless, we report the results of IPW estimation in Table 2.8 and Table 2.9 below.

Table 2.8: The Effect of Consumption and Income on SWB in Russia (Inverse Probability-Weighted Estimates)

	(1)	(2)	(3)
(Log) Household Income	0.130**		0.106**
	(12.124)		(9.643)
(Log) Household Consumption		0.163**	0.119**
		(9.542)	(6.770)
Age	0.001	0.001	0.001
	(0.736)	(0.960)	(0.923)
Age^2	0.000	0.000	0.000
	(1.480)	(1.690)	(1.642)
Health	0.140**	0.144**	0.141**
	(7.471)	(7.720)	(7.591)
Married	0.041	0.035	0.041
	(0.470)	(0.399)	(0.461)
Divorced	-0.103	-0.102	-0.091
	(-1.204)	(-1.193)	(-1.061)
Widow(er)	-0.229*	-0.241**	-0.217*
	(-2.549)	(-2.706)	(-2.402)
Employed	0.163**	0.175**	0.158**
	(6.110)	(6.615)	(5.957)
Family Size	-0.028**	-0.027*	-0.041**
	(-2.713)	(-2.575)	(-3.861)
Observations	33501	33501	33501
Groups	7865	7865	7865
R^2	0.514	0.513	0.516

* $p<0.05$, ** $p<0.01$. t-statistics (robust to clustering within households) in parentheses. Round dummies included but not reported.

[44] As discussed in e.g. Robins and Rotnitzky (1995) and Robins, Rotnitzky, and Zhao (1995), the case of non-monotone attrition, where attriters are allowed to (re)enter in subsequent periods is rather more involved, requiring additional assumptions, and a much more complicated estimation procedure.

Table 2.9: The Effect of Consumption, Income and Wealth on SWB in Russia (Inverse Probability-Weighted Regressions)

	Durables		RentDebt		OverallDebt		Savings		Combined	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
(Log) Household Income	0.106***	0.106***	0.118***	0.116***	0.119***	0.115***	0.162***	0.164***	0.118***	0.116***
	(10.488)	(10.504)	(10.088)	(9.894)	(6.784)	(6.524)	(3.369)	(3.383)	(10.088)	(9.926)
(Log) Household Consumption	0.119***	0.112***	0.121***	0.121***	0.110***	0.112***	0.113	0.116	0.121***	0.115***
	(6.966)	(6.373)	(6.110)	(6.082)	(3.951)	(4.032)	(1.279)	(1.315)	(6.110)	(5.720)
(Log) Durable Stock		0.009*								0.007
		(1.748)								(1.104)
(Log) Rent Arrears				-0.010***						-0.010***
				(-3.195)						(-3.187)
(Log) Debt						-0.011***				
						(-3.305)				
Savings								-0.060		
								(-0.810)		
Standard Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	33501	33501	24848	24848	14926	14926	3853	3853	24848	24848
Groups	7865	7865	5678	5678	3503	3503	2084	2084	5678	5678
R^2	0.517	0.517	0.532	0.532	0.559	0.560	0.778	0.778	0.532	0.532

* p<0.10, ** p<0.05, *** p<0.01.
t-statistics (robust to clustering within panelids) in parentheses.

As can be seen, the substantive results shown in Table 2.8 and Table 2.9 are essentially unchanged compared to those of Table 2.2 and Table 2.3. While some of the control variables become insignificant, the only substantive changes in the coefficients of interest are that the coefficient on the subjective savings variable in column (8) and the coefficient on the durable stock in column (10) of Table 2.9 become insignificant. The former is hardly surprising given the loss of almost 80% of the original rounds 12-13 sample, and in the latter case the coefficients on income, consumption, and rent arrears remain significant, despite losing around 2/3 of the original sample. In all other cases, the coefficients on income, consumption, and each of the individual wealth variables retain their significance, and approximate magnitudes compared with the main regressions presented earlier.

2.6 Why Does Money Matter?

The results of the previous section suggested that neither the standard theoretical account of how money matters to SWB nor the standard empirical approach to estimating its importance may be adequate. In contrast to both the consumption focus of the former and the income focus of the latter, consumption, income, and wealth each appear to exert independent effects on SWB in Russia. Moreover, in line with earlier results in Headey and Wooden

(2004) and Headey, Muffels, and Wooden (2008), we found that focusing on income alone appears to understate the importance of money to SWB (here by around 80%).

Having argued that consumption, income and wealth have separate effects on SWB, we now go on to address two further issues arising from this. Section 2.6.1 employs a system of equations approach to identify some of the channels though which such effects may operate. It presents evidence that status and anticipation concerns appear to account for part of the effect of our pecuniary variables on SWB. Section 2.6.2 adopts an alternative, interaction-based approach to argue both that status concerns are important in accounting for the effect of income on SWB; and that this leads to systematic variation in the importance of household income, depending on the degree to which that income is *earned* (either by the particular individual concerned or collectively by members of the household). Altogether, this evidence serves both to strengthen the existing case that income and wealth matter to SWB independently of consumption, and to illustrate the potential importance of accounting for the channels by which they do so.

2.6.1 Why Does Money Matter?

As noted above, this section uses a system of equations approach to examine the extent to which status concerns and anticipatory emotions can explain the effects of income and wealth on SWB identified in Section 2.5. It proceeds in three stages, first describing our empirical strategy, then setting out the basic results, and finally addressing issues of interpretation and robustness.

Empirical Strategy

If we assume that the m_{it} variables (and indeed, the x_{it} variables too) may affect SWB either directly, or through the channels displayed in Figure 2.1, we can respecify equation (2.3) as a system of equations of the form:

$$u_{it} = e_{it}\gamma_e + s_{it}\gamma_s + f_{it}\gamma_f + m'_{it}\lambda_u + z'_{it}\mu_u + \alpha_{ui} + \omega_{uit} \quad (2.9)$$

$$e_{it} = m'_{it}\lambda_e + z'_{it}\mu_e + \alpha_{ei} + \omega_{eit} \quad (2.10)$$

$$s_{it} = m'_{it}\lambda_s + z'_{it}\mu_s + \alpha_{si} + \omega_{sit} \quad (2.11)$$

$$f_{it} = m'_{it}\lambda_f + z'_{it}\mu_f + \alpha_{fi} + \omega_{fit} \quad (2.12)$$

where e_{it} is self-esteem, s_{it} is status, and f_{it} is anticipation. If we have explicit measures of the e_{it}, s_{it} and f_{it} variables, we can estimate these equations consistently with fixed effects under the assumption that the errors from equations (2.10)-(2.12) are independent of those from equation (2.9): $\text{Cov}(\omega_{jit}, \omega_{uit}) = 0$, $\forall\ j \in \{e, s, f\}$. (This assumption does not seem unreasonable given that the fixed effects should remove most of the correlation between the errors, though it remains possible that there are time-varying unobservables that could affect both intermediate outcomes and SWB.[45])

The RLMS does in fact contain a number of variables that could serve as useful proxies for status and anticipation variables, s_{it} and f_{it}. Unfortunately, they do not allow us to distinguish convincingly between status and self-esteem effects, so we are forced to drop the e_{it} equation. Any effects that operate through self-esteem will be absorbed either in the status channel coefficients (λ_s and γ_s) or the "direct" effect (λ_u). The variables we use as status proxies are individuals' subjective perceptions (on a 9-point scale) of:

(1) s_{1it}: the level of respect in which they are held;

(2) s_{2it}: level of rights/power they have; and

(3) s_{3it}: their economic rank.

For anticipation proxies, we employ:

(1) f_{1it}: answers to the question "Do you think that in the next 12 months you and your family will live better than today or worse?" (ranging

[45] The alternative would be to seek identification via exclusion restrictions, but that would require us to find variables that could affect the intermediate outcomes (self-esteem, status, and security) without having any direct effect on SWB. It is difficult to imagine what such variables might be.

from "much better" to "much worse" on a 5-point scale, with a mid-point of "no change"); and

(2) f_{2it} answers to the question "How concerned are you about the possibility that you might not be able to provide yourself with the bare essentials in the next 12 months?" (ranging from "very concerned" to "not at all concerned" on a 5-point scale).

These are not, by any means, perfect measures of the subjective outcomes we are interested in. The status proxies are primarily externally rather than internally focused (the first appears to be aimed at capturing the degree of respect accorded one by others, while all three focus on perceived rank) and are unlikely to effectively capture any sort of self-esteem effect. Our security measures, too, are somewhat imperfect as assessments of individuals' future prospects: f_{1it} focuses on *changes* in individuals' prospects rather than an absolute prediction of their future situation; and while f_{2it} does provide a prediction about an absolute outcome, its focus is primarily negative, and it will capture anxiety about the future far more effectively than it does hope. Nonetheless, we make do with what we have.

Employing these variables gives us the following operationalisation of the system (2.9)-(2.12):

$$u_{it} = s_{1it}\gamma_{s1} + s_{2it}\gamma_{s2} + s_{3it}\gamma_{s3} + f_{1it}\gamma_{f1} + f_{2it}\gamma_{f2}$$

$$+ m'_{it}\lambda_u + z'_{it}\mu_u + \alpha_{ui} + \omega_{uit} \tag{2.13}$$

$$s_{1it} = m'_{it}\lambda_{s1} + z'_{it}\mu_{s1} + \alpha_{s1i} + \omega_{s1it} \tag{2.14}$$

$$s_{2it} = m'_{it}\lambda_{s2} + z'_{it}\mu_{s2} + \alpha_{s2i} + \omega_{s2it} \tag{2.15}$$

$$s_{3it} = m'_{it}\lambda_{s3} + z'_{it}\mu_{s3} + \alpha_{s3i} + \omega_{s3it} \tag{2.16}$$

$$f_{1it} = m'_{it}\lambda_{f1} + z'_{it}\mu_{f1} + \alpha_{f1i} + \omega_{f1it} \tag{2.17}$$

$$f_{2it} = m'_{it}\lambda_{f2} + z'_{it}\mu_{f2} + \alpha_{f2i} + \omega_{f2it} \tag{2.18}$$

Results

Table 2.10 reports the results of estimating equations (2.13)-(2.18) with household consumption and income as the m_{it}s and standard controls as the z_{it}s. Table 2.11 displays the same regressions, but now with the durable stock and rent debt included in the m_{it}s. The results suggest three broad conclusions. First, as we can see from column (1) of both tables, the mediator variables for status and anticipation have significant effects of the predicted sign: that is, they each appear to be important in explaining SWB.[46] Second, as we can see in the remaining columns (2)-(6), the pecuniary variables are almost all strongly associated with both status and anticipation. The only exception to this pattern is that a household's durable stock does not appear

[46] The coefficients are not directly comparable due to the different units in which they are measured.

to be related to either individuals' concern for the future or their expected future prospects, which was probably to be expected given the nature focus of the available anticipation proxies on either changes in prospects or negative prospects.

Finally, the mediator variables also appear to be associated with a number of our controls: in particular, employment correlates both with individuals' sense of status, and their level of concern about the future; health is strongly associated with all of the mediator variables; and marital status appears to be associated with significant changes in concern for the future. Perhaps most surprising of these associations is the fact that marriage is correlated with *increased* concern for the future in Table 2.10 (recall that high concern is indicated by low values, so the negative coefficients imply increased concern), which perhaps reflects the fact that one has more to be concerned about than just oneself. There also appear to be fairly clear age patterns to the variables, with expectations about the future decreasing, and status increasing, with age (both at a diminishing rate). We would not wish to place too much weight on these results, particularly in instances such as health, where causation could conceivably run either way. Nonetheless, the patterns appear plausible, and contribute to the overall impression that status and anticipation are indeed important.

To get a better sense of exactly how important each channel is in account-

Table 2.10: The Mediating Effect of Security and Status on SWB in Russia

	LS	Concern	Hope	Respect	Power	Econ Rank
(Log) Household Income	0.054***	0.089***	0.075***	0.056***	0.068***	0.109***
	(8.560)	(11.897)	(12.091)	(4.906)	(6.744)	(12.090)
(Log) Household Consumption	0.064***	0.078***	0.057***	0.092***	0.078***	0.173***
	(6.902)	(7.035)	(6.422)	(5.408)	(5.141)	(12.629)
Concern About Future	0.155***					
	(32.906)					
Hope For Future	0.162***					
	(28.856)					
Respect	0.036***					
	(12.361)					
Power	0.031***					
	(8.491)					
Economic Rank	0.112***					
	(26.425)					
Employed	0.091***	0.102***	0.009	0.136***	0.109***	0.109***
	(6.500)	(6.007)	(0.692)	(5.378)	(4.766)	(5.251)
Age	-0.031***	-0.051***	-0.038***	0.047**	0.023	-0.017
	(-3.219)	(-4.049)	(-3.686)	(2.565)	(1.366)	(-1.154)
Age^2	0.000***	0.001***	0.000***	-0.000*	-0.000**	0.000**
	(2.740)	(9.667)	(9.282)	(-1.702)	(-2.047)	(2.141)
Health	0.105***	0.091***	0.099***	0.088***	0.161***	0.149***
	(12.260)	(8.947)	(12.336)	(5.628)	(11.538)	(11.941)
Married	0.126***	-0.103***	0.045	-0.051	-0.033	0.028
	(3.956)	(-2.986)	(1.598)	(-1.005)	(-0.687)	(0.659)
Divorced	-0.051	-0.129***	-0.039	-0.056	-0.014	-0.095*
	(-1.378)	(-3.073)	(-1.122)	(-0.852)	(-0.234)	(-1.780)
Widow(er)	-0.067*	-0.226***	-0.050	0.001	-0.040	-0.130**
	(-1.677)	(-4.936)	(-1.299)	(0.008)	(-0.640)	(-2.244)
Family Size	-0.021***	-0.048***	-0.013**	-0.012	0.022**	-0.018**
	(-3.586)	(-7.012)	(-2.294)	(-1.111)	(2.255)	(-2.008)
Observations	62897	63130	63130	63130	63130	63130
Groups	17846	17874	17874	17874	17874	17874
R^2	0.277	0.092	0.184	0.000	0.056	0.157

* p<0.10, ** p<0.05, *** p<0.01. t-statistics (robust to clustering within panelids) in parentheses.

ing for the importance of different variables, we can eliminate the s_{it} and f_{it} variables from the system (2.13)-(2.18), and decompose the total effect of each pecuniary variable as:

$$\beta = \lambda_{s1}\gamma_{s1} + \lambda_{s2}\gamma_{s2} + \lambda_{s3}\gamma_{s3} + \lambda_{f1}\gamma_{f1} + \lambda_{f2}\gamma_{f2} + \lambda_u \qquad (2.19)$$

Intuitively, (2.19) says that the total SWB effect β of m_{it} can be decomposed into the sum of its status effects (the $\lambda_s\gamma_s$ terms) anticipation effects (the $\lambda_f\gamma_f$ terms), and a direct (or, perhaps more accurately "residual") effect, λ_u. Each of the indirect effects is in turn composed of two parts: the effect of m_{it}

Table 2.11: The Mediating Effect of Security and Status on SWB in Russia

	LS	Concern	Hope	Respect	Power	Econ Rank
(Log) Household Income	0.055***	0.086***	0.070***	0.049***	0.060***	0.102***
	(8.148)	(10.505)	(10.793)	(4.037)	(5.450)	(10.404)
(Log) Household Consumption	0.061***	0.085***	0.057***	0.071***	0.069***	0.165***
	(5.851)	(6.783)	(5.847)	(3.768)	(4.015)	(10.600)
(Log) Durable Stock	0.015***	0.001	-0.001	0.018***	0.016***	0.023***
	(3.857)	(0.308)	(-0.290)	(2.703)	(2.693)	(3.843)
(Log) Rent Arrears	-0.009***	-0.011***	-0.008***	-0.007**	-0.014***	-0.016***
	(-5.213)	(-5.159)	(-4.644)	(-2.138)	(-4.761)	(-6.031)
Concern About Future	0.149***					
	(29.330)					
Hope For Future	0.165***					
	(26.483)					
Respect	0.036***					
	(11.327)					
Power	0.031***					
	(7.767)					
Economic Rank	0.110***					
	(23.521)					
Employed	0.096***	0.112***	0.015	0.127***	0.099***	0.130***
	(6.223)	(6.038)	(1.027)	(4.781)	(4.042)	(5.872)
Age	-0.043***	-0.056***	-0.040***	0.058***	0.038**	-0.015
	(-3.922)	(-3.405)	(-3.490)	(2.921)	(2.064)	(-0.944)
Age^2	0.000***	0.001***	0.000***	-0.000***	-0.000*	0.000***
	(3.973)	(10.246)	(7.459)	(-3.240)	(-1.778)	(3.413)
Health	0.112***	0.077***	0.086***	0.089***	0.149***	0.144***
	(11.962)	(7.000)	(10.216)	(5.386)	(9.843)	(10.683)
Married	0.137***	-0.048	0.073**	0.030	-0.037	0.067
	(3.879)	(-1.249)	(2.400)	(0.558)	(-0.708)	(1.477)
Divorced	-0.043	-0.074	-0.011	0.000	0.001	-0.055
	(-1.009)	(-1.549)	(-0.283)	(0.002)	(0.009)	(-0.952)
Widow(er)	-0.097**	-0.133***	0.001	0.085	-0.047	-0.087
	(-2.183)	(-2.600)	(0.014)	(1.054)	(-0.665)	(-1.387)
Family Size	-0.025***	-0.043***	-0.013**	-0.016	0.024**	-0.020**
	(-3.837)	(-5.766)	(-2.048)	(-1.465)	(2.366)	(-2.182)
Observations	54616	54784	54784	54784	54784	54784
Groups	16447	16466	16466	16466	16466	16466
R^2	0.270	0.072	0.213	0.000	0.000	0.092

* p<0.10, ** p<0.05, *** p<0.01. t-statistics (robust to clustering within panelids) in parentheses.

on status or anticipation (via the λ_s or λ_f) and the effect of that channel on SWB (via γ_s or γ_f).

We use this equation to calculate both the total effect of each m_{it} variable via all channels and to decompose this into the shares accounted for by each of the measured mediator variables. The decomposition corresponding to the results in Table 2.11 is displayed graphically in Figure 2.5, and appears to confirm that status and anticipation are important in accounting for the impact of our pecuniary variables.[47] Collectively, these channels account for

[47] The decomposition of income and consumption corresponding to the results in Ta-

around 40% of the effect of income, consumption, and rent debt — although they account for only around 20% of the effect of the durable stock. In the case of consumption and rent debt, the contributions of status and anticipation are approximately equal. In contrast, anticipation is more important to the effect of income, but entirely unimportant to the effect of durables.

In all cases, there remain large residual effects, which are unfortunately difficult to interpret. These could reflect either actual direct effects of these variables on SWB, effects that operate via unmeasured channels (such as self-esteem), or simply measurement error in the mediator variables (which, as we noted earlier, do not precisely capture the notions we are concerned with). For the latter reasons, the displayed residual effect are more likely to reflect upper bounds on the "direct" effect than unbiased estimates of it.

Even setting aside the large residual effects, there remain a number of concerns with the interpretation of these results, which the following sections seek to address. The first centres on the issue of causation, the second on the potential for survey artifacts to drive the results, and the third on the possibility of attrition affecting the results. We address these in turn below.

ble 2.10 are virtually identical, and are therefore omitted for brevity.

Figure 2.5: Decompositions of Consumption, Income and Wealth into Status and Anticipation Effects

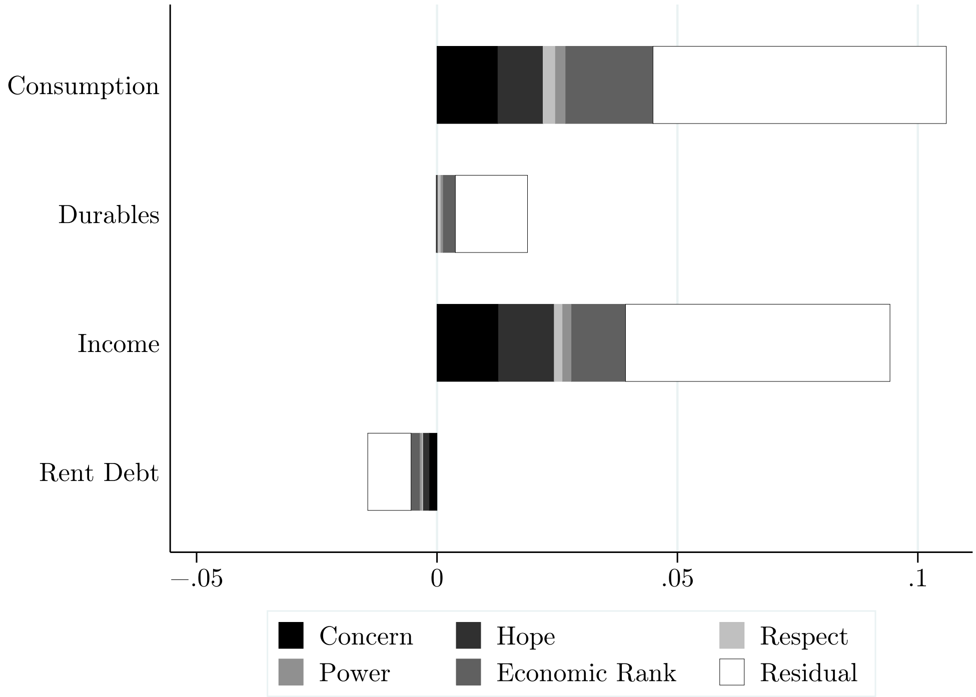

Bars illustrate both the overall effect of each pecuniary variable on SWB, and the contribution of each status or anticipation channel to that effect.

Causation

For the most part, it seems reasonable to assume that the direction of causation here runs from from changes in economic outcomes to changes in the subjective variables, rather than the other way around. While independent changes in subjective variables (particularly changes for the worse) are likely to elicit *subsequent* behavioural responses, these should not greatly affect our results, because subjective assessments should quickly readjust to any resulting change in circumstances. (For example, if my expectations about my future prospects change for the worse, I may subsequently begin to work more, and earn more; but the very fact that I am now earning more should then feed back into my future expectations, and mask the original causal channel.)

There are, however, two specific exceptions to this. First, it seems that consumption is more likely to *reflect* individuals' expectations about the future than to *affect* them. If we anticipate a comfortable future, we are likely to consume more now than if we anticipate a worse one. This suggests that the suggested influence of consumption on concern and hope for the future may well be spurious. On the other hand, we are likely to see the opposite effect with wealth: if our prospects deteriorate (improve), we may respond by saving more (less). The resulting negative association between our sense of security and our wealth could potentially offset any presumed *positive* causal link from wealth to security.

Survey Artefacts

Another potential problem in the interpretation of these results is that the association between the status and anticipation measures that we use could be inflated by priming effects in the survey. With the exception of the concern about the future variable, f_{2it}, all of the other questions appear immediately prior to the general life-satisfaction question in the RLMS. We test the idea that this may have biased our results using an alternative life-satisfaction question asked in rounds 12 and 13. This question asked respondents "How satisfied or unsatisfied are you with: your life in general at present?", with possible answers "absolutely satisfied", "mostly satisfied", "yes and no", "not very satisfied", and "absolutely not satisfied".

Table 2.12 reports the results of regressions of the same form as column (1) of Table 2.10. The first column reproduces the relevant part of of Table 2.10 for comparison. The second and third columns show the results of the same regression run on the sub-sample for which the alternative life-satisfaction answers u^*_{it} are available, first with the original dependent variable u_{it}, and then with the alternative u^*_{it}. Comparing columns (1) and (2), we see that simply restricting the sample to the two final rounds results in a significant fall in the coefficients and significance levels of the respect, economic rank, and concern variables — though all, with the exception of the respect variable, remain significant. Comparing columns (2) and (3), we find that using the alternative measure of life-satisfaction results in a further drop in most coefficients and their significance. Though again, all except the respect variable remain significant, the coefficients on respect, economic rank, and future prospects fall by close to 50%. On the other hand, the power variable, and the concern about the future variable (the latter of which should not suffer from priming bias) remain closer to their original coefficients. These results suggest that the importance of the subjective variables may well be overstated in our main regressions. Nonetheless, the fact that their coefficients remain sizeable and significant even with the alternative measure u^*_{it} provides evidence that they are still important, even if less so than we might otherwise have thought.

To account for the potential biasing effects of reverse causation and priming, we recalculate the decompositions displayed in Figure 2.5 above with two adjustments. First, we restrict the effect of consumption to operate only via status and residual effects. Second, we scale down the effect of each indirect channel to SWB, using the ratio of its coefficients in columns (2) and (3) of Table 2.12. Figure 2.6 displays the adjusted decompositions.

In contrast to the previous results, indirect channels now account for only $\frac{1}{3}$ of the effect of Rent Debt (split evenly between status and anticipation), $\frac{1}{4}$ of the effect of income (of which the majority is still due to anticipation), $\frac{1}{5}$ of the effect of consumption (now attributable entirely to status), and $\frac{1}{6}$ of the effect of durables (also attributable entirely to status). However, more so than previously, we can now be confident that the estimated effects operating via each channel are indeed lower bounds — and consequently that such channels are indeed important, even if the majority of the effect of our pecuniary variables is direct or unexplained.

One final point of note is that, if our concerns about reverse causation

Table 2.12: Testing for the presence of survey artefacts

	(1)	(2)	(3)
Concern About Future	0.155***	0.077***	0.048***
	(33.065)	(4.426)	(2.613)
Hope For Future	0.162***	0.188***	0.109***
	(28.809)	(6.710)	(3.682)
Respect	0.036***	0.007	0.000
	(12.533)	(0.456)	(0.023)
Power	0.031***	0.038**	0.038**
	(8.554)	(2.327)	(2.350)
Economic Rank	0.112***	0.106***	0.053***
	(26.458)	(5.571)	(2.983)
Observations	62922	8066	8066
Groups	17847	5491	5491
R^2	0.282	0.165	0.038

* $p<0.10$, ** $p<0.05$, *** $p<0.01$. t-statistics (robust to clustering within panelids) in parentheses.

from anticipation to current consumption are true, then this will require us to revise downwards our estimates of the total effect of consumption on SWB. Indeed, as Figure 2.6 shows, its total effect is now slightly less than that of income, and the total effect of the two together is now only around $0.160y_{it}$.[48] Comparing this with our original income-only estimate of $0.116y_{it}$ suggests that it may in fact only understate the effect of money by around 40%.

[48]This figure comes from recalculating equation (2.19) with the $\lambda_f\gamma_y$ components omitted for the consumption β.

Figure 2.6: Adjusted Decompositions of Consumption, Income and Wealth into Status and Anticipation Effects

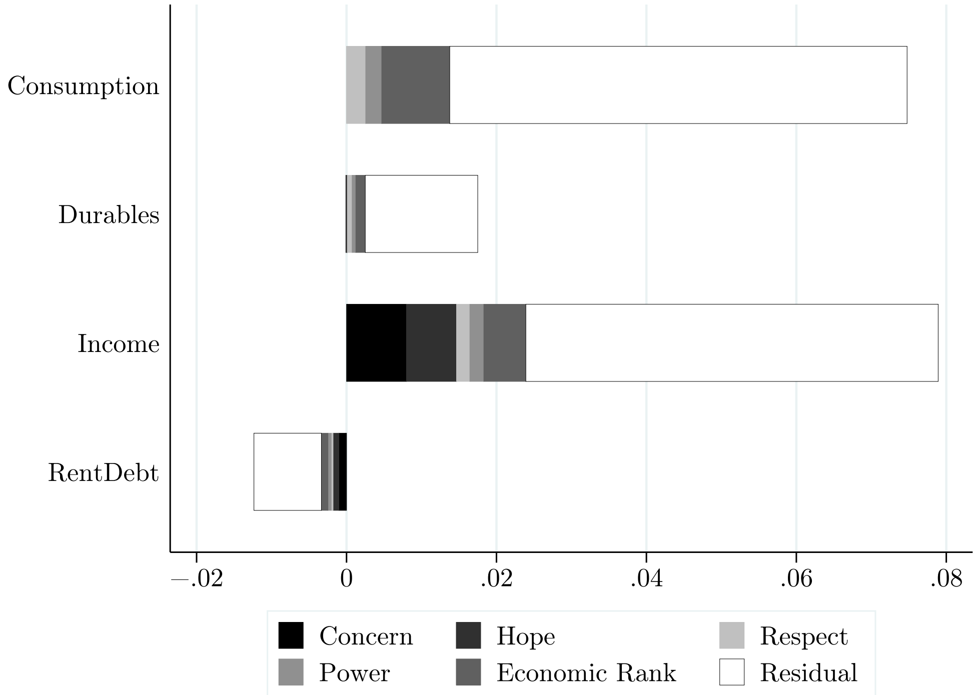

Attrition

As with the analysis of Section 2.5, attrition remains a potential source of bias in the analysis. However, after conducting both the tests and corrections outlined in Section 2.5.3, the results appear to remain robust. (We do not report the results here to conserve space.)

Conclusion

Altogether, the evidence of this section suggests that both status and anticipation concerns are important in accounting for the effects of income, consumption and wealth on SWB. As already noted, this strengthens the case that the effects of income and wealth are indeed independent of consumption. It also opens up the possibility (which we explore in the next section) that there could be systematic variation in the importance of different types of consumption, income and wealth in different circumstances. Nonetheless, there are some weaknesses in the analysis. Our initial estimates of the importance of the status and anticipation channels may well be biased upwards by survey artefacts, and there remains a significant unexplained "residual" to be accounted for.

2.6.2 Variation in the Effects of Consumption, Income, & Wealth

Section 2.6.1 suggested that money affects well-being partly through indirect channels. This opens up the possibility that different types of income, consumption or wealth may have different hedonic returns. As a specific example of this, this section considers the hypothesis that earned income may be more important than unearned income as a source of status or self-esteem, and therefore of global SWB. Again, we proceed in three stages, first setting out our empirical strategy, moving on to the basic results, and then discussing issues of interpretation and robustness.

Empirical Strategy

To operationalise the hypothesis that earned income will have a greater effect on status/self-esteem than unearned income, we begin by re-expressing the system of equations in equation (2.9)-equation (2.12) in the following reduced form (combining the self-esteem and status channels for simplicity):

$$u_{it} = m_{it}'\beta + y_{it}(\lambda_s\gamma_s + \lambda_f\gamma_f + \lambda_u) + z_{it}'\theta + \varepsilon_{it} \tag{2.20}$$

where the m_{it} is a vector of pecuniary variables other than income, which has been separated out as y_{it}, and everything else is defined as before. Our hypothesis can then be operationalised as assuming that the impact of income on status/self-esteem — that is, the coefficient λ_s— depends on the proportion of income that is earned:[49]

$$\lambda_s = \lambda_{s0} + \frac{Y_{it}^{earned}}{Y_{it}}\phi \tag{2.21}$$

Substituting this into (2.20) and simplifying then yields:

$$u_{it} = m_{it}'\beta + \frac{Y_{it}^{earned}}{Y_{it}}y_{it}\phi\gamma_s + z_{it}'\theta + \eta_i + \varepsilon_{it} \tag{2.22}$$

[49] If the equation for s_{it} were linear in Y_{it} rather than its log, y_{it}, then we could simply enter Y_{it}^{earned} and $Y_{it}^{unearned}$ separately. However, the non-linearity makes such a strategy infeasible, and leads us to pursue the interaction approach set out here.

where we again include y_{it} in the m_{it} vector, and $\beta_y = \lambda_{s0}$ from equation (2.21). If the coefficient on the the interaction term, $\frac{Y_{it}^{earned}}{Y_{it}} y_{it}$, is positive and significant, this suggests both that status/self-esteem is relevant to the importance of income (i.e. $\gamma_s > 0$) and that $\frac{Y_{it}^{earned}}{Y_{it}}$ does indeed mediate this effect (i.e. $\phi > 0$) and cause systematic variation in the importance of y_{it}.[50]

Results

We estimate this equation with three alternative measures of Y_{it}^{earned}: *individual* income; non-transfer income; and a vector comprised of both measures. In all cases, we also include the proportion of earned income, $\frac{Y_{it}^{earned}}{Y_{it}}$, hours worked, and hours worked squared in the regressions as additional controls.[51] The results of regressions using income and consumption as the relevant m variables are shown in Table 2.13.[52]

In the first two specifications, both the proportion of earned income alone, and the interaction terms attract significant positive coefficients, with the individual income interaction being slightly smaller than the non-transfer income interaction. Both effects are maintained when the two variables are

[50] Given the results of the system estimation in the previous section, we set aside the possibility that $\phi\gamma_s$ could be significantly positive due to both components being negative.

[51] Including hours worked ensures that the effect of earning more income is not attenuated for individuals who have to work longer hours to do so. As it happens, the coefficients on the hours worked terms are always tiny and insignificant, and so are not reported.

[52] We have also run the regressions including the durable stock and rent arrears variables: the results were virtually identical.

included together (in column (6)). Consistent with our hypothesis, this suggests that holding total household income constant, people are happier the greater the proportion that is earned or which they are in some sense responsible for.

Is the Earned Income Effect Driven By Status?

Of course, one might plausibly think that this relationship is driven by something other than our hypothesised link between earned income and status. Two other possibilities are:

(1) that higher *individual* income gives the earner greater bargaining power within the household, and enables them to gain a greater share of household consumption; and

(2) that *individual* income correlates with desirable job characteristics

Neither (1) nor (2) can explain the higher coefficient on *earned*, as opposed to *individual* income. Moreover, we have already dismissed the idea that job characteristics could account for the importance of income above. Just to be sure, we conduct two (unreported) tests of these hypotheses. The first hypothesis suggests that the same amount of household consumption should yield more SWB the higher is individual income. We can thus test it by including an interaction between the proportion of individual income and the log of household consumption. However, the interaction term in

Table 2.13: The Effect of Income Source on the Importance of Income to Life-Satisfaction in Russia

	Individual		Earned		Combined	
	(1)	(2)	(3)	(4)	(5)	(6)
(Log) Household Income	0.096***	0.103***	0.093***	0.063***	0.096***	0.075***
	(15.529)	(16.510)	(15.374)	(8.171)	(15.529)	(9.621)
(Log) Household Consumption	0.113***	0.110***	0.113***	0.107***	0.113***	0.105***
	(12.377)	(12.011)	(12.569)	(11.906)	(12.377)	(11.402)
Proportion Individual Income		0.034***				0.025***
		(5.718)				(5.742)
Proportion Individual x (Log) HH Income		0.005***				0.004***
		(5.792)				(5.832)
Proportion Earned Income				0.109***		0.109***
				(5.394)		(4.881)
Proporion Earned x (Log) HH Income				0.037***		0.034***
				(6.237)		(5.731)
Employed	0.117***	0.115***	0.120***	0.113***	0.117***	0.108***
	(7.017)	(6.872)	(7.338)	(6.906)	(7.017)	(6.465)
Observations	76383	76383	78409	78409	76383	76383
Groups	19023	19023	19235	19235	19023	19023
R^2	0.150	0.149	0.158	0.157	0.150	0.149

* $p<0.10$, ** $p<0.05$, *** $p<0.01$. t-statistics (robust to clustering within panelids) in parentheses. Standard controls, hours worked, and squared hours worked included but not reported.

such an equation is insignificant, suggesting that this is not the source of the effect. The second explanation suggests that the effect should disappear if we control for job characteristics. However, including measures of managerial status, job risk, and professional expertise does not alter the coefficients.

In contrast, we can test more directly the initial hypothesis that earned income is important specifically because of its effect on *status*, and when we do so, the results are encouraging. This time, instead of substituting equation (2.21) into (2.20), we substitute it directly into equations (2.14)-(2.16) and test the status effects directly. That is, we run regressions identical to the previous two, but now with the three status measures (power, respect, and economic rank) as the dependent variables. The (unreported) results support our original hypothesis: in all cases but one, the interaction terms on individual and earned income are significant and have the expected positive sign, suggesting that the interaction effect is indeed mediated via status. (The only insignificant interaction occurs for household earned income in the respect equation, and does not greatly affect the thrust of our conclusion.)

Attrition

Again, the possibility arises that the results here may be affected by attrition. In line with the analysis presented above, we conduct tests to detect attrition bias, and run IPW regressions to correct for the possibility. In neither case do the (unreported) results affect the conclusions above.

2.7 Conclusion

The analysis of this chapter suggests three broad conclusions. First, our results confirm previous suggestions that estimates of the effect of money on well-being that focus solely on income are likely to underestimate its effect substantially (Headey, Muffels, and Wooden, 2008; Headey and Wooden, 2004). In contrast to previous work, however, our fixed effects estimates avoid any potential for biases arising out of omitted personality traits (which, as Ferrer-i-Carbonell and Frijters, 2004, show, are likely to correlate with both economic outcomes and SWB).

Second, in contrast to the standard interpretation of this result (see e.g. Clark, Frijters, and Shields, 2008), which views it as consistent with income acting simply as a noisy measure of consumption, we go beyond previous work in arguing that income, consumption and wealth each affect well-being independently, and that the standard model is incomplete in this regard. The presence of income and wealth in the current utility function may have implications for both welfare analysis and individual decision-making.

For example, the evidence that more than just consumption matters to SWB may help to inform ongoing policy debates about the most appropriate means of measuring poverty, inequality or material well-being more broadly. In particular, it would imply that the frequent assertion that consumption is a theoretically preferable measure of well-being (see e.g. Meyer and Sullivan, 2003, 2007) be overstated, and suggest something of a middle ground in the contest between income and consumption-based measures.

The presence of income as a direct argument in the utility function also creates a preference for income-smoothing (for essentially the same reasons as the traditional consumption smoothing result), meaning that income volatility is likely to be more harmful to well-being than the standard consumption-based theory would predict. On the other hand, the presence of wealth in the utility function creates a competing preference for "front-loaded" income streams. The proofs of these results are straightforward, and relegated to Appendix A.1.

The presence of wealth may also have important implications for the optimal behaviour of those with high debt levels, and in particular for their behaviour under risk (see e.g. Karelis, 2007, though he focuses primarily on the effects of convexities at low levels of consumption). While risk-preferences cannot be read directly off the shape of the SWB function,[53] convexities in the SWB function will increase the likelihood of risk-loving behaviour.

[53] A consistent decision-utility function can be any monotonic transform of an SWB function. If this transform is sufficiently concave (convex), then it may result in risk-averse (risk-loving) behaviour, regardless of the curvature of the SWB function itself.

Third, there is evidence that at least part of the effect on life satisfaction of both pecuniary and other variables is mediated via perceptions of status and concerns about the future (although a sizeable residual effect typically remains to be explained). The suggestion that the impact of income is mediated by status or self-esteem suggests that different sources of income may have different hedonic returns, and we present evidence that Russians do in fact gain greater SWB from income for which they are somehow responsible (either as individuals or as a household) than income for which they are not. This may have implications for cost-benefit analysis of alternative welfare policies, such as "workfare" or conditional transfers.

Meanwhile, the evidence that anticipation matters lends support to models such as those of Loewenstein (1987), Caplin and Leahy (2001), and Köszegi (2005), which incorporate future consumption possibilities into the current utility function.[54] More broadly, it may also be thought to lend weight to the old Jevonian perspective (see e.g. Jevons, 1888, 1905) on intertemporal decision-making — which argued that saving and investment decisions were rooted in the effect of anticipatory emotions on *current* utility, rather than (or more plausibly, in addition to) weighing the interests of

[54] One interesting question suggested by this is whether individual-level variation in the strength of the wealth effect on current utility could partly explain the variation in individuals' "propensity to plan" for the future, which has been documented by Ameriks, Caplin, and Leahy (2003). They present evidence that these varying propensities cannot be due simply to variation in discount rates.

the present against the future, as emphasised by Von Böhm-Bawerk (1889), Fisher (1930), and in the standard discounted utility model of Samuelson (1937).[55] To the extent that the future matters to SWB *now* and not merely to SWB when it occurs, this suggests, amongst other things, that the rate at which individuals actually discount future utility may be much higher than the standard model would suggest. Indeed, as Jevons (1888) argued, it could provide a motive for saving even if individuals weighted their future SWB at zero.

The Russian experience over the 10 years from 1994-2004 has provided us with a unique natural experiment to examine the impact of consumption, income, and wealth on SWB — and the results from our analysis here seem fairly clear. Nonetheless, as with any experiment (natural or otherwise) there remains a question about how the results are likely to scale beyond the immediate context. To see whether our conclusions are likely to carry over to more stable environments, and/or different cultural contexts, the next chapter turns to examine British data from 1996-2005 as a point of comparison.

[55] Of course, these perspectives are not mutually exclusive. Frederick, Loewenstein, and O'Donoghue (2002) discuss the contrast between them, and its potential implications, in more detail.

Chapter 3

Britain

The previous chapter used Russian evidence to argue that standard approaches to thinking about the importance of money to SWB are flawed in two distinct ways. First, the standard empirical strategy of using income as the independent variable of interest in SWB regressions understates the importance of money to SWB. Second, the standard theoretical model, which assumes that current consumption is the sole pecuniary argument in the instantaneous utility function is incomplete: income and wealth matter too, in ways that are likely to be important to welfare analysis.

Nonetheless, given the volatility of Russian transition experience, one might doubt the applicability of these results more broadly. Perhaps the trauma of transition has engendered in the Russian populace a widespread anxiety about their economic security, and it is this anxiety that causes

Russians to place so much weight on income and wealth as buffers against future misfortune. Or perhaps Russian culture focuses more than others on income and wealth as markers of status. If so, then it could still be the case that they are relatively unimportant elsewhere, under more "normal" economic circumstances.

On the other hand, these same considerations would even more strongly impugn the results of income-only SWB regressions. If, outside of the rather special Russian case, income is not in fact very important to SWB, except via its effect on consumption, then income regressions may understate the importance of money to SWB to an even greater extent than the Russian evidence would suggest: in the Russian case income measures at least have the virtue of capturing something that individuals appear to care about *independently*, rather than being merely noisy measures of consumption.[1]

To see whether these concerns are well-founded, this chapter analyses the experience of a country that did not face such a volatile economic situation in the recent past; and which might reasonably be presumed to differ culturally from Russia: Britain from 1996 to 2005. We find that the key points from our analysis of the Russian case stand: income regressions continue to understate the importance of money (to an even greater extent than in Russia); and

[1]Moreover, to the extent that Russians were less able than Britons to smooth their consumption over the period, the value of income as a consumption measure would also be greater (see e.g. Stillman, 2001, for the argument that Russians smoothed consumption relatively little over much of the early RLMS sample period).

the data still reject the standard consumption model (SCM) of how money matters to SWB. Nonetheless, we do find some significant differences between the two countries, and these differences temper our conclusions somewhat. In particular, while consumption and wealth appear to exert similar effects on SWB in both countries, income is much less important to SWB in Britain than it is in Russia. Though there are a number of potential reasons for this difference, which is most convincing is not yet clear.

The British data also allow us to probe in more depth the precise way in which individuals' wealth holdings affect their SWB — an issue towards which conflicting approaches have so far been taken in the literature (cf. Headey and Wooden, 2004 and Headey, Muffels, and Wooden, 2008 with Brown, Taylor, and Wheatley Price, 2005), but which the paucity of the Russian wealth data largely precluded us from examining. They appear to do so in a way that may be initially surprising: it is not an individuals' overall wealth that matters to SWB (indeed, if we assumed this were the case, we would mistakenly conclude that wealth did not matter at all); instead, individuals appear to evaluate different types of wealth (such as housing, mortgage debt, savings, investments, and non-mortgage debt) separately, and also to value some of these wealth types more than others. The pattern of these differences suggests that wealth is important because individuals' future (anticipated) consumption matters to their current well-being (rather than having a direct effect, or because it contributes to individuals' sense of status or self-esteem). In particular, wealth that is more closely connected to specific, near-term consumption is more important to SWB than is non-specific savings or long-term wealth holdings. Collectively, the findings lend support to behavioural models that incorporate anticipatory consumption into the utility function (see e.g. Loewenstein, 1987; Elster and Loewenstein, 1992; Caplin and Leahy, 2001; Köszegi, 2005), and the general view that individual motivations for the accumulation of wealth do not reside solely in the rational weighing of current against future utility implicit in the standard discounted utility model.

The chapter proceeds as follows. Section 3.1 gives a brief overview of the British Household Panel Survey (BHPS). Section 3.2 investigates the importance of income to SWB in Britain. Section 3.3, the most substantive portion of the chapter, examines the importance of wealth. Section 3.4 concludes.

3.1 The BHPS

We employ waves 7-10 and 12-15 of the British Household Panel Survey (BHPS), a nationally representative survey tracking approximately 5,500 households and around 10,000 individuals each year from 1996-2005. (Wave 11, carried out in 2001 is excluded because the life-satisfaction question we use as a dependent variable throughout is not available.) The key variables we employ are described briefly below. A more detailed account of the constructed variables can be found in Appendix A.4.

In each year adult household members were asked: "how satisfied or dissatisfied are you with your life overall", with possible answers ranging from "not satisfied at all" to "completely satisfied" on a 7-point scale.[2]

In addition, the BHPS also includes a number of questions asking individuals and households about their consumption and income. We follow Deaton and Zaidi (2002) in constructing a household consumption measure consisting of four components: food expenditure and non-food expenditure, (imputed) residential rent, and a measure of the user cost of durables. The details of the procedures used to construct these measures are detailed in Appendix A.4. As a household income measure, we use the household's income in the last month (semi-annualised for comparability with the consumption measure). All measures are deflated using monthly CPI data.

Finally, the BHPS also includes a number of wealth measures. In all

[2]Although the phrasing might appear somewhat ambiguous as to the time frame concerned, the context makes clear that *current* satisfaction is the focus. The question is preceded by the following statement:

> Here are some questions about how you feel about your life. Please tick the number which you feel best describes how dissatisfied or satisfied you are with the following aspects of your current situation.

Respondents are then asked a series of questions about the extent to which they are satisfied with various aspects of their lives (including health, financial situation, job and relationships), which culminate in the overall life satisfaction question above.

rounds, it included questions about the value of one's own home, other property and mortgage debt. However, the other property measure is reported only as a categorical variables until wave 10, so we have useable data only from that point on. In addition, in waves 10 and 15, the BHPS included an additional module of wealth questions which garnered information regarding the savings, investments and non-mortgage debt of each household. Savings is a stock variable including money held in savings or deposit accounts (with a bank, post office or building society), National Savings Bank (Post Office), and TESSAs and ISAs. Investments are comprised of money held in National Savings Certificates, Premium Bonds, Unit or Investment Trusts, Personal Equity Plans, UK or foreign shares, National Savings Bonds, or other investments. Finally, non-mortgage debt consists of money owed on hire purchase agreements, personal loans (from banks, building societies, or other financial institutions), credit cards (including store cards), catalogue or mail order purchase agreements, DSS Social Fund loans, any other loans from private individuals, overdrafts, student loans, or other (non-mortgage) debt).

Figure 3.1 displays the evolution of the mean of SWB, and of (logged) household values for the income, consumption, and own home housing wealth variables over the sample; while Figure 3.2 displays the values for savings, investments, other property and non-mortgage debt. These figures confirm that there has been much less overall movement in most variables than was the case in Russia. There is evidence of some degree of common movements in the wealth variables: after an initial fall, mean logged housing wealth rises over the sample, as does the mean logged value of other property; while savings, investments, and non-mortgage debt each fall over the sample for which they are available. Nonetheless, it seems that Britain should provide a good test of the effect of consumption, income, and wealth in "normal" times.

3.2 The Importance of Income to SWB in Britain

This section focuses on the importance of income and consumption to SWB in Britain, and attempts to answer two specific questions: whether income regressions understate the importance of money to SWB in Britain; and whether, contrary to the standard consumption model, current income contributes to SWB independently of current consumption. In doing so, it follows essentially the same strategy as that adopted in Section 2.5, running regressions of the form:

$$u_{it} = m'_{it}\beta + z'_{it}\theta + \eta_i + \varepsilon_{it} \tag{3.1}$$

Figure 3.1: The Evolution of Log Household Income, Consumption, Housing Wealth, and Life-Satisfaction in Britain

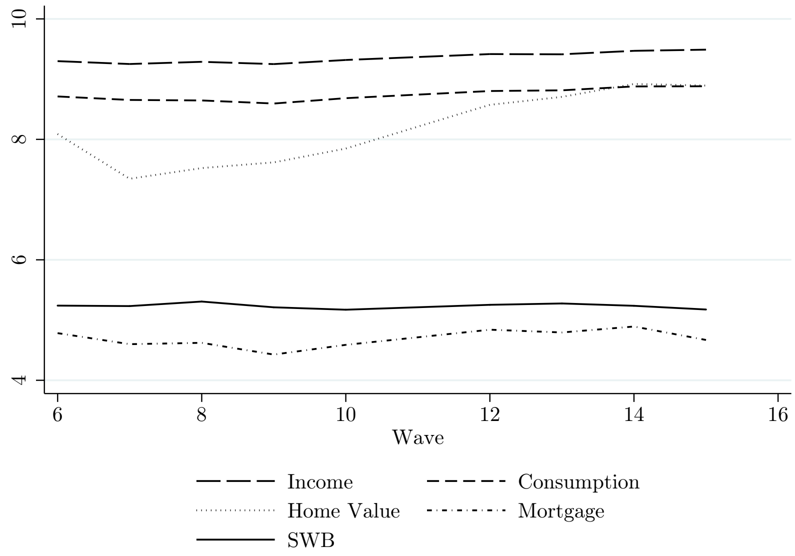

Figure 3.2: The Evolution of Non-Housing Wealth in Britain

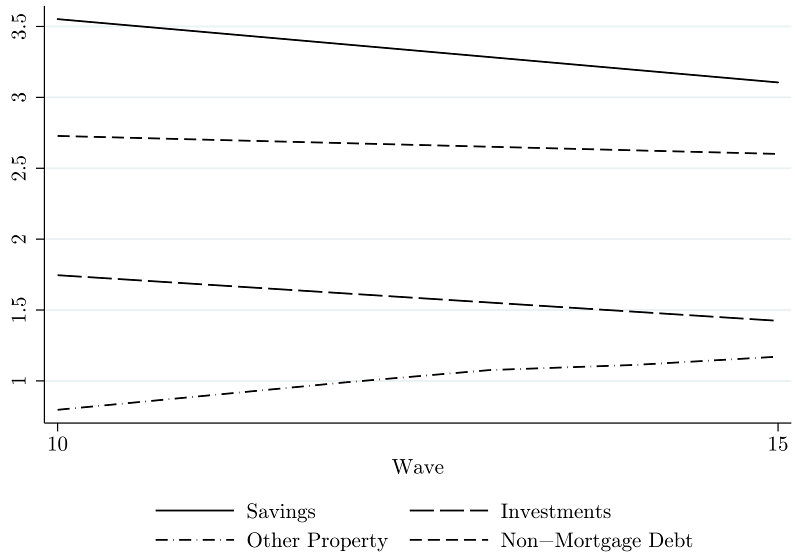

Throughout, we use as the dependent variable individual answers to the 7-point life-satisfaction question described in Section 3.1 above, and we employ again a standard set of control variables z_{it}, comprised of age, age squared, health, marital status (dummies for each of married, non-married couple, divorced, separated and widow(er), with never married as the omitted category), an employment dummy, and family size. We also include a full complement of wave dummies to allow for common shocks.

We test three alternative models by varying the components of m_{it}: (1) the standard income regression employed in most of the SWB literature in which $m_{it} = y_{it}$; (2) the consumption-based model of standard economic theory in which $m_{it} = c_{it}$; and (3) an alternative model that allows both consumption and income to affect current SWB and sets $m_{it} = (y_{it}, c_{it})'$. We leave aside for now the fourth model from Section 2.5, where we also allowed wealth to enter into the equation. Variants of this model will occupy us in much greater detail in in Section 3.3. The preliminary analysis favours model (3) over both models (1) and (2): income *does* seem to matter to SWB independently of consumption. We argue that this effect is unlikely to be due to measurement error in consumption, reverse causation, or omitted job characteristics (Section 3.3 further suggests that it is unlikely to be an unmeasured wealth effect); however, in contrast to our previous results, income does appear to matter much less to SWB in Britain than in Russia.

A corollary of this is that, while the standard income regression of model (1) greatly understates the importance of money to SWB, the consumption model (2) does so to a much smaller extent.

3.2.1 Basic Results

Table 3.1 displays the results of models (1)-(3). In all specifications, the control variables have effects broadly consistent with previous findings: being in a stable relationship has a strong positive effect (slightly larger for the unmarried than the married, but not significantly so), while divorce, widowhood, and separation have large negative effects relative to having never been married. The effect of being separated appears to be significantly *more* negative than being divorced, suggesting that divorce may actually be an improvement on separation.[3] Being employed has a significant positive effect over and above any monetary gains it induces, and health too is very important. The coefficients also suggest that, controlling for these factors, SWB declines with age. The usual positive coefficient on the squared term is insignificant here, indicating that the usual u-shaped age profile seems to be absent (though given the likely sensitivity of the age coefficients to the

[3]This fits with the argument advanced by Gilbert (2006), that the anticipation of an uncertain negative event may be worse than the actual experience of it, because our standard psychological defense mechanisms cannot begin to operate until the uncertainty is resolved. Only once the event actually happens can we begin adjusting to it and making sense of it in a way that allows us to cope with it.

controls included we would not want to over-interpret this).

Table 3.1: The Effect of Income and Consumption on Life Satisfaction in Britain, 1996-2005

	(1)	(2)	(3)
(Log) Household Income	0.016***		0.012**
	(2.765)		(1.966)
(Log) Household Consumption		0.110***	0.106***
		(6.195)	(5.920)
Age	-0.027**	-0.028**	-0.028**
	(-2.164)	(-2.211)	(-2.230)
$\text{Age}^2/10$	0.001	0.001	0.001
	(1.227)	(1.574)	(1.565)
Health	0.216***	0.216***	0.216***
	(34.915)	(34.834)	(34.849)
Married	0.095***	0.101***	0.099***
	(3.264)	(3.467)	(3.425)
Couple	0.121***	0.133***	0.131***
	(4.989)	(5.513)	(5.396)
Widowed	-0.199***	-0.193***	-0.192***
	(-3.377)	(-3.291)	(-3.271)
Divorced	-0.114**	-0.104**	-0.104**
	(-2.491)	(-2.271)	(-2.270)
Separated	-0.305***	-0.298***	-0.297***
	(-6.137)	(-5.990)	(-5.971)
Employed	0.084***	0.087***	0.082***
	(6.012)	(6.371)	(5.889)
Household Size	-0.023***	-0.034***	-0.036***
	(-3.731)	(-5.120)	(-5.423)
Observations	110744	110744	110744
Groups	22918	22918	22918
R^2	0.029	0.029	0.029

* $p<0.10$, ** $p<0.05$, *** $p<0.01$. t-statistics (robust to clustering within pids) in parentheses.

Turning to the substantive coefficients of interest, we see that the effect of income in column (1) is quite small, both when compared to the coefficients on the controls, and in relation to the consumption coefficient in column (2): in fact the latter is almost seven times larger. Yet while the results of the first two columns are consistent with the idea that income may merely be a noisy proxy for consumption, the results in column (3) throws some doubt on the adequacy of model (2): when included together, both income and consumption are still significant, and their effects are barely different from those of the first two columns.

Taking these estimates at face value suggests two broad conclusions. First, that consumption and income *both* affect SWB. And second, that the standard income regression of model (1) severely understates the importance of money to SWB. The precise degree of this understatement will depend on how income is allocated between consumption in different time periods; but we can get a rough idea of the total effect assuming that all income is consumed in the period it is received. This implies that the total effect is $y_{it}(\hat{\beta}_y + \hat{\beta}_c) = 0.118 y_{it}$, which is 7.4 times larger than the measured effect of (log) income from column (1) of $0.016 y_{it}$.

3.2.2 Is the Effect of Income Independent?

There are nonetheless a number of potential reasons why one might doubt these conclusions, and in particular the claim that the effect of income is independent of the effect of consumption. The following subsections address the two most obvious of these in turn: measurement error in consumption, and omitted variables bias. They argue that neither challenges our basic point.[4]

[4] For reasons set out in Section 2.5.2 we do not expect reverse causation to play a significant role in driving the positive coefficient on income: any such causal link should be largely absorbed within our fixed effects, or should operate with sufficient lags as to

Measurement Error in Consumption

As we noted in the previous chapter, one potential reason to doubt the conclusion that the income effect is truly distinct from the consumption effect is that the former may simply be picking up measurement error in the latter. While we cannot conclusively rule out such an explanation, the one test we can conduct suggests that it is unlikely to drive the result. As we noted in Section 3.2.2, if measured income were truly orthogonal to SWB except for its effect on consumption, then it would be a valid instrument for it. To examine the plausibility of this suggestion, we estimate an IV equation, instrumenting for consumption with income. The resulting coefficient on consumption is 0.370, almost 3.5 times larger than the initial consumption estimate of 0.110.

Under the standard classical errors in variables assumptions, this would imply that measurement error constitutes somewhere in the vicinity 70% of the variation in the measured deviations from mean consumption.[5] The implausibility of this estimate suggests that income does indeed have an independent, albeit small, effect on SWB in Britain.

avoid contaminating our estimates.

[5]Calculated from the standard formula as $\sigma_u^2/(\sigma_{c^*}^2 + \sigma_u^2) = (\beta_{c^*}/\hat{\beta}_c^{FE} - 1)/(\beta_{c^*}/\hat{\beta}_c^{FE})$.

Omitted Job Characteristics

Another possible reason to doubt the independent income coefficient is that it may simply be proxying for desirable job characteristics that happen to be correlated with it. To test this hypothesis we include five such potential variables in the baseline regression of model (3) (run on the employed sub-sample): a dummy indicating whether the individual holds a managerial position; two measures of professional status (RG Class and Socio-Economic Group); and two subjective satisfaction with work variables (satisfaction with job security, and satisfaction with work itself). Table 3.2 shows the results of these regressions. With the exception of the manager dummy, all the additional variables exert a statistically significant and positive effect on SWB. However, in no case do they affect the point estimates or significance of the income and consumption coefficients.[6]

3.3 Wealth

This section moves on to examine the effect of wealth on SWB within the BHPS. We build on existing work in this area in three main ways. First, we

[6]It is perhaps noteworthy that the income coefficient is both higher, and more significant within the employed sub-sample regressions here than it is in the full sample regressions of Table 3.1. Running additional regressions on the non-employed sub-sample confirms that the effect appears smaller and less precisely measured. However, this imprecision also means that the difference between the two is insignificant.

Table 3.2: The Effect of Income, Consumption and Job Characteristics on Life Satisfaction in Britain, 1996-2005

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
(Log) Household Income	0.027***	0.027***	0.026***	0.026***	0.024**	0.029***	0.027***
	(2.834)	(2.789)	(2.739)	(2.686)	(2.486)	(3.064)	(2.815)
(Log) Household Consumption	0.111***	0.111***	0.111***	0.110***	0.109***	0.108***	0.107***
	(5.315)	(5.304)	(5.309)	(5.290)	(5.246)	(5.283)	(5.240)
Manager		0.009					
		(0.752)					
RG Class			0.015**				
			(2.554)				
Socio-Economic Group				0.005**			
				(2.485)			
Satisfaction With Job Security					0.049***		0.032***
					(15.731)		(10.072)
Satisfaction With Work Itself						0.099***	0.090***
						(25.330)	(22.725)
Observations	67486	67486	67486	67486	67486	67486	67486
Groups	15775	15775	15775	15775	15775	15775	15775
R^2	0.071	0.072	0.071	0.071	0.088	0.126	0.131

* p<0.10, ** p<0.05, *** p<0.01. t-statistics (robust to clustering within pids) in parentheses. All standard controls except the Employment Dummy (Age, Age Squared, Health, Marital Status, Household Size, and Wave Dummies) included but not reported.

seek to clarify the issue of *how* exactly wealth matters to SWB — an issue to which previous work has adopted somewhat conflicting approaches. Headey and Wooden (2004) and Headey, Muffels, and Wooden (2008) assume that it is (log) net wealth that matters to SWB, and find significant effects of this variable in Australia (Headey and Wooden, 2004; Headey, Muffels, and Wooden, 2008), Germany, the Netherlands, Britain, and Hungary (Headey, Muffels, and Wooden, 2008). In contrast, Brown, Taylor, and Wheatley Price (2005) assume that different types of assets and debts affect SWB separately, and find that while short term debt and savings have significant effects on SWB in Britain, housing wealth and investments have no discernable effect. We resolve this tension by comparing the predictive ability of a range of different models (including a number of models not addressed in any previous work), and find broad support for the disaggregated approach taken in Brown, Taylor, and Wheatley Price (2005): contra Headey and Wooden (2004) and Headey, Muffels, and Wooden (2008) it is not households' overall net position that affects SWB; instead, individuals evaluate different types of wealth (such as their own home wealth, mortgage and non-mortgage debt, savings, and investments) separately. Moreover, in contrast to suggestions made in other parts of the behavioural economics literature (e.g. Thaler, 1985, 1999; Thaler and Johnson, 1990) they appear to do so even when this is not in their interests.

Second, most of the existing work has been unable (often due to data limitations) to control for unobserved heterogeneity. This is true of all the estimations in Headey and Wooden (2004) and Brown, Taylor, and Wheatley Price (2005). If, as seems likely, greater wealth is correlated with unobserved personality traits that also increase SWB, then failure to control for such factors may lead us to overestimate the impact of wealth on SWB (see e.g. Ferrer-i-Carbonell and Frijters, 2004). Moreover, while Headey, Muffels, and Wooden (2008) do estimate some fixed effects equations to investigate the effect of wealth in the Netherlands and Hungary, these use standard-of-living satisfaction as the dependent variable, which will surely also overstate the impact of wealth to global SWB (as indeed it does in their OLS specifications).[7] In line with previous results, Hausman tests conclusively reject the random effects assumptions in all the specifications we have run here, and we thus report fixed effects estimates throughout. In contrast to the results of Headey and Wooden (2004) and Headey, Muffels, and Wooden (2008), once we control for fixed effects, net wealth no longer seems to have a significant effect on SWB. Moreover, while we confirm the robustness of Brown, Taylor, and Wheatley Price's result that short term savings and debt affect SWB, we find some evidence that housing wealth also matters.

Finally, we also attempt to make some headway in determining *why* wealth matters to SWB, an issue with the existing literature has left largely untouched. Our findings suggest that it is the types of wealth that are most likely to affect consumption in the short-term that have the greatest effect on SWB, consistent with the idea that wealth matters because the anticipation of future consumption affects current utility, rather than because of any effect on status or self-esteem (or because of measurement error in current consumption). This is in keeping with Senik's (2004; 2008) finding that individuals' future prospects are important to SWB, and provides further support to behavioural models that attempt to incorporate the value of future consumption into current utility (e.g. Caplin and Leahy, 2001; Köszegi,

[7] It is not entirely clear why they report life-satisfaction equations as their preferred OLS specification, but only standard-of-living satisfaction equations for their fixed effects.

2005).

Section 3.3.1 addresses the joint questions of whether and how wealth matters to SWB. Section 3.3.2 then examines the issue of why wealth appears to be important.

3.3.1 How Does Wealth Matter to SWB?

Wealth could potentially affect SWB in a variety of ways. Drawing on the behavioural economics literature on prospect theory and mental accounts, this section develops and tests the following four hypotheses about the nature of wealth's impact on SWB.

(1) Do different asset classes, such as housing, short-term, and long-term wealth have distinct effects on SWB?

(2) Do assets and debts have distinct effects on SWB (within asset classes)?

(3) Do individuals frame their wealth holdings so as to maximise SWB?

(4) Are individuals debt averse?

Only if the answer to the first three questions is negative would we conclude that it is simply a household's net position that matters to SWB.

Mental Accounting

Rather than thinking about their overall net wealth, individuals may instead think about different types of wealth as belonging to distinct mental accounts that each affect SWB separately. Housing wealth may then affect SWB independently of investments, for example, or savings may affect it independently of credit-card debt.

Some support for such an hypothesis can be found in the behavioural economics literature on "mental accounting" (see generally Thaler, 1980, 1985, 1990, 1999). In particular, Shefrin and Thaler's (1988; 1992) behavioural life-cycle (BLC) model suggests that such accounts play an important role in individuals' inter-temporal consumption-savings decisions. The details of the model need not concern us greatly here. The key point is that both its predictions and the empirical evidence seem to suggest that individuals have different propensities to consume out of different types of wealth (defined for each wealth type j, as $\partial C_t/\partial W_{jt}$) and in particular, that while individuals are quite happy to consume out of current assets, they are less willing to consume out of housing wealth (e.g. by borrowing against housing to fund current consumption), or long-term investments and pensions. This stands in contrast to the standard model of inter-temporal decision-making, which predicts that an individuals' consumption in each period should (absent complications such as credit constraints) be a direct function of her lifetime wealth.

Whether or not these three distinct *asset classes* (housing, long-term, and short-term wealth) affect SWB separately constitutes the first dimension along which we distinguish different possible wealth models, and corresponds to question (1) above. As it happens, these asset classes map neatly onto the components of wealth surveyed in the BHPS, which asks individuals about their home value, mortgage debt, investments, savings accounts, and non-mortgage debt. The structure of the BHPS questions also suggests an additional way (corresponding to question (2) above) in which individuals could separate their wealth holdings into distinct mental accounts: by treating assets and debts separately.

The combined possibilities of segregating asset classes (SAC) and/or segregating assets and debts (SAD) yield four candidate mental account structures, which are displayed in Figure 3.3. In the top left we have the Overall Net Wealth model, in which there is neither AC nor AD structure. In the bottom left, we have the fully disaggregated SAC-SAD model in which home value, mortgage debt, investments, savings, and non-mortgage debt can each contribute separately to SWB. The other two models (SAC and SAD) each constitute midpoints between these extremes.

Figure 3.3: Possible Mental Account Structures

Segregates Assets and Debts (SAD)?		Segregates Asset Classes (SAC)? No	Segregates Asset Classes (SAC)? Yes
	No	Total Net Wealth	Net Housing Wealth Net Investments Net Savings
	Yes	Total Assets Total Debt	Housing Assets Mortgage Debt Investment Assets Gross Savings Short-Term Debt

Hedonic Editing

One interesting possibility raised by the possible existence of separate mental accounts is that individuals could change the way they frame their wealth holdings in order to maximize their utility (see e.g. Thaler, 1985, 1999; Thaler and Johnson, 1990). The idea that they could do so is known as the "Hedonic Editing" hypothesis, and has found some support in existing experimental studies. To our knowledge, it has not been tested in the SWB literature. However, it may have interesting implications in the SWB context. For example, if individuals do not frame their wealth holdings so as to maximise their SWB, then (for reasons we explain below) they may derive significant disutility from debt, even if their overall net position is positive.

For s-shaped functional forms of the sort we assumed in the previous chapter[8] (and which correspond to the standard prospect theory value function), it will make sense to combine two mental accounts if either both accounts are in debt, or the two constitute a "mixed asset" (i.e. one account is a positive wealth holding and the other is a smaller debt).

Intuitively, £1000 of credit card debt seems much worse considered alone than if one either adds it to £100,000 of mortgage debt, or subtracts it from a £100,000 asset portfolio — at which point it hardly seems to matter at all. In both cases, we would be better off (at least emotionally) to combine the two accounts. The converse is true for assets: £1000 in a savings account is likely to seem more important considered separately than if considered together with either a £100,000 asset portfolio, or a £100,000 mortgage debt (the latter is an example of a "mixed debt"). In these cases, we would be (emotionally) best to think about the asset separately.[9]

[8] See also Equation (2.4.1) in Section 2.4.1 above.

[9] As e.g. Linville and Fischer (1991) note, the extent to which the final conclusion will apply to mixed losses generally will actually depend on both the precise shape of the value function, and the sizes of the losses and gains. If individuals are sufficiently loss averse, the utility of an asset in its own mental account may be less that the utility of (mentally) using it to reduce debt.

More formally, preferences for combining debts, and for segregating assets follow from the fact that, for a value function $V(\cdot)$, with diminishing marginal sensitivity to assets and debts, the following hold by definition:

$$V(d_1 + d_2) > V(d_1) + V(d_2) \tag{3.2}$$

$$V(a_1) + V(a_2) > V(a_1 + a_2) \tag{3.3}$$

where $d_i < 0 < a_j$, $\forall\, i, j$.

The desire to aggregate "mixed assets" (i.e. small debts coupled with larger assets) follows from the fact that, assuming (weak) loss aversion (defined here as $V(-x) \geq -V(x)$, $\forall\, x \geq 0$):

$$V(a + d) > V(a) + V(d) \tag{3.4}$$

In this vein, Thaler (1985, 1999) and Thaler and Johnson (1990) propose an "hedonic editing" (HE) model whereby individuals may combine existing mental accounts in order to maximise their utility, as just set out (and displayed in Table 3.3).[10] Thaler and Johnson (1990) have also posited an alternative "quasi-hedonic editing" (QHE) model which suggests that individuals are unable to combine multiple losses. The existing evidence is generally taken to be consistent with this latter model, over both the no editing, and the original HE hypothesis. (Indeed, the QHE hypothesis was originally formulated to accommodate apparent deviations from the HE story found in practice.) For reasons I will not go into here, it seems to me that this evidence is rather weaker than is usually supposed, especially in relation to the crucial prediction that individuals combine mixed gains, which would

---

holds for all $|d| < a$.

For "mixed losses", where $|d| \geq a$, the situation is more complex. Here, weak loss aversion implies that there will be a $k \geq 0$, such that:

$$V(a + d) > V(a) + V(d) \tag{3.5}$$

for $|d| - a < k$. It is only if $V(\cdot)$ is symmetric (i.e. $V(-x) = -V(x)$), that

$$V(a) + V(d) > V(a + d) \tag{3.6}$$

for all $|d| > a$, and mixed debts should not be combined (though, as we suggest below, a symmetric value function is quite consistent with the data, and implies that $k = 0$).

[10] We set aside the possibility that individuals might segregate the baseline accounts further. This seems justified on two grounds. First, as Thaler (1999) notes, there must be some limits to individuals' ability to engage in such segregation: otherwise they would segregate every individual gain into the smallest units imaginable, and could avoid diminishing marginal utility entirely! Moreover, Thaler and Johnson (1990) find that empirically, individuals do not tend to further segregate accounts that are separate in the initial framing.

distinguish QHE from a simple no editing model.[11] In any event, it seems worthwhile to revisit the question here, given the quite different context to those in which the HE and QHE models have been tested to date (which e.g. have exclusively involved flows of gains and losses rather than wealth stocks).

Table 3.3: Predictions of Alternative Models of Mental Account Editing

	No Editing	Hedonic Editing	Quasi-Hedonic Editing
(1) Multiple Assets	No Change	No Change	No Change
(2) Multiple Debts	No Change	Combine	No Change
(3) Mixed Asset	No Change	Combine	Combine
(4) Mixed Debt	No Change	No Change*	No Change*

*As footnote 9 notes, in the presences of loss aversion, this prediction is slightly ambiguous, but will hold for sufficiently large debts.

While the predictions of the various editing models are fairly straightforward when there are only two mental accounts that could be combined, matters become more complicated when there may be up to five (as in the SAC-SAD model, which has a separate mental account for each wealth type in the BHPS). If we knew the shape of the sub-utility functions for each

[11] Very briefly, the basic point is that such evidence typically relies on individuals' expressed preferences about whether they would rather experience a given pair of events on the same day, or on subsequent days. A preference for same-day experiences is then taken to reflect a preference for combining outcomes, and a preference for subsequent-day experiences is taken to reflect a preference for segregating outcomes. The problem in interpreting these studies is that the responses seem likely to be confounded by time-preference (e.g. a preference to experience event on the same day could be due to impatience rather than a desire to mentally combine their effects), and are consequently unable to provide clear evidence for any particular hypothesis.

mental wealth account, we could calculate the optimal framing directly (by simply comparing the SWB levels attainable under each possible framing and choosing the best one). However, because these functions are unknown — indeed, they are precisely what we are trying to estimate — this strategy is difficult to implement here.

Instead, we proceed by making three intuitive simplifying assumptions about how accounts may be combined. We assume that an account with a negative balance may be combined either with assets in the same asset class (if the underlying account structure distinguishes assets and debts) or with debts in other asset classes. However, given the evidence that individuals have some difficulty combining mental accounts at all, we assume that no account can be combined both across and within asset classes. (So for example, the SAC-SAD model would not allow us to combine short-term debt with housing wealth.) In addition to the null (no editing) model, these assumptions are consistent with three distinct models of editing:

(1) *Asset-Debt (AD) Editing* : debts and assets are combined *within* an asset class iff the resulting net position would be positive;

(2) *Debt-Debt (DD) Editing*: multiple debts are always combined, but debts are never combined with assets; or

(3) *Asset-Debt-Debt (ADD) Editing*: debts and assets are combined within

an asset class iff the resulting net position would be positive, and then any remaining debts are combined.[12]

AD editing corresponds to the QHE hypothesis (both rule out the possibility of combining multiple debts), while ADD editing corresponds to the original HE model. DD editing has no equivalent in the existing literature.

Depending on the underlying mental account structure, only some of these types of editing may be possible. In fact, the only mental account structure that could allow any of the three editing types is the fully disaggregated SAC-SAD model; the Total Net Wealth (TNW) model allows no editing at all (there is only one mental account to begin with); the SAC model automatically combines assets and debts in the same class, and can therefore allow only DD editing; while the SAD model has a single Total Debt account and so can only allow AD editing.

Debt Aversion

A final issue in the specification of the wealth sub-utility function(s) is whether individuals are debt averse: that is, whether debts reduce SWB more than assets lift it. The phenomenon has been fairly robustly demonstrated as a pervasive factor in decision-making, and appears to explain a number of

[12] The possibility of first combining all debts and then combining total debt with one or more asset accounts is ruled out by the restriction that assets and debts cannot be combined across the underlying asset classes.

anomalies not accounted for under standard theory, such as risk-aversion over small stakes (Rabin, 2000) and the endowment effect (Kahneman, Knetsch, and Thaler, 1991).

Nonetheless, its existence as an *experienced* utility phenomenon is potentially more contentious. Wilson and Gilbert (2005) and Kermer et al. (2006) argue that loss-averse decision-making may in fact be attributable to systematic "affective forecasting errors", and present evidence that, while individuals consistently tend to overestimate the hedonic impact of both positive *and* negative events, this overestimation is asymmetric: we overestimate the effect of negative events more than positive ones. If this were true, we could observe loss-averse decision-making, without losses actually having a greater impact on our SWB. In contrast, Di Tella, Haisken-De New, and MacCulloch (2007) find evidence consistent with the idea that declines in income have a greater effect than equivalent increases, and Brown, Taylor, and Wheatley Price (2005) find evidence of debt aversion in the BHPS. We revisit their findings in the context of a broader set of possible models, and making explicit allowance for the possibility that unobserved personality traits could potentially affect their cross-sectional results (e.g. if such traits were more strongly correlated with debt than asset accumulation).

Empirical Strategy

The first three dimensions along which we distinguish possible wealth models (asset class segregation, asset debt segregation and hedonic editing) combine to generate a range of possible models, which are summarized in Figure 3.4 as models 1-4C. Each of these have corresponding starred versions 1*-4C* that allow for debt-aversion. We also include, as model 0, the null model in which wealth has no effect, for a total set of 19 possible models. For any of these models that include more than one mental account (i.e. all but models 0 and 1), there will also exist a number of nested sub-models, in which one or more of the included wealth types has no effect on SWB. (For example, the SAC-SAD models will each include sub-models where, say, home value and mortgage debt are restricted to have no effect, but investments, savings, and non-mortgage debt are allowed to influence SWB.)

Figure 3.4: Overview of Possible Mental Account Models

Account	Editing Possibilities Permitted			
Structure	None	AD	DD	ADD
None	0			
TNW	1, 1*			
SAD	2, 2*	2A, 2A*		
SAC	3, 3*		3B, 3B*	
SAC-SAD	4, 4*	4A, 4A*	4B, 4B*	4C, 4C*

Our task in what follows is to select the model that best describes the effect of wealth on SWB (if indeed there is any). In order to do so, we first need to define each model precisely in terms of the variables (and restrictions) it involves. We then outline our strategy for selecting between the alternatives.

Variable Definitions

In the basic cases (models 1, 2, 3 and 4, with no editing or debt-aversion) the relevant variables are just those set out in 3.3, transformed as follows:

$$a_{it} = \begin{cases} \ln(A_{it}+1) & \text{if } A_{it} \geq 0 \\ -\ln(1-A_{it}) & \text{if } A_{it} < 0 \end{cases} \tag{3.7}$$

To allow for debt-aversion, we simply let the coefficients on positive and negative wealth vary. (If the former is significantly less than the latter, then this provides evidence in favour of loss aversion.) This means different things in the context of SAD and non-SAD models. In the former, allowing for loss aversion or not is achieved by allowing the coefficients on the asset and debt measures within each asset class to vary, or by restricting them to be equal. In the second case, where the model combines assets and debts to form net wealth accounts (by asset class) allowing for loss aversion requires us to modify the variable definitions above, instead defining two separate variables as follows:

$$a_{it}^{+} = \begin{cases} \ln(A_{it}+1) & \text{if } A_{it} \geq 0 \\ 0 & \text{if } A_{it} < 0 \end{cases} \tag{3.8}$$

$$a_{it}^{-} = \begin{cases} 0 & \text{if } A_{it} \geq 0 \\ -\ln(1-A_{it}) & \text{if } A_{it} < 0 \end{cases} \tag{3.9}$$

Finally, to allow for different types of hedonic editing, we redefine the variables in (3.7) as follows.

(1) For AD editing, we combine assets A_k^+ and debts A_k^- within an asset class k, whenever the total would be positive; otherwise we leave them as is. This gives us two variables (within each asset class), defined as:

$$a_k^+ = \ln(1 + A_k^+ - \mathbb{1}(A_k^+ \geq A_k^+)A_k^-) \tag{3.10}$$

$$a_k^- = -\mathbb{1}(A_k^+ < A_k^-)\ln(1 - A_k^-) \tag{3.11}$$

(2) For DD editing, we leave assets as they are, but combine any debts (or negative wealth) to create a new total debt variable:

$$a_k^+ = \ln(1 + \mathbb{1}(A_k \geq 0)A_k) \tag{3.12}$$

$$a^- = -\ln\left(1 - \sum_{k:A_k<0} A_k\right) \tag{3.13}$$

(3) Finally, for ADD editing, we first combine assets and debts within each class (if the total would be positive), and then combine any remaining debts into a total debt variable:

$$a_k^+ = \ln(1 + A_k^+ - \mathbb{1}(A_k^+ \geq A_k^+)A_k^-) \tag{3.14}$$

$$a^- = -\ln\left(1 - \sum_{k:A_k^+ < A_k^-} A_k^-\right) \tag{3.15}$$

In the context of DD and ADD editing models, *non*-loss aversion models require that we restrict the coefficients on all *positive* assets to be equal.[13]

Model Selection

With these variable definitions in place, the task of choosing a preferred model proceeds in two stages. In the first, we follow a general-to-specific (GETS) strategy *within* each applicable model type (1-4C*) to identify the best candidate model of its class. In the second, we compare the best candidate models from each class in terms of information criteria to determine which performs better.

The GETS approach we adopt is a simplified version of that outlined in e.g. Hendry and Krolzig (2002). For each model class, we begin by fitting

[13] This is because combining debts across asset classes requires us to assume that *their* effects are equal. Consequently, equating the impact of assets and debts within asset classes implies that the impact of assets is equal across asset classes. Formally, if j and k denote asset classes: $(\beta_{a_j^-} = \beta_{a_k^-}, \forall\, j,k \;\&\; \beta_{a_k^+} = \beta_{a_k^-}, \forall\, k) \Rightarrow \beta_{a_j^+} = \beta_{a_k^+}, \forall\, j,k$.

the most general model possible, including all the wealth variables we have available. Whichever of these are significant are retained. We then explore multiple reduction paths, beginning by dropping one or more insignificant wealth variables, and iterating this process until either (a) only significant wealth variables remain or (b) no significant wealth variables remain.[14] If the reduction paths converge on a single model, then this is treated as the best candidate model of its class. If more than one model emerges, we compare these in terms of information criteria (described below) to determine the best candidate model. The exception to this strategy is if no variables are significant in any stage of the equation. In this case the reduced model will be identical to the baseline model 0, where wealth has no effect. In such cases we retain the initial model in the second stage, but only for illustrative purposes (it would make no substantive difference to the results if it were simply dropped at this stage).

In the second stage, we take the best candidate model from each relevant class and compare these in terms of two information criteria.[15] Intuitively,

[14]The standard for significance is kept fairly lax at this stage, retaining any variables with an absolute t-statistic greater than one. This is primarily useful for illustrative purposes when comparing models, as models including such variables will almost inevitably be dominated in the second stage of the analysis described below.

[15]This mimics the terminal step in GETS strategy advocated by Hendry and Krolzig (2002). While the General to Specific approach typically precedes selection by information criteria by a stage of encompassing testing (see e.g. Hendry and Krolzig, 2002), where models that cannot explain the results of other remaining models are excluded, because the models we are concerned with here are all essentially transformations of the same variables, we do not anticipate that this should be a significant problem. We nonetheless treat the issue (albeit informally) in the discussion that follows.

such criteria select a preferred model by trading off its ability to explain the data (better goodness of fit) against the its complexity (lower parsimony). As one might expect, this trade-off can be made in a variety of ways, and there is an extensive literature comparing the performance of alternative criteria and attempting to analyse their costs and benefits. Here we focus on the two most common: the (Corrected) Akaike Information Criterion (AICc); and Schwarz's Bayesian Information Criterion (BIC).

The AICc (see e.g. Akaike, 1974; Burnham and Anderson, 2003) is defined as $\text{AICc} = -2\ln(L) + 2k\frac{N}{N-k-1}$, where L is the likelihood, N is the sample size, and k the number of estimated model parameters.[16] The BIC (see e.g. Schwarz, 1978) is defined as $\text{BIC} = -2\ln(L) + \ln(N)k$. Both criteria select models with the *lowest* criterion value, based on a combination of their goodness of fit (the $-2\ln(L)$ term) and a "penalty" for model complexity, which helps to prevent over-fitting (the latter term, which in both cases is a function of the number of parameters, k). The penalty applied by the BIC is typically rather higher than that applied by the AICc.[17] However, it

[16]The original AIC proposed by Akaike (1974) was formulated as an asymptotic approximation to the Kullback-Liebler distance (Kullback and Leibler, 1951), and was defined as $\text{AIC} = -2\ln(L) + 2k$. The corrected version in the main text adjusts the second term to account for the small sample bias. (Though, given the size of our sample relative to the number of parameters, there is little difference between the two here.)

[17]This means that the AIC tends to select more complicated models (with a correspondingly higher risk of over-fitting), while the BIC tends to select less complex models (with a correspondingly higher risk of omitting potentially relevant variables from the analysis). How we assess this trade-off will depend on how complex we think that the reality we are trying to approximate is, relative to the set of models we consider. If we expect that model complexity is unrelated to the prior probability of a given model being closer to the

turns out that the difference between the two does not matter greatly for our purposes, and quite similar patterns emerge regardless of which we choose to focus on.

Results

This section presents the results of our model selection strategy, and draws six substantive conclusions: (1) wealth matters; (2) assets in different classes affect SWB separately; (3) assets and debts affect SWB separately; (4) individuals do not appear to engage in any systematic form of hedonic editing; (5) there is little evidence that individuals are debt averse; and (6) the most important forms of wealth to SWB are those most closely linked to consumption in the near future. From a methodological perspective, the evidence also supports our initial suspicion that misspecification of the sub-utility function in wealth could lead to incorrect inference: had we simply assumed that net wealth was the appropriate argument in the SWB function, and ignored the possibility that individuals might operate distinct mental accounts, we would have incorrectly concluded that wealth was unimportant to SWB.

We address the evidence for each of the above claims in more detail below. In the first part, we discuss our main results, derived from regressions run on

truth (or, more strongly, that a simpler model is likely to *better* approximate the truth) this will tend to favour the BIC. On the other hand, the more complex we think reality is, the more we will tend to favour the AIC. See e.g. Burnham and Anderson (2004).

the sub-sample of the BHPS for which all wealth data are available (waves 10 and 15). We then demonstrate that these results are robust in the larger BHPS sub-sample in which only housing wealth is available. Finally, we address a slight anomaly in the main results, which sees the the income coefficient increase and the consumption coefficient lose significance when we are restricted to only the 2000 and 2005 samples. This change appears to be due to the stock market crash of 2000, and to disappear once the main impact of the crash receded. While it does not therefore seem to significantly undermine the results of Section 3.2, it nonetheless highlights the crucial importance of context to the determinants of SWB.

All Wealth Sample

Table 3.4 shows the results of the final candidate models in which different asset classes are not segregated (No Wealth, Total Net Worth and SAD models: 0-2A*). A few points stand out. First, the coefficients on household income are much larger than in the full sample regressions of Section 3.2, while the coefficients on household consumption have fallen and become insignificant. (Both of these changes occur even in the baseline No Wealth model, and the coefficients change little as we include various wealth variables, so it is not the introduction of the wealth variables *per se* that is responsible.) We shall address both of these issues in more detail shortly, but for now we focus on the wealth coefficients. It is immediately apparent that in none of the models are any of these coefficients significant: if indeed wealth is important to SWB, the non-SAC models do not capture its effect.

In contrast, Table 3.5 displays the results of the SAC models (SAC and SAC-SAD models, 3-4C); and here, the results look more promising (though the income and consumption coefficients maintain their odd behaviour). With the exception of three models (3*, 3B and 4A*) all other models emerge with at least one significant wealth variable. These results further suggest that it is the savings and non-mortgage debt variables that are most likely to be important to SWB. Savings is significant in all the remaining equations except 3B*, with coefficients ranging from 0.004 in model 3 to 0.007 in model 4; and non-mortgage debt is significant in all the remaining equations except the final editing models 4B and 4C*, with coefficients ranging from -0.004 in model 3, to -0.009 in model 4*. On the other hand, housing is insignificant in all the models, and mortgage debt only achieves significance when its effect is constrained to equal that of non-mortgage debt in model 3B* (SAC DD editing with debt aversion). Investments are never significant, and end up being dropped from every equation.

While these results are suggestive, we turn to information criteria to ultimately determine whether any of the models dominate the null No Wealth model, and if so which model is most plausible. These criteria are displayed

Table 3.4: The Effect of All Wealth Sources on Life Satisfaction in Britain (Non-SAC Models)

	Model						
	0	1*	1	2*	2	2A*	2A
(Log) Household Income	0.049**	0.048**	0.048**	0.049**	0.049**	0.048**	0.048**
	(2.458)	(2.439)	(2.427)	(2.482)	(2.464)	(2.437)	(2.436)
(Log) Household Consumption	0.066	0.067	0.059	0.064	0.060	0.066	0.060
	(1.056)	(1.007)	(0.899)	(0.971)	(0.955)	(0.986)	(0.916)
(Log) Positive Wealth		-0.003	0.002	0.001	0.003	-0.002	0.001
		(-0.445)	(0.510)	(0.224)	(0.766)	(-0.284)	(0.349)
(Log) NegativeWealth		-0.009	-0.002	-0.003	-0.003	-0.005	-0.001
		(-0.884)	(-0.510)	(-0.831)	(-0.766)	(-0.766)	(-0.349)
Observations	23719	23719	23719	23719	23719	23719	23719
Groups	17002	17002	17002	17002	17002	17002	17002
R^2	0.001	0.001	0.001	0.001	0.001	0.001	0.001
AIC	32057.83	32058.17	32058.88	32059.23	32057.60	32059.54	32059.36
BIC	32162.79	32179.28	32171.91	32180.35	32170.64	32180.65	32172.40

* $p<0.10$, ** $p<0.05$, *** $p<0.01$. t-statistics (robust to clustering within pids) in parentheses.

Table 3.5: The Effect of All Wealth Sources on Life Satisfaction in Britain (SAC Models)

	Model									
	3*	3	3B*	3B	4*	4	4A*	4A	4B	4C*
(Log) Household Income	0.048**	0.048**	0.048**	0.047**	0.048**	0.048**	0.048**	0.048**	0.049**	0.049**
	(2.421)	(2.416)	(2.441)	(2.389)	(2.404)	(2.398)	(2.404)	(2.405)	(2.456)	(2.467)
(Log) Household Consumption	0.061	0.061	0.066	0.051	0.058	0.059	0.061	0.061	0.066	0.067
	(0.979)	(0.987)	(0.980)	(0.802)	(0.931)	(0.944)	(0.976)	(0.977)	(1.060)	(0.991)
(Log) Savings		0.004*		0.003	0.006*	0.007***	0.005	0.005**	0.005*	0.006*
		(1.902)		(1.622)	(1.713)	(2.776)	(1.606)	(2.266)	(1.667)	(1.772)
(Log) Non-Mortgage Debt	-0.006	-0.004*	-0.008*	-0.003	-0.009**	-0.007***	-0.005	-0.005**		-0.003
	(-1.246)	(-1.902)	(-1.900)	(-1.622)	(-2.291)	(-2.776)	(-1.365)	(-2.266)		(-1.088)
(Log) Housing Wealth				0.003						-0.000
				(1.622)						
(Log) Housing Debt			-0.008*	-0.003						
			(-1.900)	(-1.622)						(-1.088)
Observations	23719	23719	23719	23719	23719	23719	23719	23719	23719	23719
Groups	17002	17002	17002	17002	17002	17002	17002	17002	17002	17002
R^2	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001	0.001
AIC	32049.92	32048.54	32044.68	32051.49	32035.34	32034.65	32045.27	32043.27	32050.51	32049.04
BIC	32171.03	32161.58	32173.86	32164.53	32156.45	32147.69	32166.38	32156.31	32171.62	32178.22

* $p<0.10$, ** $p<0.05$, *** $p<0.01$. t-statistics (robust to clustering within pids) in parentheses.

for each model in Figures 3.5 and 3.6. In both cases, we normalise the criteria, expressing them as deviations from the baseline model 0 (so the values for the baseline model are always equal to zero, and negative values are preferred). Although there are some differences in the ranking of the non-preferred models depending on whether we focus on the AICc or the BIC, both criteria select model 4 (SAC-SAD, No Editing, No Debt Aversion) as the preferred choice. Pairwise comparisons confirm that (1) wealth matters (all the SAC models dominate the baseline model 0 under the AICc, and 4 models do so under the BIC); (2) the SAC models (3*, 3, 4*, 4, 4A* and 4A) dominate their corresponding non-SAC counterparts (1*, 1, 2*, 2, 2A* and 2A); (3) the best SAD models (2,4* and 4) dominate the corresponding non-SAD models (1,3* and 3); (4) there is no suggestion of any sort of systematic hedonic editing (2*, and 2 dominate 2A* and 2A*; 3* and 3 dominate 3B* and 3B; and 4* and 4 each dominate their respective editing counterparts 4A* and 4C*, and 4A and 4B); (5) the preferred equal coefficient model (4) dominates the model that allows for debt aversion (4*) (and an F-test of the restriction that the coefficients on savings and debt are equal in 4* fails to reject the hypothesis); and finally (6) the short-term savings and non-mortgage debt variables are the most (indeed only!) important wealth variables to SWB.

Figure 3.5: AICc Statistics for All Wealth Models

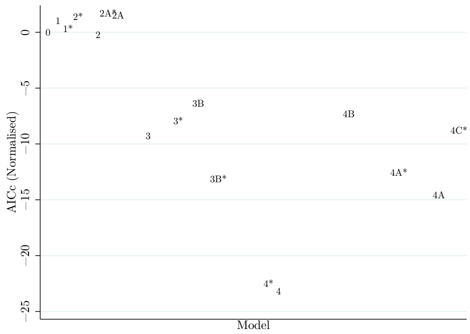

Figure 3.6: BIC Statistics for All Wealth Models

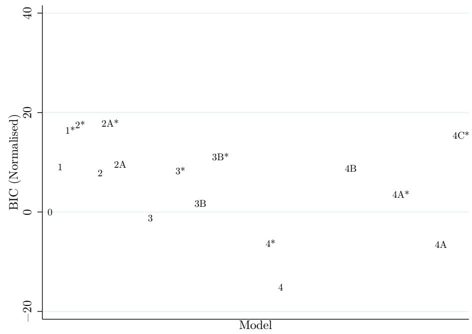

AICc and BIC statistics for each model run on the All Wealth sample. Models are defined as in Figure 3.4: 0=No Wealth; 1=Net Wealth; 2=SAD; 3=SAC; 4=SAC-SAD; A=AD Editing; B=DD Editing; C=ADD Editing; *=Debt Aversion. Statistics are expressed as deviations from the baseline (No Wealth) model 0. Smaller values of the statistic indicate a preferred model according to the criterion; and negative values indicate that a model outperforms the baseline model 0. Both statistics prefer model 4.

Own Home Housing Wealth

As a check on these conclusions, we have also conducted a similar analysis using the (larger) sub-sample of the BHPS for which we have only information about the household's own home housing wealth (i.e. excluding other real estate wealth).[18] While this data includes only one asset class, and thus does not allow us to distinguish explicitly between SAC and non-SAC models, in almost all other respects, the analysis confirms the main conclusions set out above. The only exception is that, perhaps due to having more rounds of data, both the value of one's own home, and the level of a household's outstanding mortgage debt, now attract significant coefficients (positive and

[18]We also experimented with using the sub-sample for which we have only non-real property wealth; unsurprisingly, the results were almost identical to those just presented, and we do not discuss them any further here.

negative, respectively).

Table 3.6 shows the results of the final candidate models for each model class in the Own Home Housing Wealth sample. As already noted, having data on only one asset class (housing) means we cannot explicitly distinguish between SAC and non-SAC models, so for the purposes of this analysis, we treat models 1 and 3, and models 2 and 4 (and their respective variants) as identical. (Technically of course, the models we test *do* segregate housing assets from other assets, and assume that the impact of other assets is zero. Consequently, if we find a model that outperforms the baseline model 0 in this sample, then this may provide limited support for an SAC structure. However, because we cannot compare such a model to a more comprehensive one, this evidence is at best suggestive.)

Looking at the results, we notice that both the size and significance of the consumption coefficient have returned, but that now it is the income coefficient that has become insignificant (apparently confirming the relative magnitudes of the two which we estimated in Section 3.2). Shifting our focus to wealth, we see that only two model pairs attract significant coefficients on any wealth variables: the No Editing SAD models $\{2, 4\}$, and $\{2*, 4*\}$, in which assets and debts enter separately. In both cases (whether we restrict the coefficients or not) the coefficients on both home value and mortgage debt are 0.006 and -0.006 respectively, and an F-test verifies that these coefficients are not significantly different from each other. In contrast, the net wealth specifications in models $\{1, 3\}$ and $\{1*, 3*\}$ would suggest that wealth does not enter into the SWB function at all, and the hedonic editing models $\{2A, 4A\}$, and $\{2A*, 4A*\}$ also fail to attract any significant coefficients.

The information criteria displayed in Figures 3.7 and 3.8 support the intuitive conclusions suggested by Table 3.6. Again, although there are some differences in the ranking of the non-preferred models depending on whether we look at the AICc or the BIC, both criteria rank the model pair $\{2, 4\}$ as the preferred choice. More generally, pairwise comparisons accord with those from our main regressions, suggesting that: (1) wealth matters ($\{2, 4\}$ dominates the baseline model on both criteria and $\{2*, 4*\}$ also does on the AICc); (3) SAD models ($\{2, 4\}$ and $\{2*, 4*\}$ dominate the corresponding non-SAD models ($\{1, 3\}$ and $\{1*, 3*\}$ respectively); (4) there is no evidence of systematic hedonic editing ($\{2, 4\}$ and $\{2*, 4*\}$ dominate their AD editing counterparts $\{2A, 4A\}$ and $\{2A*, 4A*\}$); and finally (5) there is little evidence of debt aversion in housing wealth ($\{2, 4\}$ dominates $\{2*, 4*\}$, and the BIC even selects 0 over the latter pair!). On the other hand, we now have some evidence from the larger sample that housing wealth may indeed matter to SWB, even though it did not show up in the wave 10 and 15 sample.

Table 3.6: The Effect of Income, Consumption and (Own Home) Housing Wealth on Life Satisfaction in Britain

	Model						
	0	{1*,3*}	{1,3}	{2*,4*}	{2,4}	{2A*,4A*}	{2A,4A}
(Log) Household Income	0.009	0.009	0.009	0.010	0.010	0.009	0.009
	(1.416)	(1.423)	(1.402)	(1.506)	(1.484)	(1.418)	(1.405)
(Log) Household Consumption	0.120***	0.120***	0.119***	0.113***	0.111***	0.120***	0.119***
	(6.077)	(5.637)	(5.647)	(5.364)	(5.563)	(5.617)	(5.603)
(Log) Positive Housing Wealth		-0.000	0.000	0.006**	0.006***	0.000	0.000
		(-0.053)	(0.255)	(2.272)	(3.238)	(0.007)	(0.150)
(Log) Negative Housing Wealth		-0.003	-0.000	-0.006***	-0.006***	-0.003	-0.000
		(-0.646)	(0.255)	(-3.266)	(3.238)	(-0.587)	(0.150)
Observations	95288	95288	95288	95288	95288	95288	95288
Groups	21757	21757	21757	21757	21757	21757	21757
R^2	0.003	0.003	0.003	0.004	0.004	0.003	0.003
AIC	225217.44	225219.34	225220.91	225204.27	225202.43	225221.02	225219.41
BIC	225406.74	225418.10	225429.13	225412.50	225401.19	225429.24	225418.16

* p<0.10, ** p<0.05, *** p<0.01. t-statistics (robust to clustering within pids) in parentheses. Models are defined as in Figure 3.4: 0=No Wealth; 1=Net Wealth; 2=SAD; 3=SAC; 4=SACD-SAD; A=CADI Editing; *=Loss Aversion (unstarred models restrict the asset and debt coefficients to sum to zero)

Figure 3.7: AICc Statistics for Own Home Housing Wealth Models

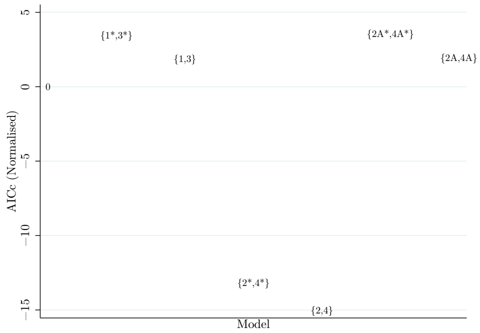

Figure 3.8: BIC Statistics for Own Home Housing Wealth Models

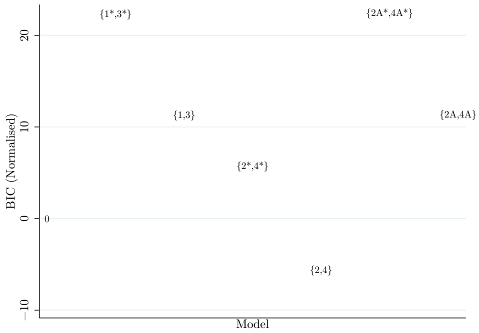

AICc and BIC statistics for each model run on the Own Home Housing Wealth sample. Models are defined as in Figure 3.4: 0=No Wealth; 1=Net Wealth; 2=SAD; 3=SAC; 4=SAC-SAD; A=CADI Editing; *=Loss Aversion. Statistics are expressed as deviations from the baseline (No Wealth) model 0. Smaller values of the statistic indicate a preferred model according to the criterion; and negative values indicate that a model outperforms the baseline model 0. Both statistics prefer the model pair {2,4}.

Consumption and Income in the All Wealth Sub-sample

As we noted above, the coefficients on income and consumption in our main results appeared to conflict with those obtained in Section 3.2. In particular, while its size only decreased slightly, the coefficient on consumption became insignificant, while the income coefficient approximately tripled. There are at least three possible reasons why this might occur. However, we argue that the most plausible is that the fall is due to the stock market crash of 2000. Once we account for this, consumption regains its former significance and the coefficient on income falls to around its previous level.

If consumption were merely picking up an omitted wealth effect in Table 3.1, then the inclusion of wealth variables could lead the consumption coefficient to become insignificant. However, this is unable to explain the fact that the coefficient becomes insignificant even in column (1) of Table 3.4) and that its size and standard error remains virtually constant even following the inclusion of the wealth variables in the subsequent models. Instead, the effect seems to stem from restricting the sample to only waves 10 and 15.

Another possibility is that the coefficient becomes insignificant due to a loss of real variation in our measures, and a concomitant increase in measurement error, as a result of restricting the sample to only two periods (a problem that may be particularly acute in a fixed effects context). However, this too seems unable to account for the patterns in the data. Table 3.7 supports the notion that running fixed effects on only two *contiguous* waves of data will tend to increase the noise-to-signal ratio in the consumption measure to such an extent that its coefficient becomes insignificant. Nonetheless, changes in household consumption over a 5-year period should not suffer to the same extent from this measurement error problem. Table 3.8 appears to confirm this, generating significant coefficients on consumption for all 5 year pairs except 2000 and 2005.

The conclusion that the change is not due to measurement error is further supported by Table 3.9. When we run regressions including 2005 data and data from each of the other years individually, it is only in the 2000/2005 and 2004/2005 samples that the consumption coefficient becomes insignificant. If measurement error were driving the result, we would expect the consumption coefficients to decline monotonically from the 1996/2005 pair to the 2004/2005 pair as the level of real variation decreases relative to measurement error. In particular, we would expect the 2002/2005, and 2003/2005 samples to also have insignificant consumption coefficients given its insignificance in the 2000/2005 sample.

A more substantive hypothesis is that the effect is related to the stock market crash of 2000, which saw the NASDAQ, Dow Jones and the FTSE shed respectively 72%, 47% and 46% of their value, beginning in September 2000 (and lasting until the latter stages of 2002 or early 2003). The BHPS surveys are mostly carried out in September and October of each year (though in wave 10, a significant number were still being completed up until May 2001) and these events could be expected to loom large in the minds of many respondents. On the one hand, this could potentially swamp the measured effect of consumption changes over the period; on the other, if a severe negative wealth shock reduces individuals' ability to sustain their current levels of consumption, higher consumption could be associated with greater anxiety about the future. This latter explanation has the virtue that it may also account for the increase in the size and significance of the income coefficient, in line with Senik's (2004) argument that the future may become more salient in times of high volatility, and the idea that part of the value of

Table 3.7: Life Satisfaction Income and Consumption Equations Using Contiguous Wave Pairs)

	Sample Years Included							
	96-97	97-98	98-99	99-00	00-02	02-03	03-04	04-05
(Log) Household Income	0.037	0.006	-0.014	0.016	0.028*	0.008	0.017	0.032**
	(1.560)	(0.250)	(-0.696)	(0.984)	(1.732)	(0.494)	(1.129)	(2.125)
(Log) Household Consumption	0.072	0.123**	0.107*	0.024	0.068	0.042	0.062	0.050
	(1.121)	(1.979)	(1.826)	(0.447)	(1.396)	(0.864)	(1.360)	(0.879)
Observations	18321	20078	23156	26433	27705	28800	28006	27257
Groups	10542	11112	14379	14999	18051	16206	15771	15236
R^2	0.075	0.001	0.008	0.071	0.000	0.010	0.040	0.121

* $p<0.10$, ** $p<0.05$, *** $p<0.01$. t-statistics (robust to clustering within pids) in parentheses.

Table 3.8: Life Satisfaction Income and Consumption Equations Using Data from 5-Year Wave Pairs

	Sample Years Included			
	97-02	98-03	99-04	00-05
(Log) Household Income	-0.001	0.012	0.009	0.030*
	(-0.063)	(0.516)	(0.543)	(1.713)
(Log) Household Consumption	0.159***	0.136***	0.076*	0.064
	(3.002)	(2.581)	(1.672)	(1.167)
Observations	24640	24238	26924	26766
Groups	18381	18039	18213	18463
R^2	0.001	0.060	0.051	0.000

* p<0.10, ** p<0.05, *** p<0.01. t-statistics (robust to clustering within pids) in parentheses.

income is the security it provides in this respect.

Preliminary support for this hypothesis can be found in Table 3.9, where significant coefficients on consumption for most of the samples in which 2005 data are included suggest that the driver is unlikely to be an event that occurred in 2005. Rather, the regressions displayed in Table 3.10 suggest that, as we suspect, it is most likely to be an event related to the 2000 wave. Not only is consumption insignificant in most of the samples including the year 2000, but the Wave 10 dummy is also negative and significant in the majority of them.

In an attempt to test the crash hypothesis more directly, we adopt two strategies. First, we divide the sample by response date, on the theory that the crash is likely to loom larger in the minds of those respondents who answered early in wave 10. Second, we divide the sample into investors and

Table 3.9: Life Satisfaction Income and Consumption Equations Using Data from 2005 with Various Other Waves

	Sample Years Included							
	96-05	97-05	98-05	99-05	00-05	02-05	03-05	04-05
(Log) Household Income	0.021	0.016	0.023	0.039**	0.030*	-0.003	0.012	0.032**
	(0.908)	(0.706)	(1.006)	(2.250)	(1.713)	(-0.200)	(0.988)	(2.125)
(Log) Household Consumption	0.127**	0.216***	0.156**	0.105**	0.064	0.138***	0.101*	0.050
	(2.021)	(3.869)	(2.447)	(2.054)	(1.167)	(2.682)	(1.883)	(0.879)
Wave 15	0.771	0.543	0.146	-0.154	0.372**	-0.100	-0.100	-0.095**
	(1.309)	(1.304)	(0.359)	(-0.746)	(2.315)	(-1.052)	(-1.306)	(-2.562)
Observations	21732	23701	23489	26779	26766	28051	27861	27257
Groups	16346	18259	17880	18465	18463	17053	16319	15236
R^2	0.001	0.000	0.000	0.127	0.000	0.038	0.003	0.121

* p<0.10, ** p<0.05, *** p<0.01. t-statistics (robust to clustering within pids) in parentheses.

Table 3.10: Life Satisfaction Income and Consumption Equations Using Data from 2000 with Various Other Waves

	Sample Years Included							
	96-00	97-00	98-00	99-00	00-02	00-03	00-04	00-05
(Log) Household Income	0.014	-0.016	-0.004	0.016	0.028*	0.026	0.014	0.030*
	(0.627)	(-0.803)	(-0.202)	(0.984)	(1.732)	(1.640)	(0.822)	(1.713)
(Log) Household Consumption	0.142***	0.258***	0.147***	0.024	0.068	0.034	-0.011	0.064
	(2.634)	(4.623)	(2.606)	(0.447)	(1.396)	(0.684)	(-0.235)	(1.167)
Wave 10	-0.573**	0.096	-0.258**	-0.060**	-0.108*	-0.088	-0.316**	-0.372**
	(-2.410)	(0.523)	(-2.350)	(-2.036)	(-1.734)	(-0.991)	(-2.552)	(-2.315)
Observations	21386	23355	23143	26433	27705	27515	26911	26766
Groups	14665	15399	14863	14999	18051	18282	18226	18463
R^2	0.020	0.000	0.036	0.071	0.000	0.053	0.000	0.000

* p<0.10, ** p<0.05, *** p<0.01. t-statistics (robust to clustering within pids) in parentheses.

non-investors, on the theory that those who invest are likely to respond more strongly to the crash than those who do not.

Table 3.11: The Effect of Wave 10 Response Date on SWB Regressions, 2000 & 2005

	(1)	(2)	(3)	(4)
(Log) Household Income	0.071**	0.010	0.065**	0.065**
	(2.234)	(0.488)	(2.066)	(2.102)
Income x Late Wave 10			-0.050	-0.050
			(-1.360)	(-1.385)
(Log) Household Consumption	-0.012	0.109*	0.002	
	(-0.127)	(1.705)	(0.021)	
Consumption x Late Wave 10			0.102	0.104*
			(0.959)	(1.682)
Early Wave 10	-0.439		-0.365**	-0.366**
	(-1.010)		(-2.239)	(-2.256)
Late Wave 10		-0.304*	-0.334**	-0.334**
		(-1.680)	(-1.976)	(-1.977)
Observations	8687	18079	26766	26766
Groups	5476	12987	18463	18463
R^2	0.002	0.000	0.000	0.000

* p<0.10, ** p<0.05, *** p<0.01. t-statistics (robust to clustering within pids) in parentheses.

Table 3.11 displays the results of regressions (using Waves 10 and 15 data) on the sub-sample that responded to the Wave 10 questionnaire in September 2000 (column (1)), on the sub-sample that responded in the remainder of Wave 10, between October 2000 and May 2001 (column (2)), and then with combined response samples, but including interaction terms for response date (September or Later) with income and consumption. As we can see comparing columns (1) and (2), it is essentially the September responders who are responsible for both the increase in the income coefficient, and the fall in the consumption coefficient (and its significance). Indeed, restricting our analysis to only late Wave 10 responders returns both coefficients to their previous magnitudes (although with predictable reductions in significance resulting from the reduced sample size). Turning to columns (3) and (4), we see essentially the same pattern when we include response date interaction terms on income and consumption: late response appears to increase the consumption coefficient and decrease the income estimate (although only the former interaction attains significance at the 10% level, and only if we restrict the effect of consumption to zero for early responders. Broadly similar patterns can be found by splitting the sample differently (including splitting it into more than two response categories), but the differences between September and later response dates appears to be the most stark. Taken together, this evidence suggests that the crash is indeed likely to be the driving force behind the anomalous results. In particular, Senik's (2004) argument that the future becomes more important relative to the present in times of high volatility, is consistent both with the fall in the consumption coefficient *and* with the marked increase in the income coefficient for the early responders.

Table 3.12 displays the results of regressions dividing the Wave 10 and 15 data into another two sub-samples: those without any investments in Wave 10 (column (1)), those holding such investments (column (2)), and a regression on the combined sub-samples with interaction terms between Wave 10 investor status and, respectively, income, consumption, and the Wave 10 dummy. The pattern of results is again consistent with our crash-related

Table 3.12: The Effect of Wave 10 Investor Status on SWB Regressions, 2000 & 2005

	(1)	(2)	(3)
(Log) Household Income	0.028	0.083**	0.034*
	(1.441)	(2.020)	(1.779)
Income x Wave 10 Investor			0.027
			(0.618)
(Log) Household Consumption	0.090	-0.016	0.109*
	(1.321)	(-0.182)	(1.645)
Consumption x Wave 10 Investor			-0.159
			(-1.539)
Wave 10	-0.461**	-0.137	-0.376**
	(-2.425)	(-0.454)	(-2.309)
Wave 10 x Wave 10 Investor			-0.040
			(-1.105)
Observations	16169	10251	26420
Groups	10477	7844	18321
R^2	0.001	0.000	0.001

* p<0.10, ** p<0.05, *** p<0.01. t-statistics (robust to clustering within pids) in parentheses.

volatility hypothesis, though the difficulties in obtaining significant coefficients means that the evidence they provide cannot be conclusive. Comparing columns (1) and (2) we see that the coefficients on income and consumption change in the way predicted, with the income coefficient higher and the consumption coefficient much lower (indeed negative, though insignificantly so) for investors than non-investors. Nonetheless, the consumption coefficient remains insignificant in the non-investing sample, and the interaction terms, while correctly signed, also fail to attain significance.

Altogether, the evidence is consistent with the conjecture that both the reduction in the consumption coefficient and the increase in the coefficient on income in the Waves 10 and 15 samples are anomalies generated by the stock market crash of 2000, and that they are is not representative of the effect of the two variables under more normal circumstances. Interestingly however, the pattern of the changes also provides some support to Senik's (2004) argument that, in volatile times, individuals may become more forward-oriented. In such a context, not only may income become more important as a hedge against future uncertainty, but also, in the context of a negative shock to wealth, high current consumption may in fact have a negative effect, particularly if individuals now perceive themselves to be over-consuming relative to what their new level of wealth can sustain.

A remaining concern of course, is that the crash may also affect the coefficients obtained in the wealth equations. In particular, the stated value of investments in 2000 may provide a poor measure of individuals' confidence in their ability to secure future consumption. For this reason, we attempt to check the robustness of the results reported in the All Wealth specifications above, by running regressions on the sub-sample of respondents who were interviewed for Wave 10 after September 2000 (on the theory that they appeared to be least affected by the crash). Unfortunately, the loss of observations meant that, while the point estimates were very similar to the main results, we were unable to obtain significant coefficients on any of the wealth variables. The similarity of the point estimates lends support to our previous conclusions, although there is a legitimate concern that the crash may have inflated the effect of savings and debt. On the other hand, the housing results set out earlier are unlikely to be greatly affected by the crash.

3.3.2 Why Does Wealth Matter to SWB?

The previous section presented evidence that wealth does appear to affect SWB independently of consumption and income. This section attempts to identify potential reasons for this effect, and to examine the possible implications of this. It suggests that wealth is important to SWB in Britain primarily because it is associated with anticipatory consumption, and that, as a result, the wealth types that will exert the most effect on SWB are those most closely linked to consumption of specific goods or services in the near future. It proceeds in three stages: first outlining the predictions of different theories about why wealth may matter to SWB, and assessing the evidence of the previous section against these predictions, before going on to test the idea that specific types of savings may have greater or lesser impact on SWB.

Distinguishing Alternative Explanations for the Importance of Wealth

Having determined that wealth is important to SWB, and *how* its effects manifest themselves, an obvious next question is *why* it should have such effects. Section 2.1 canvassed four potential channels through which wealth could potentially affect subjective well-being, independently of its contribution to (current) consumption: status; self-esteem; anticipation; and intrinsic or direct importance.

Identifying one or more plausible channels through which the effect of wealth might operate may be important for at least three reasons. First, identifying plausible *substantive* channels through which wealth affects SWB will buttress the argument that wealth does indeed matter, and that we are not merely picking up measurement error in consumption. Second, the reason(s) why wealth matters may suggest significant differences in the importance of different *types* of wealth. Finally, the channel(s) via which wealth affects SWB may be interesting independently of our specific interest in wealth: if, as we later argue, wealth is important because it has an effect on anticipatory consumption, then this will lend support to decision models that incorporate such concerns (see e.g. Loewenstein, 1987; Elster and Loewenstein, 1992; Caplin and Leahy, 2001; Köszegi, 2005), and will also suggest that *other* factors that affect anticipatory consumption are also likely to affect SWB.

The second reason just noted also suggests a way of potentially distinguishing between different explanations regarding the importance of wealth. The four theories of why wealth may matter to SWB each make slightly different predictions about which *types* of wealth, if any, are likely to have the greatest impact on SWB. A summary of the predictions of different theories can be found in Table 3.13, below. We expand upon these in detail in the following subsections, dealing with the predictions of the anticipatory consumption, status, self-esteem, and "intrinsic" theories in turn.

Table 3.13: Factors Affecting the Hedonic Impact of Different Wealth Types

Theory	**Relevant Factors**	**Affecting the Importance of...**
(1) Self-Esteem	None Systematic	N/A
(2) Status	Visibility (+)	Housing (+)
(3) Anticipatory	Temporal Proximity (+)	Home, Short-Term Savings, Debt (+)
Consumption	Ease of Imagining (+)	Home, Specific Savings (+)
	Uncertainty (-)	Investments (-)
(4) Direct	None Systematic	N/A

Anticipatory Consumption

The most detailed set of predictions is provided by the (Jevonian) anticipatory consumption explanation. If we assume, following the likes of Caplin and Leahy (2001) and Köszegi (2005), that current utility from anticipation, u_t^a is some function of the distribution of future consumption possibilities, $u_t^a = v(\Phi_t(c_{t+1}, \ldots, c_T))$, then the impact of a given wealth portfolio on utility from anticipation will depend the nature of the consumption possibilities it permits. In a decision context, Elster and Loewenstein (1992) argue that utility from the anticipation of future consumption depends (positively) on at least four attributes: (1) the amount and attractiveness of that consumption; (2) its temporal proximity; (3) the ease or vividness with which it can be imagined; and (4) the probability that it will occur. (1) just predicts that SWB will be increasing in wealth. However, as shown in Table 3.13, considerations (2)-(4) suggest a number of hypotheses about the types of wealth that are likely to be most or least important to SWB.

(2) suggests that wealth types most closely tied to consumption in the near future are likely to be more important to current SWB. In particular: (a) the value of one's home may be important given its obvious link to future housing consumption;[19] and (b) short-term, liquid assets, such as cash savings or short term debt, should have a greater impact than illiquid, or long-term assets such as pensions accounts. (3) suggests that savings for specific purposes (and in particular ones that are easy to imagine and likely to occur soon) may have a greater impact than savings held for no particular reason, or for purposes that are more difficult to imagine. Finally, (4) suggests that high risk assets will have lower SWB returns (for a given expected financial return). However, because such risk should be priced into the assets already, it is unlikely to be detectable in SWB regressions.

Status

The status theory of the importance of wealth claims that wealth is important because we care about others perceptions of us, and wealth makes them

[19] There is an interesting issue here as to whether we would expect changes in home value that arise solely due to price changes (rather than real improvements) to increase SWB. On the one hand, it does not seem to change individuals' future housing consumption possibilities in any real sense. On the other, if price increases relax credit constraints by allowing individuals to borrow more against the value of their home, then this may affect future consumption possibilities. We will not address these issues in any great detail, however.

accord us greater status. The fact that this theory is rooted in the perceptions of *others* suggests that the status we receive from wealth will depend in part on how visible that wealth is: more visible types of wealth should have greater impact on SWB than less visible wealth holdings (and, in the extreme, that "invisible" wealth should have no effect at all). We would expect this to manifest itself primarily in a greater importance being attached to housing, as the most visible of wealth types, in contrast to investments, savings, or debt.

Self-Esteem and Intrinsic Importance

In contrast, the remaining theories (based on self-esteem and intrinsic importance) seem to predict little in the way of systematic differences in the importance of different wealth types. The self-esteem theory, to the extent that it differs from the status-based account, essentially suggests that in the same way that current earnings may represent a "flow" measurement of self-worth, *accumulated* earnings (i.e wealth) may reflect the corresponding "stock". On this theory, it should not matter greatly to what purpose this stock is put, or in what form it is held, and major differences between different types of wealth would provide evidence against its being an important factor.

Similarly, on the intrinsic importance account, there is no particular reason to suspect that wealth held in particular form should be any more important than another. Consequently, this account does not predict any specific differences on this score. On the other hand, in contrast to the self-esteem account, the existence of such differences would not be inconsistent with the direct importance theory, so the presence of any such differences would be only weak evidence against it.

Wealth and Anticipatory Consumption

The broad pattern of the results in the previous section suggested that the most important wealth types from an SWB perspective are short-term savings and debt. We also found some evidence that the value of one's home and mortgage debt could significantly affect SWB, though this only occurred in the larger sample for which other wealth measures were unavailable. Investments, meanwhile, appear to have little discernible effect. Collectively, these results appear to provide support for the theory that wealth matters primarily because individuals care about their future consumption, rather than because it contributes to status (which would suggest only housing should matter), or self esteem (which would suggest that investments should matter too). However, because there is some doubt about both the effects of housing wealth (which becomes insignificant in the all wealth specification) and the non-effect of investments (which may be due the 2000 stock market crash), this evidence is potentially subject to doubt.

This section therefore attempts a further test of the anticipatory consumption theory, based on the hypotheses that future consumption opportunities are likely to have a greater hedonic impact the sooner they are expected to occur, and the more easily imaginable they are. We investigate, and find support for both of these hypotheses below, utilising survey information on the intended time-horizons and primary purpose of households' savings.

In each wave, respondents who reported saving were asked the following question:

> Would you say your savings are mainly long term savings for the future or mainly short term savings for things you need now and for unexpected events?

with possible responses falling into three categories: "mainly long term", "mainly short term", or "both equally". In addition they were also asked what they were saving for, with the (free) responses later being categorised into 12 main groups (plus a residual "other" category). The distribution of responses across both these questions (for Waves 10 and 15 only) is displayed in Figure 3.9.[20]

[20]As can be seen, many purposes tend to be associated with particular time horizons: savings for holidays, special events, cars, and home improvements tend to be focused on the short term, while savings for old age, house purchases, and children tends to be more long-term focused. In addition, we can also see that most individuals appear to save with no particular purpose in mind.

Figure 3.9: Long Term vs. Short Term Savings by Primary Savings Purpose, 2000 & 2005

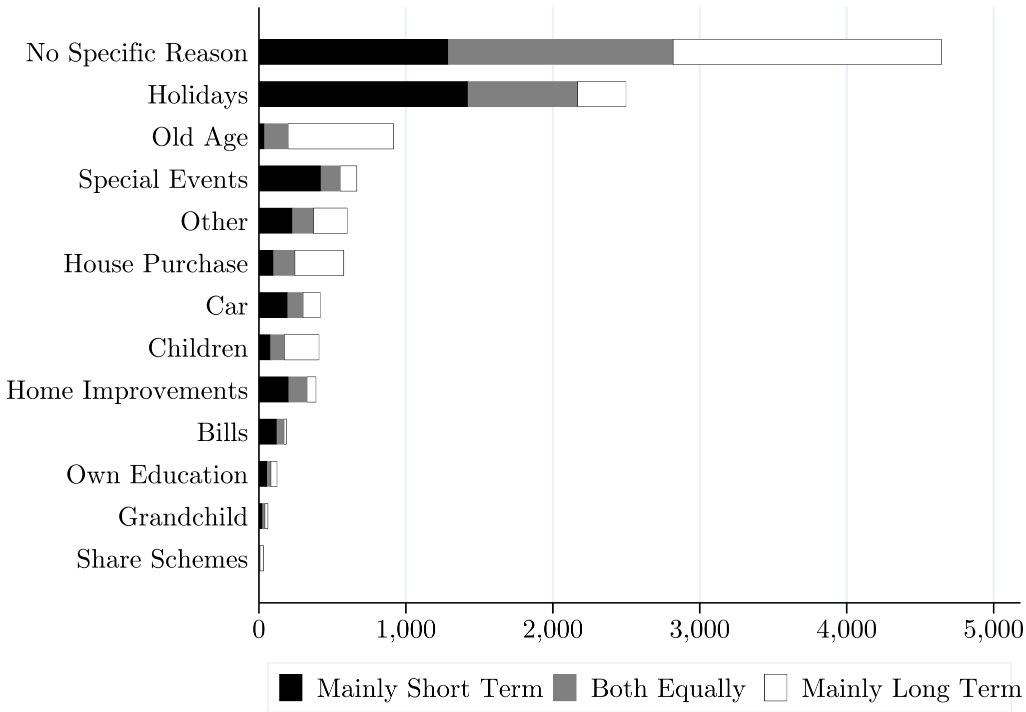

We test the time-horizon and specificity hypotheses by interacting responses to these questions with the original log savings variable (restricting attention to the sub-sample that reports saving). The results, presented in Table 3.14 appear to confirm both hypotheses. Columns (1) to (3) test the time-horizon hypothesis, by including interactions between savings and dummy variables for the categorical responses to the time horizon question. In column (1), "Mainly Long Term" is the omitted category, and coefficients on the remaining interaction terms (mixed, and short-term savings) have the expected signs. While the difference between mixed, and mainly long term savings is insignificant (and an F-test also indicates that the difference between mixed, and short term savings is insignificant), short term savings has a sizeably and significantly greater effect than long term savings. This is further confirmed in column (2), when we collapse the omitted category to include both mainly long term, and mixed long and short term responses: the coefficient on the short term savings interaction remains both substantively and statistically significant. Column (3) restricts the effects of long-term and mixed saving to be zero: and demonstrates that short term savings have a significantly positive effect.[21]

Columns (4) and (5) test the hypothesis that saving for specific, easily imagined items has a greater hedonic impact than savings for less easily imagined, or entirely non-specific purposes. In order to do so, we interact savings with a dummy variable that takes the value of 1 if the primary purpose of savings is for Holidays, a Car, a House, Home Improvements, or Special Events, and 0 otherwise (savings for Old Age, Children, Bills, Share Schemes, Own Education, Grandchildren, No Specific Reason, or Other). Both when we include the main effect of savings (column (4)), and when we constrain it to

[21] We also experimented with treating the categorical responses as a continuous time horizon variable, coded 1 for long term, 2 for mixed, and 3 for short term savings. Again, the results found a significantly larger impact of short term savings on SWB. However, the implicit restriction that the difference between long term and mixed savings is the same as the difference between mixed and short term savings appeared to make the model perform less well overall, and implied a negative, albeit insignificant, effect of mainly long term savings.

zero (column (5)), the interaction effect is sizeable and significant, supporting the hypothesis that savings for specific, easily imagined items has greater hedonic returns than non-specific savings. Moreover, the point estimates on the main savings effect suggests that any independent effect of long-term, or non-specific savings is indistinguishable from zero.

Of course, one possible concern with these results is that the two tests may simply be picking up the same effect. As we can see in Figure 3.9, savings for the specific items included in the dummy variable tend also (with the exception of savings for house purchases) to be primarily short term savings, while savings for many of the excluded purposes, such as old age, or children tend to be primarily long term savings. It is somewhat difficult to address this potential collinearity; indeed, part of the reason for expecting short term savings to have a greater impact is because we think it is in a sense more immediate, and therefore easier to imagine, so the hypotheses themselves are not intended to be entirely independent. Nonetheless, in column (6) we provide a (somewhat inconclusive) attempt to address this concern, separating out and estimating directly the coefficients on four distinct categories of savings: short-term specific, long-term specific, short-term non-specific, and long-term nonspecific. Though none of the individual coefficients is significant, an F-test rejects the hypothesis that the specific savings coefficients are equal to the non-specific savings coefficient. However, while the sign of the difference between long-term/mixed and short term-savings is still as predicted, an F-test cannot reject the hypothesis that the effects are identical, conditional on savings specificity. Although not conclusive, this suggests that the observed difference between the effects of different savings types may be due more to savings being for specific items, rather than their short-term nature.

Measurement Error

A final possibility is that the wealth coefficients are simply picking measurement error in the consumption variable, rather than an independent effect on SWB. This possibility is perhaps most plausible for the gross housing wealth variable, which is correlated with the consumption variable by construction.[22] On the other hand, there is no particular reason to suspect strong correlations between unmeasured consumption and changes in debt or savings — indeed, we would expect precisely the opposite: that, controlling for income, higher savings and lower debt should reduce rather than increase consumption. Nonetheless, to guard against the possibility that our results may be spurious, we run a series of regressions, displayed in Table 3.15, in which we instrument consumption with our wealth and income measures. If, conditional on consumption, the wealth variables are truly orthogonal to SWB,

[22]See Appendix A.4.

Table 3.14: The Effect of the Time Horizon and Purpose of Savings on Life Satisfaction in Britain, 2000 & 2005

	(1)	(2)	(3)	(4)	(5)	(6)
(Log) Household Income	0.031	0.031	0.031	0.044	0.044	0.041
	(0.736)	(0.755)	(0.753)	(1.055)	(1.051)	(0.956)
(Log) Household Consumption	0.125	0.126	0.125	0.135	0.135	0.121
	(1.303)	(1.308)	(1.302)	(1.372)	(1.366)	(1.217)
(Log) Gross Savings	-0.002	-0.001		-0.001		
	(-0.385)	(-0.189)		(-0.190)		
(Log) Savings x Savings Both Long and Short Term	0.003					
	(0.489)					
(Log) Savings x Savings Mainly Short Term	0.012*	0.011*	0.010*			
	(1.729)	(1.718)	(1.727)			
(Log) Savings x Saving for a Specific Item				0.011**	0.011**	
				(1.977)	(2.035)	
(Log) Short Term Non-Specific Savings						-0.003
						(-0.341)
(Log) Long Term Non-Specific Savings						-0.001
						(-0.165)
(Log) Short Term Savings for a Specific Item						0.008
						(0.970)
(Log) Long Term Savings for a Specific Item						0.009
						(1.318)
(Log) Non-Mortgage Debt	-0.010*	-0.010*	-0.010*	-0.010*	-0.010*	-0.010*
	(-1.670)	(-1.649)	(-1.653)	(-1.809)	(-1.817)	(-1.688)
Constant	4.556*	4.560*	4.570*	4.188	4.203	4.349
	(1.675)	(1.677)	(1.680)	(1.521)	(1.525)	(1.578)
Observations	10619	10619	10619	10561	10561	10413
Groups	8492	8492	8492	8454	8454	8351
R^2	0.002	0.002	0.002	0.004	0.003	0.005
AIC	8564.97	8564.02	8562.18	8399.39	8397.56	8237.29
BIC	8688.57	8680.35	8671.24	8515.63	8506.53	8367.80

* p<0.10, ** p<0.05, *** p<0.01. t-statistics (robust to clustering within pids) in parentheses. Standard controls (Age, Age Squared, Health Status, Marital Status dummies, Employment Status, Household Size, and a Wave 15 dummy) included but not reported. "Savings Mainly Short Term", "Savings Mainly Long Term", and "Savings Both Short and Long Term"are self-reported dummy variables. "Savings Time Horizon"is a continuous variable coded as 1 for "Mainly Long Term", 2 for "Both Short and Long Term", and 3 for "Mainly Short Term" Savings. "Saving for a Specific Item"is a dummy variable coded to 1 if the respondent cites the primary reason for savings as being for Holidays, a Car, a House, Home Improvements, or Special Events, and 0 otherwise (Old Age, Children, Bills, Share Schemes, Own Education, Grandchild, No Specific Reason, or Other). "(Log) Short Term Savings for a Specific Item "is defined as log savings if savings is mainly short term and saving is for a specific item...

and can be excluded from the SWB equation, then they would be valid instruments for consumption in these regressions (see 3.2.2). Conversely, if the Sargan-Hansen test rejects the null hypothesis that the instruments are jointly valid, then this lends support to the contention that they do indeed belong in the main equation.

As column (2) shows, the validity of the instruments is strongly rejected in the housing wealth specification (which was of greatest concern to us). Moreover, while the test statistics fail to reject in the all wealth specification (column (3)), two further pieces of evidence suggest that the wealth variables do indeed have separate effects. The first is that the first-stage regressions (unreported) suggest that our suspicions about the correlations between the wealth variables and consumption are partly warranted: in particular, deviations in consumption are negatively correlated with deviations in savings, meaning that the savings effect cannot be proxying for consumption. The second is that the increase of an order of magnitude in the instrumented consumption coefficients in column (3) seems too large to be plausible. Under the classical errors in variables assumptions, it would imply that measurement error constitutes around 90% of the variation in the consumption variable.[23]

[23]See e.g. 3.2.2 for the formula used to calculate this.

Table 3.15: SWB Regressions Instrumenting Consumption with Wealth Measures

	(1)	(2)	(3)	(4)
Instruments	None	Housing Wealth	Non-Housing Wealth	All Wealth
(Log) Household Consumption	0.110***	0.198***	1.094***	1.188***
	(6.195)	(3.575)	(3.267)	(3.258)
Observations	110744	91785	14916	13434
Groups	22918	18254	7458	6717
Sargan-Hansen Test		10.85	3.62	2.62
p-value		0.004	0.163	0.269

* p<0.10, ** p<0.05, *** p<0.01. t-statistics (robust to clustering within pids) in parentheses.

3.4 Conclusion

Following on from the analysis of the previous chapter, this chapter has attempted to determine whether the conclusions reached there — that the income-based regressions standard in applied SWB work significantly understate the importance of money to SWB; and that, contrary to the standard consumption-based utility function, income and wealth contribute independently to SWB — generalise beyond the potentially rather special case of Russia.

Although the British results do differ in some important ways from those of the previous chapter, these fundamental conclusions emerge largely unscathed. In Britain, income-only regressions understate the importance of money to SWB even more strongly than they do in Russia. In contrast to the 40%-80% underestimation in Chapter 2, here, income regressions give estimates that appear to be around seven times too small. The main reason for this result is that the effect of income in Britain, though still independently significant, is much lower than the corresponding consumption estimate (cf. Russia where the two were approximately equal).

A number of possible explanations for this result exist. The hypothesis that income is especially important in more volatile economic circumstances — and that Russia's particularly traumatic transition experience may therefore be responsible for its higher effect there — finds some support in the response of Britons to the 2000 stock market crash: during *this* period the effect of income increases between 3 and 4 fold (and the relative effect of consumption falls to a similar level to the income effect). Alternatively, the relative absence of credit constraints in Britain, as compared with Russia, may serve to make current income less important to future consumption. Finally, we also found evidence in the Russian case that part of the effect of income was due to its effect on status, and it is also possible that cultural differences in this sphere could be responsible for part of the difference in results.[24]

These results temper our previous conclusions regarding the implications

[24]In contrast to the Russian results, where we found that earned income was more important to SWB than unearned income, a similar (unreported) test on the BHPS sample suggests that this does not play a role in Britain.

of the importance of income. In contrast to the Russian case, where focusing exclusively on either consumption-based or income-based measures would significantly understate well-being, in Britain, the effect of ignoring income would be minimal, lending support to claims that consumption should be the preferred informational base for such analysis, at least in some contexts. On the other hand the argument we made in Section 2.7 — that the presence of income as an independent argument in the utility function increases the welfare costs of income volatility — still seems relevant. While the analysis in this chapter has shown that income is relatively unimportant to SWB in Britain in "normal" times, we also saw its importance increase during the stock market crash — that is, income became more important precisely as volatility increased.

In contrast to the smaller coefficient on income, the results of Section 3.3 suggest that wealth is similarly important to SWB in Britain and Russia, and that it is important primarily because individuals care about their future consumption possibilities. As hypothesised by Elster and Loewenstein (1992), wealth that is linked to future consumption that is easily imaginable and likely to occur soon seems to be especially important — with short-term savings and savings for specific items playing a particularly strong role, and debt and housing also having potentially significant effects. In contrast, longer term investments appear to affect current SWB very little if at all.

These results lend further support to the Jevonian perspective recently formalised by the likes of Caplin and Leahy (2001) and Köszegi (2005), which suggests that the motivations underlying individuals' decisions to save, consume and/or take on debt are not exclusively determined by a detached weighing of the interests of their current and future selves (as the standard discounted utility model might suggest) but may instead be attributed, at least in part, to the effect of the future on current utility. This finding also comports with Senik's (2004) analysis of the RLMS. However, while she found a connection between *others'* income and SWB, via its impact on individuals' expectations about their own future income, our results suggest that it may instead be future *consumption* that matters to SWB, and that both one's current wealth and (to a lesser extent) income may provide information about this.

On the other hand, in contrast to these models, the evidence presented here suggests that current SWB is not simply a direct function of a probability distribution over our future consumption possibilities, $u_{it}^{anticipation} = v(\Phi_t(c_{t+1}, \ldots, c_T))$. Instead, we find support for an experienced utility version of the mental accounting model proposed by Thaler (1980): SWB appears to be affected not so much by households' net wealth, as *separately* by the wealth they hold in different forms. The impact of housing wealth is, for example, separate from the impact of savings wealth, and the impact of assets is separate from the effect of debts. The latter in particular seems inconsistent with the idea that future consumption matters *directly*, and that wealth is merely proxying for it. The sub-utility function implied by the mental account model ($u_{it}^{anticipation} = \beta_{savings} \ln a_{savings} - \beta_{debt} \ln a_{debt}$) does not appear to correspond to any natural form of $u_{it}^{anticipation} = v(\Phi_t(c_{t+1}, \ldots, c_T))$.

In addition to its substantive implications, the finding that the SWB function appears to have a mental account structure also serves to increase our confidence in the wealth estimates obtained from the RLMS (where the absence of detailed wealth data forced us to estimate the effects of different wealth types separately). Indeed, it suggests that attempting to estimate the effects of overall net wealth may actually have given us *worse* estimates: after all, in the BHPS it would have given the misleading impression that wealth did not matter to SWB at all.

Finally, in contrast to previous results obtained in a behavioural context by Thaler and Johnson (1990), we find little evidence that individuals frame their wealth holdings so as to maximise their SWB. This suggests that debts can have a substantial negative effect on SWB even for individuals with sizeable positive wealth holdings. On the other hand, and in contrast to the results of Brown, Taylor, and Wheatley Price (2005), we find no evidence that debt has a *larger* effect on experienced SWB than positive wealth. This difference may be due to their failure to control for unobserved personality traits, or simply to our use of slightly different SWB measures (in contrast to life-satisfaction, the General Health Questionnaire measure they use seems to have a generally more more negative focus). Consequently, to the extent that individuals *act* in a manner that suggests they are especially averse to debt, it may be that (as Kermer et al., 2006, suggest) such behaviour is due to a "forecasting errors", whereby individuals overestimate the negative effect of debt in their decision-making.

Taken together the findings here reinforce the conclusions of the previous chapter, that accounting for distinct effects of consumption, income and wealth on SWB is important from both a methodological, and a substantive perspective. In line with results of Senik (2004), and recent modelling work in behavioural economics, they seem to emphasise in particular the importance of the future to current utility: the value of current economic outcomes is related in important ways to the signals they provide about individuals' prospects going forward, and the hopes and fears people have for the future.

Part II

Reference Groups

Chapter 4

Getting Good Estimates of Social Effects

A central issue in empirical happiness research involves the identification of social or reference effects: the effect of others' outcomes on individual well-being. The standard approach to estimating such effects is to define a reference level of some outcome variable — typically as either the cell mean for a given individual (e.g. Ferrer-i-Carbonell, 2005; Kingdon and Knight, 2007), or as the predicted value from a wage equation (e.g Clark and Oswald, 1996; Senik, 2004) — and then to estimate the coefficient on this reference level in an SWB regression.

As a number of authors have noted (see e.g. Clark, Frijters, and Shields, 2008; Clark and Senik, 2008; Senik, 2009), one shortcoming of these approaches is that they require the researcher to impose a particular conception of reference group. If the reference group is misspecified, this may introduce noise or bias into measures of the reference level, and is likely to give misleading results.[1] Table 4.1 presents an overview of the sorts of variables that have been used to demarcate the reference group under both the cell mean and wage equation approaches — and gives a sense of the variety of ways in which reference levels have been defined by different authors.

As we might expect, the cell mean approaches have tended to use many fewer dimensions to define the reference group: Ferrer-i-Carbonell's (2005) use of four distinct dimensions (sex, age, education, and region) is the most of any of the cell-mean approaches, and many use only a single dimension. In contrast, Senik (2004), who focuses on the fewest dimensions of all the wage equation specifications, still uses six. Common reference variables across both types of approach include geographic area (11 studies), sex (6), education (6), age (4), and occupation (4).

One approach to better defining the reference group is simply to ask people to whom they compare themselves (see Melenberg, 1992; Knight, Song, and Gunatilaka, 2009; Clark and Senik, 2008; Senik, 2009). Knight, Song,

[1] A natural conclusion would be to assume that this will result in attenuation of the reference coefficients, in line with standard errors-in-variables arguments. However, as we shall see below, there are reasons to think that the outcomes of some reference groups could have countervailing positive and negative effects (e.g. due to information or altruism), so it is not necessarily clear which direction such "attenuation" bias would operate in.

Table 4.1: Overview of reference levels and reference group specifications

Paper	Approach	Variables defining reference group
Cappelli and Sherer (1988)	Cell Mean	Occupation
Clark and Oswald (1996)	Wage Equation	Age, Sex, Health, Region, Occupation, Industry, Education, Accident, Hours, Establishment size, Sex mix at work, Organisation type, Marital status, Temporary contract, Incentive pay, Part-time, Union membership/recognition, Supervisor, Pension, Tenure
	Cell Mean	Hours
Sloane and Williams (2000)	Wage Equation	Sex, Marital Status, Education, Experience, Full time, Temporary contract, Hours, Supervisor, Incentive pay, Union membership/recognition, Establishment size
McBride (2001)	Cell Mean	Age
Senik (2004)	Wage Equation	Education, Sex, Experience, Profession, Industry, Region
Lévy-Garboua and Montmarquette (2004)	Wage Equation	Age, Sex, Education, Experience, Nationality, Weeks worked, Part-time, Professional, Specialized work, Bilingual, Region, Marital status, Religion, Health, Leisure
Blanchflower and Oswald (2004)	Cell Mean	State
Ferrer-i-Carbonell (2005)	Cell Mean	Sex, Age, Education, Region
Luttmer (2005)	Cell Mean	Locality
Helliwell and Huang (2005)	Cell Mean	Census tract
Graham and Felton (2006)	Cell Mean	Locality
Kingdon and Knight (2007)	Cell Mean	Location, Race
Senik (2008)	Wage Equation	Sex, Education, Experience, Occupation, Industry, Region, Full-time
Knight, Song, and Gunatilaka (2009)	Cell Mean	Village

and Gunatilaka (2009) find that most respondents in rural China self-define their reference group as comprising fellow-villagers or neighbours; Clark and Senik (2008) find that individuals are most likely to compare themselves to colleagues, rather than to friends (though those who spend more time socializing are more likely to compare to the latter); while Senik (2009) finds evidence that most external comparisons tend to be local (to colleagues, parents, school friends) rather than global (i.e. to the country as a whole).

Unfortunately, most data sets (especially long-running panel data sets) do not include such questions, and thus rule out such a strategy. In this chapter we pursue an alternative, and to our knowledge, novel approach to defining the reference group. In Section 4.1 we develop a data-driven estimation strategy that allows us to estimate jointly both the reference group and reference effects. This has the advantage both of allowing us to identify relevant reference groups even when explicit information about social comparison groups is not available; and of providing a potential check on the results of self-reported comparison groups in the studies above. It also opens up the possibility of allowing for "fuzzier" reference groups than can easily be elicited from self-reports: e.g. by including neighbours, colleagues, and similar others to varying degrees, rather than in an all-or-nothing fashion.

As an illustration of both the potential importance of appropriately defining the reference group, and of the potential of our estimation strategy in this context, Section 4.2 revisits Senik's (2004) finding that, in contrast to many previous results, others' income appears to have a *positive* rather than a negative effect on SWB in Russia. (This result has since emerged in other contexts, including Eastern Europe more generally (Senik, 2008), the US (Senik, 2008), Denmark (Clark, Kristensen, and Westergaard-Nielsen, 2007), and South Africa (Kingdon and Knight, 2007).)

We first demonstrate that allowing for the possibility that individuals might use both others' income and others' *consumption* as reference points raises questions about the net effect of changes in others' circumstances on individual SWB in Russia: while others' income has a consistently positive (informational) effect, others' consumption has a consistently negative (comparison) effect, and the net effect of the two appears to depend crucially on the definition of the reference group. When we attempt to estimate the reference groups for income and consumption, we find that, contrary to Senik (2004), that comparison significantly dominates information overall.

Our estimations also serve to shed light on the nature of the reference groups themselves. The reference groups for consumption and income differ in important ways, with productive characteristics being particularly relevant to the income reference group, but having little effect on the reference group for consumption. The results generally confirm the standard empirical assumption that individuals tend to focus on the outcomes of similar others.

However, there is suggestive evidence that this will not always be the case: individuals may take the income of those older than them as a signal, and may sometimes compare themselves to dissimilar rather than similar others. Our results also throw some doubt on the wage-equation approach to defining reference levels: it is clearly not the case that all variables that reliably predict income will be relevant to defining the reference group.

4.1 Estimating the Reference Group

In attempting to determine what definition of the reference group would best explain variations in individuals' SWB, one way to proceed would be simply to estimate a set of statistical models where the researcher imposes alternative definitions of reference effects in each. We could then choose between these using standard model selection criteria (such as information criteria). However, such an approach allows us to compare only a finite set of discretely specified models, which may not capture the full range of relevant variation in the model space. In contrast, the strategy we develop below allows us to attack the question of the appropriate reference group more directly, and to do so within the context of a single parameterised model.

Of course, we cannot get by without making some assumptions about the process that generates social effects. Broadly in line with the existing approaches, we assume that the relevance of others' outcomes to SWB depends on how similar two individuals are — where we interpret "similarity" broadly, to include geographic proximity as well as more abstract concepts of "social distance" such as education, age, family size, and job status.[2]

Much of the existing literature assumes that individual SWB is affected primarily by the outcomes of similar others: the more similar two individuals are, the more weight each will assume in the other's utility function, whether via social comparison, altruism, or information effects. While this does not seem unreasonable, there may sometimes be reasons to think that *dissimilar* others' outcomes could be more important to the reference level: for example, we might care more about our "team" doing well relative to others, than our own position within the team; or we might focus on dissimilar others as sources of information if we anticipate that our own status may change in the near future. Thus, although we initially motivate our approach with reference to the standard "like-affects-like" assumption, the estimation strategy we eventually employ will allow for more general patterns of dependence.

In contrast to the standard approaches in the SWB literature, we introduce flexibility in the weight individuals may give to others' outcomes in

[2]The concept of "social distance" is a familiar one in sociology, though economists have also utilized the notion: e.g. Akerlof (1997) develops a theoretical model of social decision-making, where decisions depend in part on the decisions of similar others; while Griliches (1992) and Hautsch and Klotz (2003) employ abstract notions of technological distance in investigating the presence of R&D spillovers and innovation decisions.

two stages. First, we develop a parameterised social distance function that allows us to estimate the social distance individuals perceive between each other; second, we embed this in a kernel weighting function that determines the influence of a given individual's outcomes as a function of that social distance.

The resulting composite function is indexed by a set of λ parameters, which describe how quickly the importance of others' outcomes decays as they become more dissimilar in a given dimension (e.g. place of residence, employment status, etc.). The cell-mean model then corresponds to the special case when $\lambda \to \infty$; from a predicted wage perspective, our approach essentially replaces the standard linear predictor with a non-parametric kernel predictor, with the λ parameters defining the kernel bandwidth. Normally, we would attempt to optimize this bandwidth for its ability to predict wages (out-of-sample); here, the bandwidth is instead estimated so as to optimize its ability to predict *reference levels*.

While a variable will not be able to predict reference levels unless it also predicts wages, the converse is not true. 'The' reference level could reflect one or more of the following: what an individual expects to get in the future; what they think they deserve; what they feel they need to be a success; or simply what others around them have. There is no particular reason that any of these should perfectly coincide with the individual's actual wage, and consequently, no particular reason to think that a better wage predictor should necessarily result in a better approximation to the reference level. To take an example, we might consider productive characteristics — such as education, experience, and hours worked — relevant to what we "deserve", but other characteristics — such as race, gender, or marital status — irrelevant. Even though the latter factors might better predict one's *actual* wage, they would still be irrelevant to the reference level. Similarly, the earnings of those with longer work experience than us might be most relevant to what we will earn in the future, even though the earnings of those with similar experience would better predict our current income.

To our knowledge, this is the first attempt in the literature to estimate reference groups directly from SWB data — or indeed to directly estimate reference groups in any context. The idea of embedding a distance function in a weighting kernel is familiar from machine learning and other applications of kernel and locally-weighted regression (see e.g. Atkeson, Moore, and Schaal, 1997, for a good overview). Indeed, in the closest precursor to our estimator, Hautsch and Klotz (2003) employ a similar distance-function-embedded-in-kernel strategy to analyse neighbourhood influences on innovation decisions. However, while they allow the scale parameter of the kernel to vary, and thus attempt to estimate how quickly influence declines with distance, they still impose a predefined notion of that distance (the Mahalanobis distance).

In contrast, our approach allows us to estimate the parameters of both the kernel and the distance function jointly, and thus allows the data to tell us what dimensions of distance are most important.

The remainder of this section describes our model and estimation strategy in more detail. Section 4.1.1 motivates and sets up the basic model in a cross-sectional context. Sections 4.1.2 and 4.1.3 extend the model to account for individual-level fixed effects and cross-group comparisons respectively. Section 4.1.4 describes details of the estimation and implementation via (penalized) maximum likelihood.

4.1.1 The Basic Model

Our basic model is a generalisation of the cell-mean approach to defining the reference level. However, rather than taking the reference level as a simple average of some outcome (e.g. income, consumption, etc.) across a pre-specified group of similar others, we instead define it as a *weighted* average of all others' outcomes across the entire population. We assume that the most weight will be given to the outcomes of those who are identical across a range of characteristics (e.g. who live in the same area, are the same age, have the same amount of education etc.) and that this weight will decline the more dissimilar the individuals become. We then attempt to estimate how quickly this decline occurs for each characteristic: for example, the importance of another person's consumption might decline quickly with geographic distance, but more slowly with age and/or education. A given set of these rates-of-decline (which we denote in the model by a vector of λ parameters) will define a unique reference level for each individual in each period, and therefore allow us to estimate the effect of changes in reference levels on SWB in the usual way.

Estimating the effect of others' outcomes on SWB can be thought of as attempting to estimate the function $s(\cdot)$ in an equation of the following form (omitting time subscripts for convenience[3]):

$$u_i = x_i'\beta + s(M_{-i}) + \varepsilon_i \tag{4.1}$$

where M_{-i} is a matrix of of the relevant outcomes $(m_1', m_2', \ldots, m_N')'$ for all individuals in the population except i.[4] Our development here will focus on the case where m is scalar, but the extension to the vector case is straightforward.

We begin by assuming a weighted additive form for $s(\cdot) = \sum_{j \neq i} \rho_{ij} m_j$.

[3]We reintroduce the time dimension explicitly below, when we allow for individual-specific fixed effects.

[4]In the eventual implementation, we shall sometimes define this vector at the household level to include all households other than individual i's household, but the development here abstracts from this detail for simplicity of exposition.

In this formulation, ρ_{ij} parameterises the effect of individual j's outcome(s) m_j on individual i's utility. The two standard approaches to defining the reference level can then be interpreted as decomposing ρ_{ij} into two terms as $\rho_{ij} = \rho w_{ij}$: the constant, ρ represents the degree of comparison to an (individual-specific) reference level, $\bar{m}_i$; while the cross-individual weight, w_{ij} when multiplied by m_j, represents individual j's contribution to determining that reference level. In line with the discussion above, w_{ij} will typically be a *similarity measure*, $\text{sim}(z_i, z_j)$ that weights m_j according to the similarity of (or social distance between) i and j, in terms of a set of variables z_i and z_j, which as already noted, could include e.g. geographic location, education, employment status etc.

The cell mean approach is then equivalent to defining w_{ij} as:

$$w_{ij} = \frac{\omega_{ij}}{\sum_{k \neq i} \omega_{ik}} \tag{4.2}$$

with ω_{ij} as a binary indicator function equal to one if individuals i and j are deemed relevantly similar:

$$\omega_{ij} = \begin{cases} 1 & \text{if } z_j = z_i \\ 0 & \text{otherwise} \end{cases} \tag{4.3}$$

for some p-dimensional z chosen by the investigator (usually composed of one or more discrete variables, such as area of residence, education level, age category etc.).[5] More generally (i.e. with alternative specifications of ω_{ij}), equation (4.2), which essentially restricts the weights w_{ij} to sum to one, describes the same class of similarity measures as the Nadaraya-Watson kernel (Nadaraya, 1964; Watson, 1964) often used in kernel estimation, and the implied reference level $\bar{m}_i$, can be thought of as a kernel predictor of $\hat{m}_i$ conditional on z_i.

In contrast, the linear regression approach to estimating the reference level is equivalent to defining w_{ij} as:[6]

$$w_{ij} = z_i'(Z'Z)^{-1}z_j \tag{4.4}$$

(which reduces to the cell mean approach described by equations (4.2) and (4.3) when the components of z_i are a set of mutually exclusive dummy variables). It is perhaps less natural to think of the linear regression approach in terms of a similarity measure. Nonetheless, (4.4) is actually a Mahalanobis angle measure, which in general defines the similarity between two vectors x and y as:

$$\mathrm{sim}(x, y) = x'\Sigma^{-1}y \tag{4.5}$$

[5] In similar fashion, a cell median approach would use a weight function equal to one only for the median observation(s) in the set of relevantly similar individuals.

[6] This follows straightforwardly from the equation for the predicted value of x: $\hat{x}_i = z_i'\hat{\theta} = z_i'(Z'Z)^{-1}Z'X$.

where Σ is the covariance matrix of the distribution from which x and y are drawn. Geometrically this is equivalent to the cosine of the angle between the vectors x and y, in the space induced by the Σ^{-1} transform.[7]

As already noted, the major shortcoming of both these approaches is that they impose a pre-specified weighting scheme. This has the benefit of allowing easy identification of ρ, but at the cost of assuming the appropriateness of the specified weightings (and the reference group that they imply). An obvious way of extending the existing cell-mean approach is to retain the Nadaraya-Watson kernel, but to use a more flexible distance measure in defining the ω_{ij}. This maintains the assumption that the the reference level $\bar{m}_i$ is a weighted average of the m_{-i}, with the weights depending on the perceived social distance between i and j, as defined by the distance function $d(z_i, z_j)$.[8]

A natural choice for a such a distance function is a weighted Euclidean

[7](In this case, the angles are defined relative to the origin. However, the Mahalanobis angle is more commonly defined relative to the centroid of the data space $\text{sim}(x, y, \bar{x}) = (x - \bar{x})'\Sigma^{-1}(y - \bar{x})$.

[8]Interestingly, once we extend the "kernel prediction" approach beyond the case of categorical z, the familiar boundary problem (whereby the kernel estimates tend to predict poorly at the edges of the distribution, because the bandwidth becomes asymmetric) takes on a potentially substantive meaning. If individuals compare themselves only to actually existing others, the boundary problem is not a problem at all, because the reference level is not intended to be a *prediction* of individual i's outcome, merely an appropriately weighted average of the outcomes of others. If instead, we think aspiration levels may be based on potentially hypothetical others, or directly on the outcome individual i might *expect* given her attributes, then it may be more of an issue (though only at the edges of the distribution). Unfortunately, a formulation sufficiently flexible to allow for both possibilities seems difficult to obtain, and we err on the side of simplicity.

distance:

$$d(z_i, z_j) = \left(\sum_{k=1}^{p} \lambda_k z_{kij}^{*2} \right)^{\frac{1}{2}} \tag{4.6}$$

where z_{kij}^* is the distance between the kth elements of z_i and z_j, defined as either: $|z_{ki} - z_{kj}|$ for standard continuous or binary variables; or for categorical variables, as the (z_{ki}, z_{ki})-th element of a symmetric distance matrix D_k, that is specified by the investigator so as to reflect differences between the categories (e.g. if the categories were geographical regions, D_k could be a contiguity matrix, with the (p, q)-th element equal to 1 if regions p and q are contiguous, and zero otherwise.) The weighted distance then, is just standard Euclidean distance between z_i and z_j, but with each dimension k allowed to vary in importance according to the scaling parameter λ_k.[9,10]

[9] For the moment, it will only make sense for λ_k to be non-negative. However, we relax this restriction in Section 4.1.3 below.

[10] The Euclidean distance measure could potentially be generalised in at least two ways. The first would involve generalising from the 2-norm to the ν-norm:

$$d(z_i, z_j) = \left(\sum_{k=1}^{p} \lambda_k z_{kij}^{*\nu} \right)^{\frac{1}{\nu}} \tag{4.7}$$

However, experience in the machine learning field suggests that the λ weights have much more of an impact on performance than does ν, and our (limited and informal) simulations appeared to confirm this.

The second generalisation would be to allow for interactions between distances in different dimensions:

$$\begin{aligned} d(z_i, z_j) &= \left(z_{ij}^{*\prime} M' M z_{ij}^* \right)^{\frac{1}{2}} \\ &= \left(\sum_{k=1}^{p} \lambda_k z_{kij}^{*2} + 2 \sum_{k=2}^{p} \sum_{l=1}^{k-1} \lambda_{kl} z_{kij}^* z_{lij}^* \right)^{\frac{1}{2}} \end{aligned} \tag{4.8}$$

(where λ_k is the kth diagonal and λ_{kl} is the (k, l)-th element of $\Lambda = M'M$).

A special case of (4.8), employed by Hautsch and Klotz (2003) is the Mahalanobis

To turn the distance function in (4.6) into a weight, we simply embed it in a kernel, such as the following (modified) Gaussian kernel:[11,12]

$$\omega_{ij} = \exp\left\{-d(z_i, z_j)^2\right\} \tag{4.10}$$

(which embodies the substantive assumption that a given unit change in distance will result in a given *proportional* change in weight). Combining

distance, which sets $M'M$ equal to the inverse of the covariance matrix of the Z variables, $N(Z'Z)^{-1}$. Hautsch and Klotz (2003) argue that the Euclidean distance is inappropriate in this context, because it assumes the dimensions of z are uncorrelated and have equal variances. In contrast, we view the choice of distance metric as a substantive question, rather than a merely statistical matter of correlations and variances — indeed, this is part of the motivation for our approach: the fact that two potential z variables have equal variances does not mean that they should contribute equally to social distance, as would be implied by the use of a fixed Mahalanobis distance; nor does it seem that just because two variables happen to be positively correlated, the weight of their interaction (λ_{kl} in the second line of equation (4.8) should be reduced (or negative).

Ideally, we would estimate the additional λ_{kl} parameters explicitly; but in what follows, we are content to assume they are zero to keep things relatively simple, and avoid having to estimate the $\frac{p(p+1)}{2}$ extra coefficients entailed by equation (4.8). If we wished, it would be straightforward to incorporate concerns about correlations into our model by pre-multiplying the z variables by M (where M is defined as above to be the Cholesky factor of $N(Z'Z)^{-1}$) and then estimating the λ_k weights in the transformed space, though this would somewhat complicate interpretation of the parameters in the original space.

[11] The full Gaussian kernel would be defined as:

$$\omega_{ij} = (2\pi)^{-\frac{1}{2}} \exp\left(-d(x_i, x_j)^2/2\kappa^2\right) \tag{4.9}$$

where κ is the kernel width, and controls the scale of the kernel. However, because the scale of $d(z_i, z_j)$ is arbitrary at this point (depending on the scale of the λ_k), we can normalise $2\kappa^2$ to one without loss of generality. In the Nadarya-Watson framework, the $(2\pi)^{\frac{1}{2}}$ normalisation will cancel (as it will appear in both the numerator and the denominator of equation (4.2)), so we also omit it in (4.10).

[12] Alternative kernels are of course possible. However, the Gaussian is probably most convenient to estimate. Unlike other popular choices employed in kernel estimation, it has the advantage of being continuous over $-\infty < d(z_i, z_j) < \infty$ (cf. the Epanechikov and tricube, which abruptly reach zero for $d(z_i, z_j) \geq \kappa$, and would therefore require more complicated optimization methods). The exponential form is also useful for reasons that will become apparent in section 4.1.3 below.

(4.1), (4.2), (4.6), and (4.10) then yields the following equation, which we can estimate by maximum likelihood (e.g. assuming normally distributed errors).

$$u_i = x_i'\beta + \rho \sum_{j \neq i} \frac{\exp\left(-\sum_{k=1}^{p} \lambda_k z_{kij}^{*2}\right)}{\sum_{l \neq i} \exp\left(-\sum_{k=1}^{p} \lambda_k z_{kil}^{*2}\right)} m_j + \varepsilon_i \tag{4.11}$$

The substantive interpretation of the parameters is that the individual λ_ks represent the relative importance of the corresponding z_k variables in determining social distance — the higher is λ_k, the lower the weight of "distant" others in determining individual i's reference level — while ρ reflects the impact of the reference level on SWB in the usual way.[13]

4.1.2 Fixed Effects

Extending this formulation to the fixed effects case is relatively straightforward. The general model is now:

$$u_{it} = x_{it}'\beta + s(M_{-i,t}) + \eta_i + \varepsilon_{it} \tag{4.12}$$

[13] As already noted, the extension to the case where the m_j are vectors is conceptually straightforward, the main substantive issue is whether or not we should then allow the λ_k parameters to vary between the dimensions of m — that is, whether we assume that the same social distance weighting applies in respect of different outcomes — and this is in any event a testable restriction.

As in the standard approach, the fixed effect η_i can be removed by subtracting individual means of each term:

$$\bar{u_i} = \bar{x_i}'\beta + \bar{s}(M_{-i,t}) + \eta_i + \bar{\varepsilon}_i \tag{4.13}$$

where $\bar{s}(M_{-i,t}) = \frac{1}{T_i}\sum_{\tau=1}^{T_i} s(M_{-i,\tau})$. This gives:

$$\tilde{u_i} = \tilde{x_i}'\beta + \tilde{s}(M_{-i,t}) + \tilde{\varepsilon}_i \tag{4.14}$$

with $\tilde{s}(M_{-i,t}) = s(M_{-i,t}) - \bar{s}(M_{-i,t})$. With $s(\cdot)$ defined as in equation (4.11), this yields the full equation:

$$\tilde{u_i} = \tilde{x_i}'\beta + \rho \sum_{j\neq i}\left\{\frac{e^{-z^*_{ijt}{}'\lambda}}{\sum_{l\neq i} e^{-z^*_{ilt}{}'\lambda}} m_{jt} - \frac{1}{T_i}\sum_{\tau=1}^{T_i}\frac{e^{-z^*_{ij\tau}{}'\lambda}}{\sum_{l\neq i} e^{-z^*_{il\tau}{}'\lambda}} m_{j\tau}\right\} + \tilde{\varepsilon}_i \tag{4.15}$$

(with z_{ijt} and λ now p-vectors). Again, it is straightforward to estimate the main parameters of this model via maximum likelihood.

4.1.3 Cross-Group Comparisons

It is possible that the assumption made above (and in the approaches that the strategy above is intended to generalize) that we care less about the outcomes of those who differ from us is too simplistic. In fact, given the widespread tendency of individuals to favour in-groups over out-groups, it would not be entirely surprising to find that individual make comparisons *across* groups at least as much as, if not more than, within them.[14] For example, individuals might care more that their family, nationality, or ethnic group does well relative to other groups, than they do about their standing within the group. One way to conceptualise this would be to think of individuals as simultaneously both altruistic towards *and* envious of those similar to them: if altruism declines with social distance more slowly than envy, then they would end up comparing more to dissimilar than similar others. An informational perspective could also lead to a focus on others who are dissimilar in some respects (say, those who are older or with more experience). As this suggests, there could be some dimensions of social distance where the weights decay with distance, and some where they decay with similarity. As a example, consider that, even if individuals tend to compare primarily to *socially*-dissimilar others, their awareness of others' outcomes may well decline with *geographic* distance.

In fact, if we allow the λ_k parameters to take on negative, as well as positive values, the framework set out above already allows for this possibility — though making sense of it requires a slight reconceptualisation of

[14]Indeed, it is also possible that this tendency could be endogenous, with lower-status group-members attempting to enhance their self-image by comparing their group's outcomes to those of other groups.

our set-up. Once we allow different elements of social distance to have counteracting effects, it no longer makes a great deal of sense to think about a single measure of social distance between two individuals, encapsulated in a distance function such as (4.6). Instead, we measure distance in each dimension separately, and embed each of the distances in its own kernel, to give k dimension-specific weights:

$$\omega_{kij} = e^{-\lambda_k z_{kij}^{*2}} \tag{4.16}$$

which we then combine to obtain the overall weight:

$$\omega_{ij} = \prod_{k=1}^{K} \omega_{kij} \tag{4.17}$$

The dimension-specific weights described by equation (4.16) are depicted graphically in Figure 4.1. Except for the extended parameter space, the weighting function that results from combining equations (4.16) and (4.17) is identical to that derived in section 4.1.1 from the combination of equations (4.2), (4.6) and (4.10): the only difference lies in the interpretation.

While (4.16) and (4.17) abandon the concept of a single "social distance" between two individuals, they nonetheless retain the intuitive interpretation of the λ_ks as indexing the degree to which distance in dimension k affects

Figure 4.1: ω_{kij} as a function of λ_k and z^*_{kij}

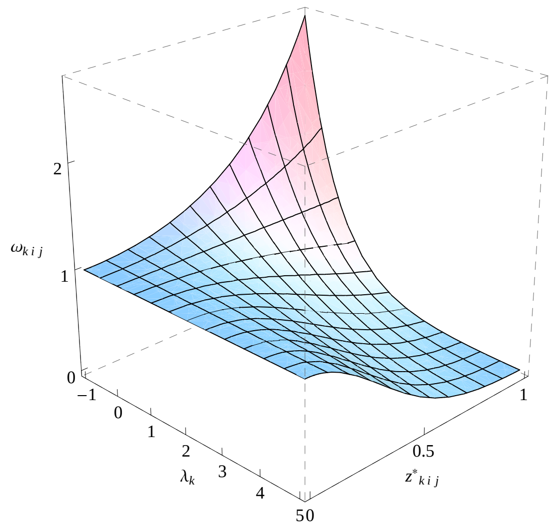

The figure shows the effect of increasing the social distance between two individuals (z^*_{kij}) for different values of λ_k, relative to identical individuals ($z_{kij} = 0$), whose weight is normalised to 1. When $\lambda_k = 0$ all individuals are weighted equally (the graph is horizontal) regardless of their social distance. When $\lambda_k > 0$ (towards the front right of the cube), individuals become (exponentially) less important the more dissimilar they are (the graph resembles the right half of a normal density) and the rate at which they do so increases as λ_k grows. When $\lambda_k < 0$ (towards the back left) the opposite is true: individuals now become (exponentially) more important the more dissimilar they are, and the rate at which they do so increases the more negative is λ_k.

how much we care about others' outcomes. Previously each individual j's outcome was assigned a distance-dependent weight from the interval $(0, 1]$, indicating its importance relative to an (hypothetical) individual identical to i in dimension k. By allowing λ_k to be negative, we now let these proportions span the entire positive line from $(0, +\infty)$. Note that changing the sign of λ_k inverts the corresponding weights: an individual that would have had a weight of, say, $\frac{1}{2}$ if $\lambda_k = c > 0$ would instead have a weight of 2 if $\lambda_k = -c$. Moreover, as (4.17) makes clear, different dimensions can now offset each other: e.g. if we compare primarily to those who are geographically close ($\lambda_{loc} > 0$), but who are of a different ethnicity ($\lambda_{eth} < 0$), then someone of a different ethnicity ($\omega_{eth,ij} > 1$), living in a different locality ($\omega_{loc,ij} < 1$) could be just as important as someone of the same ethnicity in the the same locality ($\omega_{eth,ij} = \omega_{loc,ij} = 1$).

4.1.4 Estimation

The most obvious way to estimate the parameters of equation (4.11), or its fixed effects counterpart (4.15) is to assume the that the errors are normally distributed, and to proceed via maximum likelihood (or equivalently, non-linear least squares). However, one concern with such a strategy is the possibility that a standard ML estimator of equation (4.11) or (4.15) could over-fit the reference levels.

Recall that one interpretation of the reference level estimator is as a kernel predictor of the relevant outcome. However, in contrast to a standard kernel regression, where the kernel bandwidth can be adjusted to avoid over-fitting, and to reflect a desired level of smoothing in the regression surface, we lose this ability here, because we (implicitly) estimate the bandwidth jointly with the λ parameters of the distance function. To guard against the possibility of over-fitting, we employ a penalized likelihood approach (see Van Houwelingen, 2001), adding a quadratic penalty to the likelihood of the form:[15]

$$\pi = \delta\theta' P\theta \tag{4.18}$$

where P is a diagonal matrix that standardises the parameters to which the penalty is to be applied (in our case this will generally be just the subset λ, which we standardise by the variance of the corresponding z variables) and δ is a tuning parameter. Maximising the penalized likelihood $\ell - \pi$ will yield consistent estimates of the model parameters provided δ is fixed, or $\frac{\delta}{N} \to 0$ as $N \to \infty$ (see e.g. Ghidey, Lesaffre, and Eilers, 2004).

Most approaches to selecting the optimal value of the tuning parameter (defined that which minimises the expected mean squared error of the parameter estimates) employ a cross-validation approach. However, Ueki and Fueda (2008) derive an analytic expression for the optimal parameter that allows us to estimate it much more simply. We employ their "parametric direct plug-in estimator" here, which sets the tuning parameter to:

$$\hat{\delta} = \frac{\hat{\sigma}^2 \mathrm{tr}\{(Z'Z)^{-1}\}}{\hat{\lambda}'(Z'Z)^{-1}\hat{\lambda}} \tag{4.19}$$

[15] Adding this penalty term also has the benefit of avoiding a situation where the likelihood could become flat if any of the λ parameters diverge to ∞.

(where $\hat{\sigma}^2$ and $\hat{\lambda}$ are consistent estimates of σ^2 and λ obtained from an initial unpenalised maximum likelihood fit).[16]

Denoting the negative Hessian matrices from the penalised and unpenalised likelihoods as H and I respectively, we estimate the covariance matrix as $\hat{V} = H^{-1}IH^{-1}$ (following Gray, 1992), or $\hat{V} = H^{-1}$ (following Verweij and Van Houwelingen, 1994) if, as is sometimes the case, the former is not positive definite.

4.2 Does Information Dominate Comparison in Russia?

A widespread result in empirical SWB research has been the finding that others' income has a negative effect on individual satisfaction (Clark, Frijters, and Shields, 2008, provide a recent overview of this evidence). Recently, however, a number of papers have emerged to suggest that others' income can sometimes have a *positive* effect on SWB. In a series of papers, Claudia Senik

[16] In theory, more efficient estimates could be obtained by iteratively updating the estimator of $\hat{\delta}$ using the new values of $\hat{\sigma}^2$ and $\hat{\lambda}$ until convergence. In practice, we retain Ueki and Fueda's (2008) one-step approach to keep the computational burden down.

Relative to the optimal δ, the one-step estimator is likely to under-correct (in finite samples) precisely when it is most needed. The true optimal parameter δ (as opposed to our estimate, $\hat{\delta}$) depends negatively on the true value of λ. Because we do not know the true value, we substitute a consistent estimator $\hat{\lambda}$. For a given true λ, a higher estimate should require more correction (higher optimal δ), but will lead us to think it requires less (lower $\hat{\delta}$).

has presented evidence for such positive reference effects in Russia (Senik, 2004), and more recently in the rest of Eastern Europe and the US (Senik, 2008). She argues that these results reflect an "information" effect of the sort hypothesised by Hirschman and Rothschild (1973) — whereby individuals can take others' outcomes as a signal of their own future prospects — and suggests that the relative weight of such information, vis-à-vis the more traditional comparison effect, is mediated by perceptions of mobility: in countries where economic circumstances are volatile, and incomes are changing quickly (e.g. the transition economies of Eastern Europe) or where perceptions of mobility are otherwise high (e.g. the US) people treat others' income more as a signal than as a comparator — or as Senik (2004) puts it, "information dominates comparison." Clark, Kristensen, and Westergaard-Nielsen (2007) find a similar positive influence — which they also interpret as a signalling effect — of co-worker wages in Denmark.[17]

Nonetheless, while these papers mount a strong case that others' income can indeed have a positive informational effect on SWB, it is not yet clear that this information effect necessarily dominates comparison in an all-things-considered sense. Senik (2004, 2008) and Clark, Kristensen, and Westergaard-Nielsen (2007) both construct their reference levels in a way

[17]Kingdon and Knight (2007) also find a positive effect on SWB from others' income in South Africa, but, after tests, attribute this result to altruism rather than information.

that will tend to maximize their informational, rather than their comparative content: running wage regressions of individual income on productive characteristics (the predictions of which are most likely to correspond to "what I might expect to earn in the future"). It is possible that we could construct an alternative reference level that would better capture comparison effects, and which would lead to the opposite conclusion — that comparison dominates information. In line with our results in Part I (where we argued that income and consumption may have distinct effects on SWB), it is also possible that they could have distinct *reference* effects: while others' income seems most likely to have informational content for an individual's future prospects, others' consumption seems more likely to act as a comparator.

This section revisits Senik's (2004) Russian case study to investigate these issues in more detail. Section 4.2.1 outlines Senik's (2004) methodology and results. In Section 4.2.2 we present preliminary evidence confirming the suspicions outlined above: when we include reference levels in both income and consumption, the latter has a consistently negative impact on SWB, while only the former appears to capture a positive information effect. The net effect of the information and comparison effects — that is, whether one dominates the other — appears to depend crucially on the definition of the reference group: for some reference group definitions comparison clearly dominates, while for others, information does. This case study therefore seems to provide a promising application of the estimation strategy developed in Section 4.1. When we apply our estimator to the Russian case, we find both that the reference groups for consumption and income differ quite significantly, and that the negative effect of others' consumption outweighs the positive effect of others' income — that is, comparison dominates information.

4.2.1 Senik's Methodology and Results

Senik (2004) uses data from rounds 5-9 of the RLMS[18] (conducted from December 1994 to October 2000) to investigate the effect of changes in others' income on reported life-satisfaction. Following Clark and Oswald (1996), she includes in a standard SWB regression,[19] a reference income variable derived as the predicted wage from a set of 1st-stage wage regressions.

The wage regressions themselves are run separately for each round (so as to capture changes in the returns to individual characteristics over the course of the sample. They regress the log of individual income (from both wage- and non-wage sources) on education, gender, years of experience, and a set of profession, industry, and region dummies. To account for the possibility of

[18] Described in more detail in Chapter 2.

[19] The independent variables use include income, age, health, education, Russian as mother tongue, religious belief, marital status, gender, household size, occupation, industry, region, and round dummies.

selection bias, Senik uses the Heckman (1979) maximum likelihood selection model, with age, gender, and children under 7 as selection variables. (Age and children under 7 constitute the necessary exclusion restrictions).

When the predicted income from these equations is included in a set of ordered probit SWB regressions (both in pooled cross-sectional specifications, and when using Mundlak decompositions to account for correlated unobservables), its estimated coefficient is significantly positive, in contrast to what would be predicted by the relative income hypothesis.[20] Senik argues that this is because changes in an individual's predicted wage reflect changes in the income she can expect in the future, and that anticipation of this future income influences current SWB. To confirm this interpretation, she runs a variety of further tests, showing that the positive effect of reference income is higher for the young (who have more of a future to look forward to), for those who experience greater personal income volatility, and for those who say they are most concerned about the future.

4.2.2 Consumption vs. Income

In this section, we present a series of (fixed reference-group) estimations which we believe cast a degree of doubt on Senik's (2004) conclusion that

[20] Senik (2002) obtains the same result using linear estimators of the sort we employ here.

"information dominates comparison" in Russia. In line with her results, the information content of others' *income* does appear to dominate its comparison effect; but, as we demonstrate below, the opposite is true of others' *consumption*, which has a consistently negative impact on SWB. The interesting question then becomes "which of these effects dominates?" Unfortunately, the answer to this question depends quite crucially on the definition of the reference groups, and thus appears ill-suited to a fixed reference-group approach.

Geographic Means

We begin by estimating a simple version of equation (4.12), where we use geographic means of income and consumption in place of the social effects term $s(M_{-i,t})$:

$$u_{it} = x'_{it}\beta + \bar{m}'_{-i,gt}\rho + \eta_i + \varepsilon_{it} \tag{4.20}$$

As usual, u_{it} is individual i's reported life-satisfaction at time t, x_{it} is a vector of covariates including labour force status dummies, age, age squared, health, marital status, household size and round dummies, $\bar{m}_{-i,gt}$ is (depending on the specification) the mean of either log household income, log household consumption, or both, for all households in a given geographic area g (excluding individual i's household), and η_i and ε_{it} are unobserved individual and time-specific error terms.

We estimate equation (4.20) at each of the three levels of geographic aggregation on which the RLMS contains information. At the highest level, the country is split into eight geographic *regions*: Moscow & St. Petersburg, North & Northwest, Central & Central Black Earth, Volga-Vaytski & Volga Basin, North Caucasus, Urals, Western Siberia, and Eastern Siberia & Far East. With the exception of the metropolitan centres of Moscow and St. Petersburg, these regions correspond to the official economic regions of the Russian Federation. At the lowest level, the RLMS sample is divided into *districts*. In urban areas, these correspond to census enumeration districts designed to contain around 450 individuals (Sagers, 1989). In rural areas, the districts correspond to village settlements, which have an average size of around 250 (calculated from Table 1.6 of Goskomstat, 2002).[21] At an intermediate level, are survey *sites*, which reflect neighbourhoods of similar geographic size, but varying populations: in most cases (125 out of 166 sites) sites and districts are identical, though in some urban areas a site will consist of a number of districts (up to 17, but with an average of 3).

Tables 4.2, 4.3, and 4.4 display five sets of regressions at each geographic level. In Chapter 2 we checked that the coefficients on the controls comport

[21]This appears to have remained relatively stable from the 1989 census (average 255) to the 2002 census (average 249). Cf. the census enumeration district size of 350 in rural areas (Sagers, 1989).

with standard results in the literature, and so we focus only on the pecuniary variables and the reference levels here. In all specifications where the geographic mean of others' consumption is included, it attracts a sizeable (and, in all but two cases, significant) negative coefficient, suggesting that, in contrast to Senik's income results, the comparison effect of consumption strongly dominates its information effect.

Table 4.2: Reference Effects Regressions using Region Means of Household Income and Consumption

	(1)	(2)	(3)	(4)	(5)
(Log) Household Income	0.118***		0.095***	0.095***	0.095***
	(0.006)		(0.006)	(0.006)	(0.006)
(Log) Household Consumption		0.157***	0.116***	0.117***	0.117***
		(0.009)	(0.009)	(0.009)	(0.009)
Mean Income (Region)	0.084		0.058		0.253***
	(0.063)		(0.063)		(0.086)
Mean Consumption (Region)		-0.094		-0.121	-0.341***
		(0.081)		(0.081)	(0.110)
Sum of Reference Effects					-0.088
					(0.082)
Observations	70869	70869	70869	70869	70869
Groups	14298	14298	14298	14298	14298
Log-Likelihood	-83777.18	-83825.09	-83658.75	-83657.66	-83652.09
AIC	167594.4	167690.2	167359.5	167357.3	167348.2
BIC	167777.7	167873.5	167552	167549.9	167549.9

* p<0.10, ** p<0.05, *** p<0.01. Standard Errors (robust to clustering within panelids) in parentheses.

On the other hand, the effect of others' income is less consistent across different specifications. In line with Senik's income results, the regional means of others' income attract consistently positive coefficients; and while the coefficient only becomes significant when all variables are included together in column (5), this specification is least likely to be affected by omitted variable bias, and is therefore our preferred model. (For example, when own

Table 4.3: Reference Effects Regressions using Site Means of Household Income and Consumption

	(1)	(2)	(3)	(4)	(5)
(Log) Household Income	0.121***		0.098***	0.096***	0.096***
	(0.006)		(0.006)	(0.006)	(0.006)
(Log) Household Consumption		0.162***	0.118***	0.122***	0.122***
		(0.009)	(0.009)	(0.009)	(0.009)
Mean Income (Site)	-0.041**		-0.052***		0.003
	(0.017)		(0.017)		(0.021)
Mean Consumption (Site)		-0.095***		-0.106***	-0.109***
		(0.023)		(0.023)	(0.029)
Sum of Reference Effects					-0.106***
					(0.024)
Observations	70869	70869	70869	70869	70869
Groups	14298	14298	14298	14298	14298
Log-Likelihood	-83774.23	-83811.28	-83652.48	-83640.61	-83640.6
AIC	167588.5	167662.6	167347	167323.2	167325.2
BIC	167771.8	167845.9	167539.5	167515.8	167526.9

* p<0.10, ** p<0.05, *** p<0.01. Standard Errors (robust to clustering within panelids) in parentheses.

Table 4.4: Reference Effects Regressions using District Means of Household Income and Consumption

	(1)	(2)	(3)	(4)	(5)
(Log) Household Income	0.120***		0.097***	0.096***	0.095***
	(0.006)		(0.006)	(0.006)	(0.006)
(Log) Household Consumption		0.159***	0.117***	0.119***	0.120***
		(0.009)	(0.009)	(0.009)	(0.009)
Mean Income (Census District)	-0.019**		-0.022**		0.017
	(0.010)		(0.010)		(0.015)
Mean Consumption (Census District)		-0.036***		-0.040***	-0.054***
		(0.011)		(0.011)	(0.016)
Sum of Reference Effects					-0.038***
					(0.011)
Observations	70869	70869	70869	70869	70869
Groups	14298	14298	14298	14298	14298
Log-Likelihood	-83775.53	-83817.51	-83655.63	-83648.94	-83648.08
AIC	167591.1	167675	167353.3	167339.9	167340.2
BIC	167774.4	167858.4	167545.8	167532.4	167541.9

* p<0.10, ** p<0.05, *** p<0.01. Standard Errors (robust to clustering within panelids) in parentheses.

consumption is excluded the coefficient on others' income will pick up part of its positive effect on SWB, as we can see from the slight inflation of the others' income coefficient in column (1) compared to column (3). On the other hand, when others' consumption is excluded, others' income will also pick up part of its *negative* effect, attenuating the coefficients in columns (1) and (3) when compared with column (6).)

In contrast to the regional results, the site and district level means of others' income never attract sizeable or significant positive coefficients, even in our preferred all-variable specifications. (And indeed, in the earlier specifications, they attract significant *negative* coefficients — though this seems to be due primarily to the omission of others' consumption.)

The final column of each table also shows the sum of the reference levels. We see from this that, at all three geographic levels, the estimated net effect of others' economic outcomes is negative, albeit not significantly so at the regional level. That is, the consistent and significant negative effect of others' consumption at least offsets, and usually significantly outweighs any positive information effect of others' income.

Incorporating Predicted Wages

While these results cast doubt on the conclusion that information dominates comparison *overall*, much depends on the nature of both the assumed reference group, and the assumed reference outcomes — which so far may not have been as conducive as they could be to detecting information effects. As Senik argues elsewhere (Senik, 2008) information effects may be strongest for individual income, and are likely to stem from changes in the income of those with similar productive characteristics, rather than just those in the same geographic area. To provide a slightly fairer test of of whether it is information or comparison that dominates in Russia, we rerun the specifications from the previous section, but this time with a different income referent: instead of geographic means of household income, we now substitute individuals' predicted wages from a 1st stage wage regression.

In constructing the predicted wages, we try to follow closely Senik's (2004) original approach. Because we were unable to obtain industry code data for our sample (which is slightly longer than Senik's) we are forced to omit these from the income regressions; otherwise our wage equations employ the same independent variables for both the selection and main equations, with the only difference being that we include squared terms in age and experience.

Our wage equation regresses the the log of individual income on a vector of explanatory variables including years of education, years of experience, experience squared, and dummies for gender, occupation (with unemployed as the omitted category), and region. Following Senik, we include even the unemployed and those out of the labour force in the equation unless they report no individual income — both to keep these individuals in the sample, and because such individuals can also form expectations about non-wage income. The selection equation models the probability of selection as depending on gender, age, age squared, and number of children under 7 in the household (with the final three constituting the necessary exclusion restrictions). The results of these regressions, which we run separately for each round, are shown in Table A.8 in Appendix A.5.

The results of including the predicted income from these equations in the main SWB equation are shown in Table 4.5. As above, we include the standard set of controls, with the only addition being the inclusion of the log of individual income in the set of pecuniary variables (to prevent individuals' predicted income from picking up any residual effect of own income over and above the income of the household). Again, we do not report the control coefficients, focusing attention instead on the pecuniary variables. Column (1) replicates Senik's (2004) main result: predicted individual income now attracts a large and highly significant positive coefficient (which, in line with Senik, we interpret as an information effect). This result is maintained when we include geographic means of other households' consumption in columns (2)-(4); however, in all these specifications consumption also attracts a significant negative effect. Perhaps oddly, the largest comparison effect occurs at the regional level (-0.281), and and declines as we move to the site (-.109) and district (-0.040) levels. This is the opposite of what we might expect if individuals compared primarily to those who are physically close to them, but may be due to the declines in sample size as we move to the smaller geographic regions. Smaller samples are likely to lead to noisier measures of the reference level, both because the geographic means themselves are less accurately measured, and because the smaller geographic samples are less likely to contain individuals who are similar in *other* dimensions that may be important in defining the reference group.

Table 4.5: Predicted Income & Geographic Consumption Means

	(1)	(2)	(3)	(4)
(Log) Household Income	0.081***	0.082***	0.083***	0.082***
	(0.006)	(0.006)	(0.006)	(0.006)
(Log) Household Consumption	0.115***	0.117***	0.121***	0.118***
	(0.009)	(0.009)	(0.009)	(0.009)
(Log) Individual Income	0.013***	0.012***	0.012***	0.013***
	(0.002)	(0.002)	(0.002)	(0.002)
Predicted (Log) Individual Income	0.175***	0.191***	0.178***	0.175***
	(0.019)	(0.020)	(0.020)	(0.019)
Mean Consumption (Region)		-0.281***		
		(0.083)		
Mean Consumption (Site)			-0.109***	
			(0.024)	
Mean Consumption (Census District)				-0.040***
				(0.011)
Sum of Reference Effects		-0.090	0.068**	0.136***
		(0.081)	(0.030)	(0.022)
Observations	70869	70869	70869	70869
Groups	14298	14298	14298	14298
Log-Likelihood	-83566.32	-83557.69	-83546.51	-83555.95
AIC	167176.6	167161.4	167139	167157.9
BIC	167378.3	167372.3	167349.9	167368.8

* p<0.10, ** p<0.05, *** p<0.01. Standard Errors (robust to clustering within panelids) in parentheses.

Below the main coefficients in columns (2)-(4) of the table, we display the sum of the two effects. The results here depend quite critically on the geographic level at which we calculate the consumption means: at the regional level the net effect is negative, but not significantly so, while at the site and district levels, the net effect is significantly positive. Once again, it is difficult to draw a definitive conclusion about whether information or comparison dominates overall, and even more so than before, the definition of the reference group appears to be crucial to the results. In the next section, we therefore attempt to estimate the reference groups directly.

4.2.3 Estimating the Reference Group

The previous section 4.2.2 demonstrated the sensitivity of conclusions about whether information dominates comparison to the specification of the relevant reference groups. With purely geographically-defined reference groups, the negative comparison effect of consumption consistently outweighs the positive informational effect of income. When we shift to predicting individual income on the basis of productive characteristics, however, its informational value increases. While it is still outweighed by the effect of regional consumption, the net negative effect is no longer statistically significant, and if we focus on site or district means of others' consumption, our results switch, and the net effect becomes significantly positive.

At this point, our preferred model (i.e. that which best fits the data) is the model that includes as reference levels the site-level means of household consumption, and predicted individual income (column (3) of Table 4.5). In terms of both log-likelihood and information criteria (displayed at the bottom of each table), this model outperforms all the others shown in Tables 4.2- 4.5 (i.e. it has a higher log-likelihood and lower AIC and BIC[22]). However as we noted in Section 4.1, the models we have explicitly considered span only a small part of the potential model space, and the most-relevant others for either consumption or income may not be captured in the specifications we have tried. This is a particular concern for the reference level in consumption, which until this point has been specified only by geography, without reference to broader notions of social distance or (dis)similarity.

In this subsection, we attempt to jointly estimate the reference group and reference effects, using versions of the fixed-effects social-effects estimator developed in Section 4.1. In Section 4.2.3 we focus on household-level consumption and income as potential referents, while in Section 4.2.3 we go on to consider individual income in place of the latter.

Household Level Referents

To begin with, we estimate a version of equation (4.12) using the log of household income, and the log of household consumption as the reference

[22]See Chapter 3 for a discussion of these information criteria and their role in model selection.

variables:

$$u_{it} = x_{it}'\beta + \left(\sum_{h_j \neq h_i} \omega_{h_i h_j}^{(c)} c_{h_j t} \right) \rho_c + \left(\sum_{h_j \neq h_i} \omega_{h_i h_j}^{(y)} y_{h_j t} \right) \rho_y + \eta_i + \varepsilon_{it} \quad (4.21)$$

where the weight of individual j's household, h_j, in individual i's reference level $\bar{m}_{it} \in \{\bar{c}_{it}, \bar{y}_{it}\}$ is given by the formula:

$$\omega_{h_i h_j}^{(m)} = \frac{\exp(-z_{h_i h_j t}^{*m}{}' \lambda^{(m)})}{\sum_{h_l \neq h_i} \exp(-z_{h_i h_l t}^{*m}{}' \lambda^{(m)})}, \qquad m \in \{c, y\} \quad (4.22)$$

The x_{it} include the usual controls (labour force status dummies, age, age squared, health, marital status, household size and round dummies) as well as additional controls for the household-level z_{it} variables, and the aforementioned reference variables: log household consumption, and log household income.

Because the reference variables are defined at the household level, we are limited to using household level z_{it} variables to define social distance/ (dis)similarity, $z_{h_i h_l t}^{*(m)}$. This means in particular that we cannot use gender to define the reference group, as has often been done in the existing literature. However, we are able to include household-level versions of most of the other variables commonly used to define the reference group (as set out in Table 4.1). In addition to geographic indicators, we include three other variables that attempt to capture elements of a household's social position and productive capacity: the number of working age adults in the household; their average years of education, and the age of the household head. After some preliminary testing, we elected to include different geographical variables for consumption and income: while the reference level in income appeared to be best defined by region (setting $z^{*(y)}_{region,h_ih_jt} = 0$ if households live in the same region, and 1 otherwise), the reference level in consumption was better defined using a combination of both site and region variables (setting $z^{*(c)}_{site,h_ih_jt} = 0$ if households were in the same site, 1 if they were in different sites in the same region, and 2 if they were in different regions). We have checked, in an unreported set of regressions (and using a variety of specifications) that all four of these variables are significant predictors of both log household consumption and income.[23]

The results of this estimation are shown in Table 4.6 (omitting the controls to conserve space). The top panel shows the estimated main (β) and reference (ρ) effects for consumption (column 1), income (column 2), and their sum (column 3). Consistent with our expectations from earlier results, others' consumption has a large and significant negative comparison effect, while others' income has a large and significant positive information effect.

[23]The age of the household head enters with a squared term in a standard regression, but one advantage of the non-parametric/kernel aspect of our estimation strategy is that it will capture such non-linearities automatically.

Table 4.6: Reference Group Estimations (Household Income & Consumption)

	Consumption	Income	
	c	y	$c+y$
Main Effect β	0.129***	0.096***	0.225***
	(0.009)	(0.006)	(0.010)
Reference Level ρ	-0.285***	0.120***	-0.165***
	(0.051)	(0.035)	(0.053)
Sum $(\beta+\rho)$	-0.157**	0.217***	0.060
	(0.051)	(0.035)	(0.052)
Reference Group	$e^{-\lambda^{(c)}}$	$e^{-\lambda^{(y)}}$	$e^{-\lambda^{(c)}}/e^{-\lambda^{(y)}}$
# Adult Household Members	0.668**	0.719*	0.930
	(0.125)	(0.0124)	(0.224)
Household Average Education	0.999	0.863**	1.158**
	(0.003)	(0.055)	(0.073)
Age of Household Head	0.998	1.001***	0.997
	(0.002)	(0.000)	(0.002)
Site	0.037***		
	(0.011)		
Region	0.001***	0.000***	7.409
	(0.001)	(0.000)	(12.553)
Observations	73244		
Individuals	14543		
Penalised log-likelihood	-86677.85		

* p<0.10, ** p<0.05, *** p<0.01. Standard errors (robust to clustering within panelids) in parentheses.

Both are in fact larger than their main effects — although only in the former case is the difference significant. Turning to look at the net effect of the two reference levels in column 3, we now find that it is negative, and highly significant: contra Senik (2004), comparison dominates income.

We turn now to the coefficients that define the reference groups, which are displayed in the bottom panel of Table 4.6. Rather than displaying the raw λ coefficients, we instead display them in the transformed metric $e^{-\lambda}$, which represents the *proportional* decrease in a household h_j's weight that results from a 1-unit change in social distance $z^*_{kh_ih_j}$. In this metric, a coefficient of 1 represents the null hypothesis that the corresponding variable is not important to defining the reference group (and it is against this that statistical significance is assessed). Recall also that the weightings reflect the importance of *individual* households, rather than the *collective* importance of households of a particular type (which will depend on the number of household's of that type in the population). So, even if (say) geographically distant households individually have a low weight, they may still be *collectively* important in determining the reference level, if there are sufficiently more of them in the sample than geographically proximate households. The relative collective weight of households in a given distance interval $[a, b]$ is:

$$\frac{\sum_{j \neq i: z_{h_i h_j} \in [a,b]} \exp(-z_{h_i h_j} \lambda)}{\sum_{j \neq i: z_{h_i h_j} = 0} 1} \tag{4.23}$$

In the first column, we see that geography and the number of working age adults are important in defining the consumption reference group, but that education and the age of the household head are not, despite the fact that both of these variables are important predictors of household consumption. The reference group is defined in fairly narrow geographic terms, with households from different sites in the same region having weights of 3.7% of those in the same site, and those from different regions weighing only 0.1%

($\approx 0.037^2$) of that. This, however, needs to be taken in the context of the knowledge that there are many more households outside any given site than within it, so that the *collective* weight of distant households will still be fairly large, even though they have little effect individually. In fact, the collective weight of those in the same region but not the same site is on average around 70% of the collective weight of those in the same site ($\approx 0.037 \times 434/22$), while the collective weight of those outside the region is on average around 20% of those in the same site ($\approx 0.001 \times 3198/22$). Households that differ by 1 in their number of adult members weigh 67% of those with the same number, while those who differ by two adult members will weigh $0.668^2 \approx$ 45% of that.

In contrast, as we see in column 2, all four z variables are significant in defining the reference group for income. Households in the same region, with a similar number of adult household members, and similar educational background are more important sources of information than more dissimilar households. On the other hand, age has the opposite effect: here, it is households whose head is older or younger than their own that are more important than similar households. This is consistent with the idea, put forward in 4.1.3, that people may take informational cues from those who are a few years more advanced than them. Unfortunately, however, our specification forces social distance to affect the weight other households receive in a way that is both symmetric (households with younger heads weigh as much as those with older heads) and monotone (weight increases with distance over the entire range of ages; it cannot increase and then decline), so we cannot confirm this interpretation conclusively.[24] The effect of age is also relatively small — even a 10-year age difference will lead to only a 1 percentage point increase in importance ($1.001^{10} \approx 101\%$), although again, the collective weight of those of different ages will be larger, due to the greater number of households involved.

The different results for consumption and income seem to underscore the importance of allowing the reference groups to vary, and this is confirmed in column 3, which displays the ratios of the transformed coefficients in the previous two columns (so again, a ratio of 1 corresponds to the null hypothesis of no difference between the weights for income and consumption). We see here that, while the differences for region, age and number of adults are not significant, education is significantly more important in defining the reference group for income than for consumption. Such a difference is entirely consistent with the coefficient on others' income picking up an information effect (the information value of others' outcomes will be higher when they

[24] In theory, we could adapt our esimator to allow for asymmetry (i.e. different "upward" and "downward" comparisons). Allowing for non-monotone changes in the weights is more difficult, but it would be possible instead to introduce an additional shift parameter that would allow the distance at which individuals would be maximally-weighted to vary (so that, e.g. the weight function could reach a maximum at X years older than houesehold h, and decline monotonically either side of that point).

have similar productive characteristics); whereas it is perfectly conceivable that we would compare equally to the consumption of those with more or less education.

Incorporating Individual Income

Mirroring our approach in Section 4.2.2, we now consider what happens if we take the income referent to be individual, rather than household income. The specification is of the same form as (4.21), but now with the log of individual income replacing household income as the second reference variable. The controls are identical, except that we now include the logs of both household and individual income in the set of x_{it} variables. The $z^{(c)}$ variables are also unchanged. However, we now include individual-level $z^{(y)}$ variables to match the predicted wage equation estimated in 4.2.2: years of education, years of experience, gender, profession, and region. (One nice feature of the kernel predictor is that its non-linearity removes the need for exclusion restrictions.) The results of this regression are displayed in 4.7.

The top panel once again displays the main and reference effects for consumption and income. As expected, all three of own household consumption, household income, and individual income enter positively. When we turn to the reference effects we see the now familiar pattern of a significant positive effect for income, and a significant negative effect for consumption. While the

Table 4.7: Reference Group Estimations (Individual Income & Household Consumption)

	Consumption	Income	
	c	y	$c+y$
Main Effect (Household) β_h	0.127***	0.084***	0.211***
	(0.009)	(0.006)	(0.010)
Main Effect (Individual) β_i		0.012***	
		(0.002)	
Reference Level ρ	-0.258***	0.140***	-0.118**
	(0.040)	(0.029)	(0.048)
Sum $(\beta+\rho)$	-0.131***	0.237***	0.105**
	(0.040)	(0.029)	(0.049)
Reference Group	$e^{-\lambda^{(c)}}$	$e^{-\lambda^{(y)}}$	
# Adult Household Members	0.789***		
	(0.064)		
Household Average Education	1.004		
	(0.003)		
Age of Household Head	0.999		
	(0.001)		
Site	0.021***		
	(0.008)		
Years of Education		1.008***	
		(0.002)	
Years of Experience		0.998	
		(0.002)	
Gender		0.248**	
		(0.135)	
Profession		0.033***	
		(0.034)	
Region		0.588	
		(0.262)	
Observations	70869		
Individuals	14298		
Penalised log-likelihood	-83482.107		

* p<0.10, ** p<0.05, *** p<0.01. Standard errors (robust to clustering within panelids) in parentheses.

comparison effect has reduced slightly and the income effect has increased, the net effect remains significantly negative: once again, comparison dominates information.

Turning to the λ coefficients, we see that the coefficients defining the consumption referent are virtually unchanged from Table 4.6, and so we say no more about them here. When we turn to the income λs we find that gender (which we were previously unable to include at the household level) is now very important in delineating the reference group: those of the opposite gender matter less than a quarter as much as those of the same sex. Individuals also compare themselves primarily to those in their own profession: those in different professions are weighted only 3% as much as those in the same profession. Given that there are, on average, 10 times as many individuals outside one's profession as inside it, this means that their collective weight is 30% of similar others.

On the other hand, the results once again appear to confirm that that just because particular variables reliably predict a referent outcome, this does not imply that they will also significantly affect the reference level. Neither years of experience, nor region of residence are significant in defining the reference group. Although the size of the region effect does appear to be substantively important — those in different regions individually count for only around 60% of those in the same region — individuals nonetheless appear to place significant collective weight on the outcomes of geographically distant others. Those from different regions collectively matter over four times more, on average, than those in the same location (because there are on average 7 times more individuals outside one's region as inside it).

Somewhat surprising is the coefficient on education, which has now switched sign, so that those with different education levels are more important to the reference level. Nonetheless, while this effect is statistically significant, it is substantively rather small, with a one-year difference in education resulting in a less-than-one-percentage-point increase in weight.

4.3 Conclusion

In this chapter we have made a number of novel contributions to the SWB literature on reference effects. Building on the results of the previous chapters — which argued for the importance of distinguishing between consumption and income in estimating the effect of individuals' own economic outcomes on SWB — we have shown that it may be even more important to distinguish between the effects of *others'* consumption and income.

Revisiting Senik's (2004) analysis of the RLMS, we confirm her finding that others' income has a positive effect on SWB. However, we also show that others' *consumption* has a countervailing negative effect. While, as Senik convincingly argues, others' income does seem to function as a signal of future prospects, their consumption continues to have a more traditional comparison effect.

We go on to show that which of these effects dominates depends crucially on how we define individuals' reference groups. For some definitions of the reference groups, information dominates comparison; for others, comparison dominates information; and for others still neither dominates the other. To avoid the problem of having to simply impose a particular conception of the reference group, we employ a novel estimation strategy — based on the concept of social distance — that allows us to estimate both the reference group and reference effects jointly.

In addition to allowing the data to define the most appropriate reference group, this approach also generalises existing approaches to reference-group definition in a way that allows for more much more flexibility in what factors contribute to defining reference groups, and exactly how they do so. Our results broadly support both existing practice, and previous survey-based evidence suggesting that reference groups tend to be local rather than global. However, they also temper these conclusions in at least two ways.

First, they suggest that this is only true for some dimensions of social distance. Crucially, not all variables that help to predict income or consumption levels are necessarily relevant to defining the reference groups for those variables — indeed, the reference groups differ between the two outcomes. While geographic proximity and household size are most relevant to defining the comparison group for household consumption, productive characteristics (education and profession) and gender are most relevant for the informational effect of income.

Second, even when they do suggest that comparisons are *primarily* local, they typically still allow for some element of global comparison. For example, in estimating the reference level in consumption, it emerges that, while we do compare *more* to those who reside close to us, those who reside elsewhere are not irrelevant to the reference-level, as would be assumed by a more standard cell-mean approach to defining the reference group. Of course, better reference group definitions are interesting not only in their own right, but also because they should allow us to more accurately estimate the reference effects in which we are interested, and which motivated our re-analysis of the RLMS data. The results of our reference group estimation confirm both the negative comparison effect of others' consumption, and Senik's (2004) finding of a positive informational effect of others' income. In contrast to Senik's results, however, we find consistent evidence that, once we allow the data to determine the reference group, comparison ultimately dominates information.

Part III

Heterogeneity in the Determinants of SWB

Chapter 5

Heterogeneity in the Determinants of SWB

That people have different tastes for everyday items like the types of food they eat, the books they read, or the music they listen to is obvious. Yet, as we move from the micro level of apples vs. oranges, to the macro level of nutrition, shelter, meaningful relationships, status, and fulfilling work or hobbies, it becomes more common to assume that people, in the end, all value similar things to similar degrees. This tendency to universalise can be seen in writings ranging from Maslow's (1943) hierarchy of needs, to Martha Nussbaum's (1993; 2000) attempts to flesh out capability theory (e.g. Nussbaum and Sen, 1993) with a list of basic "objective" goods, that we should value, whatever else we value.[1]

Whether, or at what level of abstraction we can begin to say that individuals' really do value the same things is the focus of Part III of this thesis. At a deeper level than apples and oranges, individuals may differ quite profoundly in what is important to them: some people may place particular importance on material well-being, others on relationships with friends and family; still others may especially value good health, or a rewarding job. One way to assess the 'value' of such outcomes to individuals is by examining their impact on subjective well-being (SWB). Yet, while the extent and structure of such differences in the determinants of SWB may have important consequences both for the conduct of SWB research, and for the conclusions we draw from it, this remains an understudied area.

Part III of this thesis has three goals. First, in Section 5.1, we provide an overview of the potential importance of accounting for individual heterogeneity in the determinants of SWB. We argue that the presence and structure of individual differences may have significant implications for both individual decision-making, and for public policy — and focus in particular on the implications of whether these differences are observable, and to whom. While there are many examples in the existing literature that examine the

[1] This latter, in particular, clearly demonstrates the tendency to universalise more at higher levels of abstraction. After all, the capability approach also makes much of the idea that different goods can differently valuable to different people in different circumstances, depending on how they are able to translate such goods into valuable functionings.

observable correlates of individual heterogeneity, the literature attempting to identify unobservable individual differences is in its infancy. While there have been some attempts to identify the full scale of both observable and unobservable differences, they have so far focused on only a small number of outcomes, such as the effect of income on financial satisfaction (Clark et al., 2005) or adaptation to marriage and divorce (Lucas et al., 2003). As such, they have been unable to address the possibility of correlations between the importance of different outcomes; and have also had difficulty distinguishing "true" variation in the determinants of SWB from variation in how individuals *report* SWB.

In Section 5.2 we employ two types of random coefficient model (the second of which is a semi-parametric model new to the economics literature) to simultaneously examine variation in seven distinct determinants of SWB in Britain. In line with previous studies by e.g. Clark et al. (2005) and Lucas et al. (2003), we find that there is indeed significant individual heterogeneity in how different life outcomes translate into reported SWB. However, in contrast to their work, we are able to demonstrate that such differences are real, and cannot be due merely to differences in how individuals report satisfaction. Focusing on multiple types of heterogeneity simultaneously also suggests that individual values tend to cluster around a materialism–non-materialism axis, with those who value e.g. consumption and relative income more tending to value employment, health, relationships, social time, and children less (and vice versa). Consistent with previous work, we also find that those at the materialistic end of the spectrum are (unobservably) less satisfied than their less materialistic counterparts. Finally, we are also able to show that the outcomes of most value to the SWB-poor are health, consumption, and employment. To improve their well-being, we should focus either on improving these outcomes, or on increasing their ability to benefit from more "social" factors.

Chapter 6 goes on to examine in more detail one of the implications of such unobservable heterogeneity — focusing on the ability of SWB research to inform individual decision-making. The presence of significant heterogeneity in the determinants of SWB will tend to reduce the usefulness of population-average coefficient estimates to both individuals and policy-makers. However, their value also depends crucially on the quality of individuals' own knowledge about what will bring them satisfaction. We therefore conduct a number of novel tests of the quality of individuals' private information, using a series of questions in the BHPS that ask individuals directly what is important to them. We find that, while the quality of individual information varies both between different domains and between individuals, we are in general surprisingly poor at knowing what will bring us satisfaction, and that incorporating the results of SWB research may therefore help to improve individual decision making.

5.1 Why is heterogeneity important?

If individuals differ significantly in what contributes to their SWB, then the sorts of questions typically asked by happiness economists — including the age-old classic of whether money makes people happy — may obscure important aspects of reality. Money may indeed buy happiness for some, but at the same time do little for others; there may even be a portion of the population for whom it brings *un*happiness (for example, if increasing income raises expectations more than the increase itself). Of course, depending on what we intend to do with the answers to these questions, a focus on average effects may be sufficient; but in other cases, awareness of individual heterogeneity may be important for both individuals and policy-makers. As we explain below, these implications depend heavily on the extent to which heterogeneity is observable to econometricians, policy-makers, and individuals.

5.1.1 Observable Heterogeneity

Observable heterogeneity is heterogeneity that we can either observe directly (e.g. we might learn that Alice responds more to an increase in consumption than Ben) or that we know correlates with something we can observe directly (e.g. we might learn that the highly-educated are more affected by unemployment than those with less education).

For policy-makers, being able to observe the correlates of individual differences may allow better targeting of policies to particular groups. Kaplow (2008) provides an extensive treatment of the implications of preference heterogeneity for optimal tax policy. The details depend on the precise nature of the heterogeneity,[2] but the basic principles are fairly straightforward: other things equal, efficiency suggests that those with a higher marginal utility of consumption should face a lower overall tax burden (because they suffer more from taxation), while a concern for equity suggests they should face a higher tax burden (because they will generally have higher utility for a given level of income).

On the other hand, if we allow factors other than consumption to also affect the level of individuals' SWB, then the equity effect could be reversed, or at least attenuated. If, as some have suggested (e.g. Diener and Oishi, 2000), materialistic individuals have a higher marginal utility of consumption, but are also generally *less* happy, then equity and efficiency concerns could align to suggest that they should face a lower tax burden. For this reason, the correlations between individual values, and levels of individual SWB may

[2] Indeed, one of Kaplow's (2008) main contributions is to show that different ways of parameterising heterogeneity can have quite different implications for the marginal utility of consumption, and thus for optimal policy settings.

be especially important to policy.

The general principles of efficiency (targeting those who gain the greatest benefit from a given outcome) and equity (targeting those who are "SWB-poor") generalise beyond tax policy. If, say, poor health, or unemployment were more psychologically damaging to those in particular age groups, then this may affect policy priorities for reducing unemployment or improving health care. Such conclusions would be reinforced (or weakened) if these same individuals were also SWB-poor (or SWB-rich).

If differences in the determinants of individual well-being are not only observable, but also *malleable*, then this may open up additional avenues for individuals or policy-makers to attempt to improve SWB. While there are obvious reasons to be wary of governments "engineering" individual tastes, information about the potential drivers of taste could nonetheless be useful in a variety of ways.

First, knowing what sorts of factors drive changes in individuals tastes could allow *individuals* to manage their own tastes more effectively, by taking steps to increase their sensitivity to positive life events, or to reduce their susceptibility to negative ones. For example, if tastes are in some sense "contagious" then perhaps we should associate more with those who have SWB-conducive tastes, and avoid those who have SWB-damaging ones.

Second, it should also be noted that not all heterogeneity in individuals' responses to particular life events will be due to differences in what they value. If we know that some individuals gain a great deal from consumption, while others gain relatively little, this could be either because they value consumption to different degrees, or because they engage in different types of consumption, with different hedonic returns. It has been argued, for example, that (at least above a certain level) consuming "experiences" tends to bring greater satisfaction than consuming "things" (see e.g. Van Boven and Gilovich, 2003; Van Boven, 2005). Not only could identifying the factors responsible for more satisfying consumption, jobs, or relationships be useful to individuals, but policies that attempt to improve the quality of such outcomes seem less inherently objectionable than policies that attempt to shape individual values themselves.

Finally, there may be some areas where government "taste-management" could be considered legitimate. A possible example is the negative externality that results from individuals' comparing their income or consumption to others': if there were potential policy levers that could reduce the negative effect of others' income, for example, then these could provide an alternative way to address the externality — and in contrast to taxation, (see e.g. Boskin and Sheshinski, 1978; Layard, 1980, 2005b) or other traditional policy instruments (see e.g. Frank, 1985) to do so in a way that does not require reducing others' well-being in the process.

As we discussed in Chapter 4, Senik (2004, 2008) has argued that economic volatility and individuals' perceptions of mobility may have such an effect — causing people to focus on others' income more as a signal of their future prospects than as a point of comparison.[3] While fostering economic volatility in the name of externality reduction seems an unappealing trade-off, fostering economic mobility is a rather more attractive one. In a similar vein, Benjamin Friedman (2005) has also argued (albeit from the more discursive perspective of an economic historian than the statistical perspective of an econometrician) that one of the major benefits of economic growth is that, when the economic pie is growing, we worry less about how much of it others are getting.

Given the potential value of identifying observable correlates of heterogeneity in the determinants of SWB, and the relative ease of investigating them — either by estimating separate equations for different sub-samples (e.g. men and women) or by including various sorts of interactions in SWB equations — it is not uncommon to find such strategies amongst the existing literature. Table 5.1 surveys some examples of the sorts of heterogeneity that have been examined to date. Some of the particularly interesting results include the findings: that right-wing people in Germany adapt less to in-

[3] While Chapter 4 presented evidence that the information effect may not dominate the comparison effect (as Senik (2004) claimed) it did not challenge her broader point: that volatility and perceptions of mobility increase the information value of others' outcomes.

come, while those on the left adapt less to status (Di Tella, Haisken-De New, and MacCulloch, 2007); that religion moderates the effect of income (and in particular mitigated the negative impact of transition) in Hungary (Lelkes, 2006); and Senik's finding (discussed earlier) that volatility and mobility can affect the extent to which individuals compare themselves to others.

5.1.2 Unobservable Heterogeneity

Of course, while running regressions on separate sub-samples, or including interactions in our models can clearly provide insights into heterogeneity that pooled analysis would miss, these strategies are inherently limited by their ability to detect only observable differences, and can therefore only place a lower bound on the amount of heterogeneity in the population as a whole. The presence of additional *unobservable* heterogeneity also has potentially important implications for both individuals and policy.

If, as suggested in the previous section, first-best policies (e.g. optimal taxation rates) would be tailored to each individual's personal preferences, then the presence of unobservable preference heterogeneity will result in sub-optimal policy settings (relative to the first-best), and creates potential scope for welfare improvements. The most obvious route to improvement is simply to work harder to find observable predictors of heterogeneity, which would

Table 5.1: Examples of existing work examining observable heterogeneity in the determinants of SWB

Paper	Focus	Approach
Clark et al. (2008)	Adaptation to changes in marital, family, and labour force status)	Separate regressions by gender
Di Tella, Haisken-De New, and MacCulloch (2007)	Adaptation to income and status	Separate regressions by gender, political views, and self-employment
Frijters, Haisken-DeNew, and Shields (2004)	Income	Separate regressions for men and women
Ferrer-i-Carbonell (2005)	Comparison Income	Separate regressions for West vs. East Germans
Lelkes (2006)	Income	Separate regressions by level of religiosity
Levinson (2009)	WTP for air quality	Interactions with environmentalism, income, health, and level of local pollution
Senik (2004)	Others' Income	Separate regressions for groups defined by age, Income volatility, and concern about the future
Senik (2008)	Others' Income	Separate regressions for different countries
Stutzer and Frey (2006)	Marriage	Separate regressions for groups defined by eventual divorce, wage differences between partners, labour specialisation, children, and differences in education

then enable us to target policy better. In this sense, empirically establishing the extent of "unobservable" heterogeneity can be seen as providing a sense of the potential gains from further research.

However, even where we cannot (yet) identify observable correlates of individual differences, the presence of unobservable heterogeneity can still affect policy design. In particular, it is important to remember that whether heterogeneity is observable is a property of our knowledge about individual differences, not an inherent property of those differences *per se* — and observability will likely vary depending on whose perspective we consider. Heterogeneity that is unobservable to econometricians and policy-makers may nonetheless be observable to the individuals involved. Other things equal, decision-making should be located with those who have more of the relevant information — which will tend to favour more decentralised decision-making, the more unobservable heterogeneity there is (though this could be outweighed by distributional or incentive concerns).[4]

A related point arises when we consider matters from an individual perspective. At an individual level, one possible response to a finding that a particular factor is or is not important to SWB *on average*, is to wonder whether such a result is relevant at all: why should any of us assume that we are *average* if we already know ourselves and our values well? The extent to which such a response is reasonable depends crucially on (a) the quality of individuals' private information about their values; and (b) the degree of unobservable variation in individual values. The better we know ourselves,

[4]Basu and Meltzer (2007) provide an application of this principle in the context of health care decision-making.

and the more ‘unique’ we are compared to (observably similar) others, the less useful information about a (sub)population average will be to us. Psychologists and behavioural economists have had a fair amount to say on (a) (and we shall revisit this issue in the following chapter); yet (with a few notable exceptions) SWB research has to date had little to say about (b).

As Clark et al. (2005) note, increasing use of panel data has allowed researchers to easily control for unobservable heterogeneity in *levels* of well-being. The importance of doing so has been amply demonstrated by (amongst others) Ferrer-i-Carbonell and Frijters (2004), who present evidence from the GSOEP that failure to control for such unobservable sources of heterogeneity can significantly bias estimates of the effects of various life-outcomes on SWB.

Nonetheless, attention to value heterogeneity *per se* — that is, to heterogeneity in the slopes or coefficients of standard linear or generalized linear models — has been more limited, and it is on the presence of such heterogeneity that Clark et al. (2005) focus. Using three waves of data from the European Community Household Panel (ECHP), they estimate a latent class ordered probit model for financial satisfaction that allows for variation both in the intercept, and in the coefficient on income.[5] Their results suggest that

[5] The latent class approach assumes that each individual falls into one of a finite number of classes. Which class an individual falls into is unobserved, but this can be treated as a missing data problem, and class membership probabilities can be estimated via the EM algorithm.

the individuals in their sample can be divided into four distinct groups, which differ in their average reported satisfaction, in the way they map income into such satisfaction, and in terms of their other observable characteristics.

At one extreme, they find a relatively small group of individuals (9% of the sample) who have both a high level of satisfaction and a large response to income. Members of this group are on average richer, older, more educated, have fewer children, are more likely to be out of the labour force, and are more likely to reside in northern European countries (specifically, the UK, Ireland, the Netherlands, Belgium, Luxembourg, or Denmark) than average.

At the other extreme, they find a sizeable group (around 30% of the sample) with both low reported satisfaction, and a relatively small response to income. Members of this group are poorer, younger, less educated, more likely to be single, and to live in Italy, Spain or Greece than average. Between these extremes there are two similarly-sized groups (again around 30% of the sample each), each somewhat intermediate in of their intercepts, slopes, and observable characteristics.

Interestingly, these results appear to contradict the widespread finding that materialistic individuals (i.e. those who have a greater response to income) have lower SWB. One possible reason for this divergence is the use of *financial* satisfaction as a dependent variable, rather than a more global cognitive or emotional measure. It is possible that those who respond most to income are both more financially satisfied, and less globally satisfied than others.

More recently, Bruin and Winkelmann (2008) have also used latent class techniques to examine heterogeneity in the degree of altruism parents have towards their (adult) children in Germany. They assume that parents may be either selfish, in which case their children's happiness has no effect on their own life satisfaction, or altruistic, in which case it has a positive effect (the size of which is to be estimated). They estimate the proportion of altruists — for whom children's SWB has a significant positive effect — at approximately 20% of the GSOEP sample.[6] They also provide evidence that predicted altruistic preferences are associated with higher transfer payments to children.

A slightly different approach to modelling heterogeneity is taken by Lucas et al. (2003), who employ a multilevel random coefficient model to examine differences in the extent to which marriage and widowhood result in lasting changes in SWB. Instead of assuming that individuals fall into a number of discrete classes, this approach instead models variation as a continuous (normally distributed) random variable. They find that while, on average, individuals revert quickly to baseline levels of happiness after marriage, there

[6] In contrast to much of the social effects literature, they assume that parents' SWB has no effect on children's' SWB, and thus ignore the possible simultaneity issues that could arise if such effects were present (see e.g. Manski, 1993).

is significant variation around this average: for some, marriage brings lasting happiness, while for others it can actually lead to reductions in SWB. Reactions to widowhood also varied significantly. While individuals on average reacted negatively to widowhood, and then partially adapted, there was significant variation in the pattern of long-term response, with some individuals even becoming happier following the death of their partners.[7]

5.2 Identifying Heterogeneity in the determinants of SWB

As Clark et al. (2005) note, a potential weakness of their analysis, and one which also afflicts other work in this area to date, is that they were unable to distinguish "true" heterogeneity in the effects of money from differences in the way individuals use self-report scales to express satisfaction. In light of this, a major contribution of this chapter is to demonstrate, not only that there are significant differences in the response of reported SWB to various life outcomes, but also that these differences are *real*, and not merely due to differences in the way individuals *report* SWB.

Using BHPS data (described in more detail in Chapter 3), we provide

[7] One possible explanation for this is that when one's partner has been terminally ill and/or debilitated for a long period, death may in fact provide something of a respite. Such an effect seems to mirror the apparent uptick in SWB transitioning from separation to divorce.

three pieces of evidence for this conclusion. The first two of these rely on allowing simultaneously for heterogeneity in the impact of multiple life outcomes (a strategy which, as far as we are aware, is novel in the literature). If all the estimated heterogeneity were merely due to (affine) translations of individuals' reporting functions, then this would imply (a) a particular structure of correlations between the effects of different outcomes; and (b) that there should be no variation in the consumption equivalents of different outcomes. We show that both of these implications are conclusively rejected by the data.

In addition, we also demonstrate that for most outcomes, there is a sizeable proportion of individuals for whom the effect of that outcome is of the opposite sign to the average (e.g. individuals for whom increased consumption has a negative effect). Because individuals' reporting functions should be monotonically increasing in "true" SWB, variation in reporting functions has difficulty accounting for such a finding. In contrast to Lucas et al. (2003), who find oppositely-signed effects for some individuals from both marriage and widowhood, we use a novel (to the economics literature) semi-parametric estimator to verify that this result is not merely an artefact of assuming normality of the random coefficients. This result also stands in contrast to Clark et al. (2005), who find that income has a positive effect in all the groups they identify (albeit in their case on financial satisfaction, rather than life satisfaction more broadly).

In addition to establishing the existence of real (non-reporting) heterogeneity in the determinants of SWB, we also make a preliminary attempt to understand some of its structure. We find that, while demographic characteristics are capable of explaining some of the variation in individual responses (and in particular that the importance of reference income, unemployment, and hours of work exhibit similar age-patterns over the life-cycle), the lion's share remains unexplained even after accounting for such factors. Understanding the causes and correlates of such residual heterogeneity thus remains an important task for future research.

On the other hand, while much of the heterogeneity we identify is unrelated to observable demographic variables, we are able to discern some underlying (semi-observable) structure in it. The contributors to individual SWB appear to cluster — at least to some extent — around an axis of materialism–non-materialism. In line with previous results (e.g. Belk, 1985; Richins and Dawson, 1992; Kasser and Ryan, 1993, 1996; Sirgy, 1998; Sirgy et al., 1998; Diener and Oishi, 2000; Van Boven, 2005), we find that those who cluster at the materialistic end tend to be less satisfied with life. However, in contrast to suggestions that this effect may arise because those who are materialistic are pursuing goals that are in some sense fundamentally unsatisfying (e.g. Van Boven, 2005), our findings are more consistent with the view put forward by e.g. Sirgy (1998) and Sirgy et al. (1998): that material goods do produce satisfaction for materialistic individuals, and that the negative effect is due to having higher material aspirations.

The remainder of the section proceeds as follows. Section 5.2.1 describes the implications of heterogeneity in individuals' reporting functions, and how we can establish whether it responsible for apparent findings of heterogeneity in the determinants of 'true' SWB. Section 5.2.2 employs the LMM approach to identify the presence and structure of individual differences and to assess the degree to which they could be due to differences in reporting. Section 5.2.3 employs a semi-parametric PGM approach to examine the robustness of our conclusions to violations of the LMM's normality assumption. Section 5.3 concludes.

5.2.1 The implications of (affine) heterogeneity in the reporting scale

As Clark et al. (2005) note, it is difficult do distinguish econometrically between true heterogeneity in the importance of any given variable to SWB, and heterogeneity in the reporting functions individuals use to translate SWB into a verbal scale. However, even if we cannot separately identify these two sources of apparent heterogeneity, this does not mean that we can say nothing useful about the structure of value heterogeneity at all. In particular, given certain assumptions about the form of the reporting function, we can: (a) test whether reporting heterogeneity alone could account for differences in the way individuals map life-outcomes into SWB; and (b) estimate the distribution of money equivalents (or indeed, any other equivalents) that represent the trade-offs between outcomes.

As in most of the literature, we assume a standard linear form for "true" SWB, but in addition, we allow the coefficients to vary between individuals and time-periods. (We shall later model this heterogeneity explicitly, but leave it unmodelled here for generality.)

$$u^*_{it} = q'_{it}\psi^*_{it} + \alpha^*_i + \epsilon^*_{it} \tag{5.1}$$

This equation represents the individual- (and time-) specific mapping from a vector of observed outcomes q_{it} into "true" SWB u^*_{it}. The further mapping from "true" SWB to reported SWB can be represented as $u_{it} = f_i(u^*_{it}) + \epsilon_{it}$. If we assume that the function $f(\cdot)$ can be reasonably approximated by an individual-specific affine transform, $f_i(x) = x\phi_i + \alpha_i$, then we have:[8,9]

$$\begin{aligned} u_{it} &= u^*_{it}\phi_i + \alpha_i + \epsilon_{it} \\ &= q'_{it}\psi_{it} + \eta_i + \varepsilon_{it} \end{aligned} \tag{5.2}$$

(where $\psi_{it} = \psi^*_{it}\phi_i$, $\eta_i = \alpha^*_i\phi_i + \alpha_i$, and $\varepsilon_{it} = \epsilon^*_{it}\phi_i + \epsilon_{it}$, and we assume, without loss of generality, that $\mathrm{E}(\phi_i) = 1$).

Estimating this equation using standard methods, we cannot identify heterogeneity in any individual component ψ^*_{pit} of ψ^*_{it} separately from heterogeneity in the reporting function induced by ϕ_i. We can, however, determine whether heterogeneity in the ϕ_i alone could account for variation in the ψ_i (which would require that $\psi^*_{pi} = \psi^*_p$ for all p).

First, and most obviously, if we make the reasonable assumption that $\phi_i \geq 0$ (i.e. that the reporting function is weakly increasing in u^*_{it}) then $\psi^*_{pi} = \psi^*_p$ will imply that $\mathrm{sgn}(\psi_{pi}) = \mathrm{sgn}(\psi_{pj})$ for all p, i, and j. Identifying variation in the sign of the estimated ψ_{pi} coefficients would therefore indicate true variation in the signs of the ψ^*_{pi}.[10]

[8]This equation assumes for simplicity that reported SWB is continuous rather than discrete. However, this assumption is not crucial to the analysis. A further mapping from u_{it} to a discrete scale can easily be introduced by treating u_{it} as a latent variable in the usual way.

[9]Oswald (2008) has recently questioned the implicit assumption underlying some SWB work that the reporting function is linear. However, he goes on to present evidence from an experiment with self-reported heights, suggesting that linearity may in fact be a reasonable assumption after all.

[10]In fact, this result relies only on the assumption that the reporting function is a

Second, if all variation were due to ϕ_i, then this would imply that $\text{Var}(\psi_{pi}) = \sigma_p^2 = \text{Var}(\psi_p^* \phi_i) = \psi_p^{*2} \sigma_\phi^2$, and therefore that the absolute effect size $\frac{|\psi_p^*|}{\sigma_p} = \frac{1}{\sigma_\phi}$ is constant for all p. Moreover, for any elements p and q, we would have $\text{Cov}(\psi_{pi}, \psi_{qi}) = \text{Cov}(\psi_{pi}^* \phi_i, \psi_{qi}^* \phi_i) = \psi_p^* \psi_q^* \sigma_\phi^2$, which implies that the correlation between the two $\rho_{pq} = \text{sgn}(\psi_p^* \psi_q^*)$. Given consistent estimators of these quantities, both of these restrictions are testable.

Assuming an affine reporting function also allows us to recover the ratios of any two (non-zero) coefficients, ψ_{pit}^* and ψ_{qit}^*: because the ϕ_is will cancel $\frac{\psi_{pit}}{\psi_{qit}} = \frac{\psi_{pit}^* \phi_i}{\psi_{qit}^* \phi_i} = \frac{\psi_{pit}^*}{\psi_{qit}^*}$. The same is true for any function of the ratios, such as the consumption multiplier equivalent (CME_p): $\exp(\psi_{pit}/\psi_{cons,it})$.[11] To the extent that we are interested in quantifying the *intra*-personal trade-offs between various life events, it is ultimately the latter quantities that are of the most interest; and as such our inability to identify ϕ_i itself may not matter greatly.[12]

positive, monotonic transform of u_{it}^*.

[11] We introduced this measure in Chapter 2. See Section 2.5.

[12] In contrast, these quantities will not be sufficient if we are interested in identifying either interpersonal trade-offs, or the correlations between the size of particular coefficients and *levels* of SWB. For example, quantifying the *inter*-personal trade-offs between (say) more money for individual i and more money (or anything else) for individual j would involve a ratio of the form $\frac{\psi_{pit}^*}{\psi_{qjt}^*}$, which is not recoverable from $\frac{\psi_{pit}}{\psi_{qjt}} = \frac{\psi_{pit}^* \phi_i}{\psi_{qjt}^* \phi_j}$ unless we know the ratio of the individuals' reporting scales $\frac{\phi_i}{\phi_j}$. In fact, if we are willing to model heterogeneity in the reporting function more explicitly, there *is* a way to identify the ϕ_is separately, so as to facilitate making such trade-offs. Though we do not pursue this approach further here, we describe the methodology briefly in Appendix A.8. As usual, identification comes at the cost of additional assumptions, and in this case, also much greater computational complexity.

While violation of the restrictions set out above would allow us to identify true variation in the ψ_i^* coefficients, it is clear that it does not, without further assumptions, allow us to identify variation in the true α_i^* (the presence of the α_i intercept parameter prevents us from using the restrictions derived above). However, if we were willing to assume that the ϕ_i and α_i parameters of the reporting function are uncorrelated with both each other and the true ψ_i^* and α_i^*, then this would allow us to interpret estimated correlations between unobserved slope heterogeneity, ψ_i, and unobserved intercept heterogeneity η_i as reflecting the presence and sign of correlations between the true ψ_i^* and α_i^*.[13]

5.2.2 The Linear Mixed Model

To estimate equation (5.2), we adopt a random coefficient framework that explicitly models individual heterogeneity as a set of continuous random variables. This sort of approach is common in a number of literatures outside economics, including e.g. education research (see e.g. Goldstein et al., 1993; Goldstein, 1997) and medical science (e.g. Laird and Ware, 1982). Previous applications in the economics literature include health economics (see Rice and Leyland, 1996; Rice and Jones, 1997) willingness to pay studies (Train,

[13]To see this, note that under these assumptions, $\mathrm{Cov}(\eta_i, \psi_{pi}) = \mathrm{Cov}(\alpha_i, \psi_p^* \phi_i) + \mathrm{Cov}(\alpha_i^*, \psi_p^* \phi_i) = \mathrm{Cov}(\alpha_i^* \phi_i, \psi_p^* \phi_i)$, which will have the sign of $\mathrm{Cov}(\alpha_i^*, \psi_p^*)$ if ϕ_i is always positive.

1998; Amador, González, and Ortúzar, 2005; Hynes, Hanley, and Scarpa, 2008), and, in SWB research specifically, Lucas et al. (2003). This approach requires us to impose some form of parametric assumption on the structure of heterogeneity. The simplest of these, embodied in the Linear Mixed Model (LMM), is the assumption that variation is multivariate Gaussian.

As we have already alluded to, the LMM differs from the latent class approach taken by Clark et al. (2005) in two main ways: first, it models variation in coefficients of (5.2) as continuous, rather than as falling into a number of discrete groups; and second, that in order to do so, it imposes a Gaussian structure on that variation.[14]

Some have criticised latent class approaches on the basis that variation in coefficients is more likely to be continuous than discrete (e.g. Magder and Zeger, 1996), and from this perspective, the LMM approach seems slightly preferable. However, allowing for continuity comes at the cost of assuming a particular distributional form for the random coefficients, which may or may not be appropriate. Ultimately, the optimal trade-off between these concerns is an empirical question, and a comparison of the predictive ability of alternative models (along the lines of Hynes, Hanley, and Scarpa, 2008, who

[14]The approach we take here also differs from that used by Clark et al. (2005) in treating the responses to SWB questions as lying on a continuous linear scale. In general, the choice of whether to model such responses as ordinal or cardinal seems to have little impact on regression results (see e.g. Ferrer-i-Carbonell and Frijters, 2004) and greatly reduces the computational burden here.

compare the predictive ability of such models in the context of recreational demand models) could be a fruitful task for future research. Our adoption of the LMM is based on a number of factors.

The main reason is simply that the LMM approach most easily allows us to estimate multiple random coefficients while maintaining computational tractability.[15] Moreover, Verbeke and Lesaffre (1996a, 1997) have demonstrated that the LMM coefficient estimates are consistent even if the random coefficient distribution is misspecified. The LMM estimates therefore provide a convenient base from which to estimate the semi-parametric PGM models we employ in Section 5.2.3. Finally, obtaining similar results to Clark et al. (2005) using alternative methods should ultimately increase our confidence in the conclusions reached.

Methodology

This section sets out the estimation strategy for the LMM we use to identify individual heterogeneity in the determinants of SWB. The model we estimate is of the form:

$$u_{it} = \ddot{x}'_{it}\beta_{it} + \alpha_i + \varepsilon_{it} \tag{5.3}$$

$$\beta_{kit} = \beta_{k0} + \Gamma_k \ddot{z}_{it} + \omega_{ki} \tag{5.4}$$

[15]We ended up abandoning early attempts to estimate similar latent class models using Stata's **gllamm** package, due to the substantial computation time involved.

$$\alpha_i \quad = \quad \alpha_0 + \bar{\ddot{x}}_i'\theta + \eta_i \tag{5.5}$$

Equation (5.3) specializes (5.2), and is for the most part a standard SWB specification. As usual, u_{it} is the reported SWB of individual i at time t, and the $\ddot{x}_{it}$ are a set of (time-varying and time-invariant) covariates, with the dots indicating that their continuous components are centred on their grand means.[16]

In the estimation below, these x variables consist of log household consumption, log household income, reference income (the calculation of which is detailed below), health, labour force status dummies (employed, unemployed and looking for work, or unemployed and not looking for work, with out of the labour force as the reference category), relationship status dummies (married, cohabiting, separated, divorced, and widowed, with never married as the reference category), a dummy indicating whether the individual has children, hours spent working, hours spent socialising, hours slept, hours slept squared (with hours spent in passive leisure such as watching TV as the omitted category),[17] age, age squared, gender (male as the omitted category), household size, education (a set of dummies indicating the

[16] Dummy variables are left uncentred. Grand mean centring serves both to aid interpretation of the constant parameters of the model (see below) and also assists in achieving convergence of the restricted ML estimator (see e.g. Kreft, De Leeuw, and Aiken, 1995).

[17] For these measures we use the calibrated time-use data described in Kan and Gershuny (2006).

highest qualification the individual has obtained: either secondary, or post-secondary, with none as the omitted category), a dummy indicating whether the individual lives in an urban area,[18] and a set of wave dummies to control for common shocks.

In contrast to the standard approach of assuming the coefficients on these variables are identical for all individuals and time periods, we allow a subset β_{kit} of the β_{it} coefficients to vary. Equation (5.4) models these as depending on a fixed constant β_{k0} (which, because we have centred the continuous components of x_{it}, represents the average effect across the entire sample, for the reference category implied by the definition of any dummy variables), on a subset of the x variables, $\ddot{z}_{it}$, and on an individual-specific random component ω_{ki}.[19]

In practice, it is difficult to allow for heterogeneity in too many coefficients simultaneously; but in order to obtain accurate money equivalents we will generally need to allow for heterogeneity in at least the coefficients whose ratios we wish to compute (as well as allowing for individual-specific constants). If we do not, any such heterogeneity will be absorbed into the (unestimated) ϕ_i, and will misleadingly cancel out in the ratio calculations. In what follows we estimate random coefficients on seven of the x variables

[18] Defined as a settlement of with 10,000 people or more.

[19] The notation in (5.4) stacks these k individual equations in the standard way, with Γ as a stacked matrix of the equation-specific γ_k's.

described above: log household consumption, reference income, health, unemployment (either looking for work, or discouraged), being in a relationship (either married or cohabiting), having children, and hours worked. For all other variables f in x_{it}, we restrict $\beta_{fit} = \beta_{f0}$ (i.e. $\gamma_f = \mathrm{Var}(\omega_{ki}) = 0$).

For each of the x_{kit} variables, we include a number of demographic interactions, to allow the effect of x_{kit} to vary with: age, age squared, gender, urban residence, and education. We also include additional interactions for the following coefficients. For reference income, we include an interaction to examine whether reference effects vary depending on whether an individual is employed. For unemployment, we include an interaction with the discouraged variable to see whether the effect varies depending on whether individuals are looking for work. We also allow the reference categories to vary; the main effect compares unemployment to being out of the labour force, but we also want to compare with the employed. We interact being in a relationship with being married, to see whether the formal status of the relationship matters; and also allow for comparisons to separation, divorce, and widowhood (as well as the baseline category of never married). Finally, we allow the effect of having children to vary depending on the number of children, including interactions with dummies for two, three, and four or more children.

The individual-specific intercepts are modelled in a similar fashion in equation (5.5). They depend on a fixed constant α_0, time-invariant covariates $\bar{\ddot{x}}_i$, and an individual-specific random component η_i. The individual-level means $\bar{\ddot{x}}_i$ are included in the equation to account for possible bias due to a correlation between the x_{it} variables and the error term ε (see e.g. Bafumi and Gelman, 2006). The rationale for this is identical to that underlying the Mundlak decomposition approach taken e.g. by Senik (2004): including the individual-level means of the x variables removes any correlation between the individual-level errors η_i and the time-varying x_{it}. Any such correlation (say, due to omitted personality traits that tend to result in both high income and greater life-satisfaction) will be absorbed into the θ coefficients (which of course suggests that we should be cautious in assigning any causal interpretation to them).[20]

Combining equations (5.3), (5.4), and (5.5) yields the full equation:

$$u_{it} = \ddot{x}'_{it}\beta_0 + \ddot{x}'_{kit}\Gamma_k \ddot{z}_{it} + \ddot{x}'_{kit}\omega_{ki} + \alpha_0 + \bar{\ddot{x}}'_i\theta + \eta_i + \varepsilon_{it} \tag{5.6}$$

To estimate the equation, we assume the errors ε_{it} are (conditionally) normal with mean zero and variance σ^2, and the random effects, ω_i and η_i are (conditionally) multivariate normal, with mean zero and (unrestricted) covariance

[20] After some experimentation, we elected to constrain θ to zero for the time dummies and the age and age squared variables, as individual means of these seemed to lack a plausible connection to the errors.

matrix Ω. We then estimate these variance components (σ^2 and Ω) jointly with the main model parameters, using Stata's **xtmixed** command.

Calculating Reference Income

To calculate the reference income, we use a wage equation approach based on that employed by Clark and Oswald (1996) on the first wave of the BHPS (and Senik, 2004, in the RLMS). For each wave, we estimate a Heckman selection model to calculate predicted wages based on age, age squared, tenure, tenure squared, and dummies for health status, region, occupation, industry, education, working hours, organisation type, type of contract, incentive pay, full time vs. part time status, managerial responsibility, and whether they have a pension. The selection equation includes non-labour income, age, age squared, dummies for children in different age groups, marital status, health, region, and education (with children, marital status, and non-labour income as the exclusion restrictions). To allow the determinants of participation and income to vary by gender, we fit separate equations for males and females. The results of these equations can be found in Appendix A.6.[21]

[21] We experimented with a variety of alternative methods for calculating the reference level, including combining males and females, using alternative independent variables, estimating a multinomial selection model to distinguish between different types of non-employment (see e.g. Bourguignon, Fournier, and Gurgand, 2007), as well as calculating cell means of individual income, household income and household consumption. None of these performed especially well in terms of attracting significant coefficients, so we elected to retain the traditional Clark and Oswald (1996) specification.

Results

The results of our LMM regression are displayed in Tables 5.2 and 5.3. Table 5.2 shows the main results of interest (i.e. for the random coefficients and interactions). The top panel shows the fixed part of the equation for each coefficient, which is composed of the main effect β_{k0}, and the interactions γ_{kz}.

Most of the standard controls employed in SWB regressions are involved in either the random coefficients or their interactions. However, the coefficients on the remaining controls (shown in column 1 of Table 5.3) have roughly the expected signs. SWB is u-shaped in age (with a minimum around 38), decreases in household size, and is also lower for those in urban areas.

Turning to the main coefficients of interest (in Table 5.2) the first, and most obvious point to note is that only some of the variables have significant effects on SWB *on average* (for the reference category of rural males with no formal qualifications). These include consumption, health, unemployment, and being in a relationship, all of which have the expected sign. While the coefficient on hours worked is insignificant, and (very) slightly positive, it is important to remember that this is relative to hours spent on the omitted category: passive leisure (e.g. TV). As we can see from the coefficients on the other time-use categories further down, hours spent working do have

Table 5.2: Heterogeneity in the Determinants of SWB in the BHPS

	Consumption		Reference Income		Health		Unemployment		In a Relationship		Kids		Hours Worked	
Fixed Part														
Main effect	0.094***	(0.023)	-0.000	(0.010)	0.189***	(0.008)	-0.393***	(0.056)	0.156***	(0.048)	0.034	(0.033)	0.017	(0.023)
Interactions w/:														
Employed			0.028	(0.023)										
Discouraged							0.225***	(0.060)						
vs. Employed							0.043	(0.032)						
Married									-0.047	(0.030)				
vs. Widowed									0.248***	(0.075)				
vs. Separated									0.341***	(0.066)				
vs. Divorced									0.098	(0.063)				
2 Children											0.060**	(0.024)		
3 Children											0.077*	(0.043)		
4+ Children											-0.012	(0.092)		
vs. Social Leisure													-0.156***	(0.028)
vs. Hours Slept													-0.951***	(0.221)
vs. Hours Slept^2													0.052***	(0.011)
Sex	0.000	(0.027)	0.010	(0.015)	-0.020	(0.014)	0.021	(0.061)	-0.054	(0.034)	0.012	(0.031)	0.027***	(0.009)
Age	-0.006	(0.004)	-0.004*	(0.002)	-0.002	(0.002)	-0.026**	(0.010)	0.008	(0.006)	-0.007	(0.007)	-0.010***	(0.001)
Age^2/10	0.000	(0.000)	0.000*	(0.000)	0.000	(0.000)	0.003**	(0.001)	-0.001*	(0.001)	0.000	(0.001)	0.001***	(0.000)
Urban	0.091***	(0.028)	0.005	(0.015)	0.012	(0.015)	-0.077	(0.068)	-0.002	(0.035)	0.032	(0.032)	-0.001	(0.008)
Secondary	-0.012	(0.037)	0.019	(0.017)	-0.017	(0.017)	-0.147**	(0.072)	-0.030	(0.046)	-0.024	(0.042)	-0.006	(0.010)
Post-Secondary	-0.014	(0.042)	0.034	(0.021)	-0.042**	(0.021)	-0.213**	(0.094)	-0.036	(0.052)	0.051	(0.047)	-0.013	(0.012)
Random Part														
Standard Deviation	0.370***	(0.030)	0.223***	(0.010)	0.259***	(0.010)	0.736***	(0.038)	0.638***	(0.027)	0.208***	(0.055)	0.107***	(0.009)
Correlations w/:														
Reference Income	-0.133	(0.106)												
Health	-0.235***	(0.083)	0.096	(0.066)										
Unemployment	-0.003	(0.130)	-0.327**	(0.122)	0.074	(0.093)								
In a Relationship	-0.135*	(0.075)	0.167**	(0.075)	0.041	(0.061)	0.004	(0.089)						
Has Kids	-0.138	(0.230)	0.040	(0.246)	0.233	(0.181)	-0.251	(0.296)	0.054	(0.164)				
Hours Worked	0.123	(0.113)	-0.496***	(0.090)	0.030	(0.083)	0.237	(0.134)	-0.157**	(0.076)	-0.610*	(0.230)		
Intercept	-0.371***	(0.043)	0.001	(0.042)	-0.410***	(0.033)	0.027	(0.057)	0.146***	(0.029)	0.129	(0.103)	-0.132***	(0.043)
Observations	55794													
Groups	11360													
Log Likelihood	-77889.996													

* $p<0.10$, ** $p<0.05$, *** $p<0.01$. Standard errors in parentheses.

Table 5.3: Heterogeneity in the Determinants of SWB in the BHPS

	Main		Intercept†	
(Log) Household Income	0.003	(0.008)	0.044**	(0.018)
Household Size	-0.070***	(0.010)	0.026	(0.016)
Urban	-0.103***	(0.033)	0.058	(0.037)
Secondary	0.128	(0.114)	-0.280**	(0.116)
Post-Secondary	-0.009	(0.123)	-0.208*	(0.125)
Age	-0.030***	(0.005)		
$\text{Age}^2/10$	0.004***	(0.001)		
Sex			0.051	(0.032)
(Log) Household Consumption			-0.071**	(0.031)
Reference Income			-0.081***	(0.018)
Health			0.379***	(0.014)
Employed			-0.082	(0.056)
Unemployed			-0.367***	(0.128)
Unemployed (Discouraged)			-0.381**	(0.150)
In a Relationship			0.198***	(0.056)
Married			0.128***	(0.044)
Widowed			0.270***	(0.082)
Divorced			-0.028	(0.075)
Separated			-0.102	(0.109)
Hours Social Leisure			0.078*	(0.045)
Hours Worked			0.067*	(0.039)
Hours Slept			-0.510	(0.332)
Hours Slept^2			0.040**	(0.016)
Has Children			-0.078*	(0.043)
2 Children			-0.043	(0.046)
3 Children			0.038	(0.073)
4 Children			0.150	(0.162)
Constant			5.304***	(0.017)
Observations	55794			
Groups	11360			
Log Likelihood	-77889.996			

* $p<0.10$, ** $p<0.05$, *** $p<0.01$. Standard errors in parentheses. † Coefficients in the intercept column are for time-invariant individual-level means, $\bar{\bar{x}}_i$. Wave dummies included but not reported.

a significantly negative effect relative to both social leisure (e.g. time spent with friends) and sleep (at least over a certain range: the significant quadratic term on hours slept implies an optimum level of sleep at around 9 hours and

10 minutes (s.e. 25 minutes), and the marginal impact of trading more work for less sleep is negative up to 9 hours). In contrast, the average effects of changes in reference income and of having children are not significantly different from zero.

The observable correlates of heterogeneity

Turning to the observable correlates of *variation* in the impact of these outcomes, we see that the demographic variables we include as interactions do indeed have some predictive ability. While the patterns vary for the difference effects, every demographic variable is a significant predictor of variation in at least one coefficient.

We focus first on what we might think of as "materialistic" coefficients: consumption and reference income. The importance of consumption is nearly twice as large for those who live in urban areas rather than rural ones, but does not appear to vary with age, gender or education. The effect of the reference income, meanwhile is u-shaped in age, reaching a minimum of -0.201^{*} (s.e. 0.119, and significant at the 10% level) at around 52.

We turn next to consider the effect of the less obviously materialistic variables: health, unemployment,[22] relationships, children, and hours worked

[22] While it might be tempting to consider the importance of employment a "materialistic" concern, recall that what this coefficient measures is the impact of unemployment that is independent of the typical materialistic concerns that might motivate it (i.e. the consumption and income it allows). It thus seems reasonable to classify the residual effect

(which is essentially the inverse of leisure). We see that importance of health appears to decline with education (though the effect is fairly small relative to its overall positive effect, and it could potentially be due to more educated individuals having higher standards). Meanwhile, the impact of unemployment is less than half as bad for those who are only marginally attached to the labour force than for those actively looking for work (-0.168 vs. -0.0393), and also shows significant variation with age and education: unemployment is worse for the more highly educated, and in age terms reaches its nadir (-1.673***, s.e. 0.507) around 51. The findings in relation to the marginally attached seem broadly consistent with the evidence on habituation to unemployment presented in Clark, Georgellis, and Sanfey (2001), but while they estimated separate equations for males and females, we find no significant difference between the two. In contrast, the effect of hours worked is the one variable that does vary significantly between genders, with women being slightly less negatively affected by it.

The effects of various types of relationship exhibit fairly standard patterns: there is no significant difference between the cohabiting and the married, and those in a relationship are much happier than the widowed and separated; interestingly, however, we do not find a significant difference be-

as non-materialistic. Indeed, some have suggested that a major reason for the negative impact of unemployment is could be the reduced social contact it entails (though our time use controls should mitigate the extent to which it picks up such factors here).

tween the divorced and the never married. There is some, albeit weak evidence of an n-shaped age pattern, with a peak of 0.429 (s.e. 0.231) at around 36. The impact of having children appears to increase until 3 children, and decline thereafter. However, it is not significantly related to any of the other demographic characteristics. The coefficient on hours worked also follows the now familiar u-shaped pattern in age, reaching a minimum (relative to social leisure) of -0.323 (s.e. 0.080) at around 48. It is interesting to note that the age pattern in the effect of reference income, unemployment, and hours worked is quite similar: the negative effects of all three are all greatest around the age of 50.[23]

Finally, we turn briefly to the coefficients defining the intercept in Table 5.3. If these could be interpreted causally, then they could be thought of as capturing the longer-term impact on SWB of the relevant outcomes — rather than the shorter-term effects reflected in the coefficients on the time-varying mean-deviations. However, as we noted above, it is more difficult to interpret the coefficients on the time-invariant $\bar{\bar{x}}_i$ variables in this way, due to possible correlations with unobserved personality traits.

For this reason, we shall not focus greatly on these coefficients, other than to draw brief attention to those that differ from the coefficients on the

[23]This may go some way towards explaining the age-patterns in SWB discussed e.g. by Easterlin (2004) and Blanchflower and Oswald (2008), although the main coefficients indicate that there remains a significant residual u-shape in age even once life-cycle changes in values are taken into account.

corresponding time-varying effects. In contrast to its main effect, mean income is positively associated with SWB (consistent with the contention that it may pick up the effect of unobserved personality traits). On the other hand higher mean consumption is associated with greater *dis*satisfaction, perhaps reflecting a negative effect of materialistic values or adaptation to higher consumption levels. The coefficient on mean reference income is also significantly negative. One possible interpretation of this is that expectations tend to be relatively fixed, and adjust only slowly to changes in others' outcomes. Interestingly, those who spend longer being widowed tend to be happier, potentially reflecting adaptation to widowhood (the effect is of a similar magnitude to the negative time-varying effect). Finally, those who have children for a longer portion of the sample appear to be significantly less happy than those who do not.

The scale of unobservable heterogeneity

The bottom panel of Table 5.2 shows the standard deviations and correlations of the random coefficients. As discussed above, these represent the residual heterogeneity in the coefficients that we are unable to explain via the demographic characteristics discussed in the previous paragraphs. Analysis of these results suggests two main conclusions.

First, there is indeed significant unobserved heterogeneity in the effect of all the variables for which we allow random coefficients. The likelihood ratio test against the simple linear regression with no random components (which is conservative in this case[24]) conclusively rejects the restriction that all the variance components are zero (χ^2_{36}-statistic = 16022, with a corresponding p-value numerically indistinguishable from 0). While we have not conducted likelihood ratio tests against all possible nested models, the variance components are all quite precisely measured, and in only one case — the effect of having children — is the lower bound of the 99% confidence interval of the estimated standard deviation less than 80% of the point estimate (in this case, it is 0.104, or 50% of the point estimate of 0.208); and even at that lower bound, the standard deviation of $\beta_{children}$ is still close to three times its average effect of 0.034.

As an indication of the overall amount of unexplained heterogeneity in the coefficients, note that the estimated standard deviation for every coefficient is larger than the point estimate on the log of household consumption, and all but one (the standard deviation of β_{work}) is larger than the effect of a one unit change in health satisfaction. This means, for example, that we would predict a greater difference in SWB between individuals a standard deviation apart in their $\beta_{child,i}$ coefficients (assuming both have children) than we would

[24]The true-distribution is not the assumed χ^2_{21} because of rather complicated restrictions on the variance components, but the tails are bounded above by this distribution (see Stram and Lee, 1994; McLachlan and Basford, 1988).

between two individuals a unit apart in subjective health satisfaction.

To illustrate the extent of variation in the coefficients graphically, Figure 5.2 displays kernel density estimates of the (marginal) distributions of the observed, unobserved, and total variation in each of the coefficients. In most of the panels, we see that while the unobserved variation density (dotted line) is quite flat, the observed variation density (solid line) is very narrow, indicating that there is much more unobserved than observed variation. This is confirmed when we look at the total variation density (dashed line) which is often visually indistinguishable from unobserved variation. Looking at the proportions of observed variance, we se that in only three cases — hours worked, children, and the intercept — can we observe more than 10% of the variation in the coefficients.

Another way to think about the amount of heterogeneity implied by these estimates is to consider the implied proportion of individuals for whom the effect of a given outcome is the opposite of its main effect. Looking at Figure 5.2, it appears that a number of the coefficients have a substantial proportion of their mass on the opposite side of zero from their average: 40% for consumption, 49% for reference income, 28% for unemployment; 11% for hours worked; 24% for health; 42% for relationships; and 44% for children.[25]

[25] Of course, simply because there are a large proportion of oppositely-coefficients, does not necessarily mean that heterogeneity is substantively important. It could be that both the average effect *and* the variation around it are small. But as we have already noted above, that does not generally seem to be the case here.

Figure 5.1: LMM Estimates of heterogeneity in the determinants of SWB

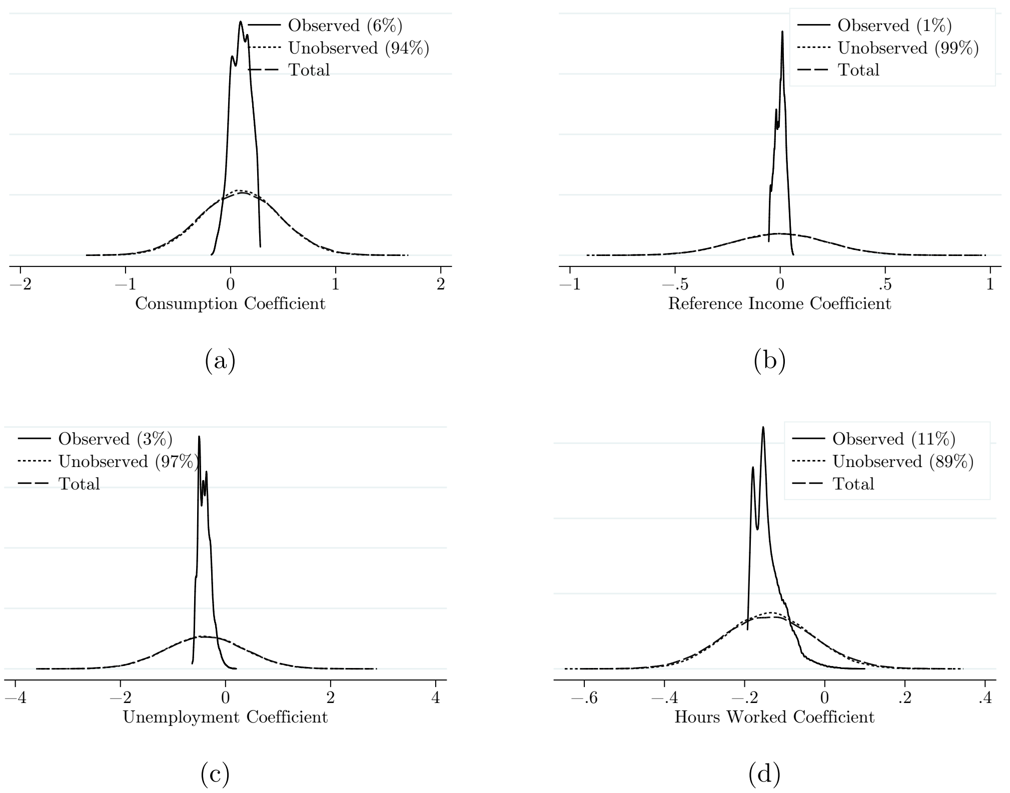

It seems reasonable to think that some of these outcomes could affect different individuals in different directions: for reasons discussed at length in Part II there are reasons to think that the effect of reference income could be ambiguous depending on whether individuals treat it primarily as a signal or as a status indicator; and there is little *a priori* reason to think that having children, or being in a relationship should have uniformly signed effects; even unemployment and hours spent working could plausibly have opposite effects for different people, depending on how much they enjoy their work, and how

Figure 5.2: Heterogeneity in the determinants of SWB (continued)

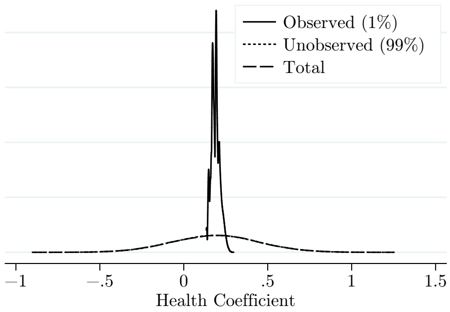

(e)

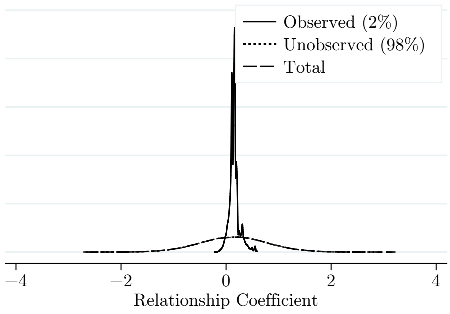

(f)

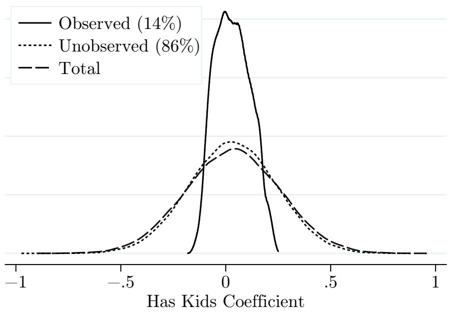

(g)

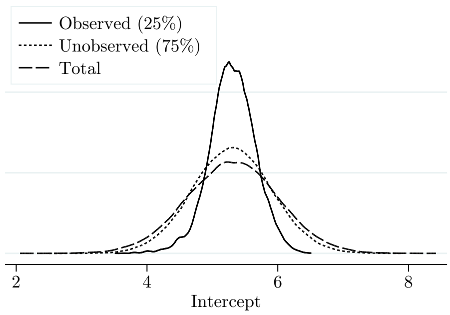

(h)

Kernel density estimates of the distribution of individual coefficients on (a) Log Household Consumption, (b) Reference Income, (c) Unemployment, (d) Hours Worked, (e) Health, (f) Being in a relationship, (g) Having Children, and (h) of the distribution of the individual intercepts. β_i coefficients for each variable are on the horizontal axis, while the vertical axis represents the density. The distribution of *observed* variation is the distribution of predicted β_i coefficients based on the interactions in the top panel of Table 5.2 (for the random coefficients) and the variables in column (2) of Table 5.3 for the intercept. The baseline categories are: “out of the labour force” for unemployment; “never married” for being in a relationship, and “hours of social leisure” for hours worked. The *unobserved* variation is simulated from the estimated multivariate normal distributions describing the random coefficients (centred at the mean of the fixed distribution). The *total* distribution is the sum of these two distributions. Where the observed variation is small relative to the unobserved variation, the unobserved and total variation densities will be very similar (and are therefore often difficult to distinguish).

much fulfillment they get out of the available alternatives. Nonetheless, it might initially be more difficult to imagine how greater consumption, and especially improved health could have a negative effect on SWB. (In particular, it seems implausible that this apparent negative relationship could hold over the entire range of outcomes — which would, after all, imply that individuals would be happiest with no consumption, or in the lowest category of health.)

One possibility is that the implied negative effects depend on exactly what we have held constant in our estimation. In particular, if increases in consumption or health come at a cost, have other negative side-effects, or result in changes in values that we have not modelled explicitly, then it is entirely possible that they could have a *net* negative effect for some individuals. Possible omitted variables are probably easier to imagine for pecuniary outcomes than for health: in line with our results in Chapter 3 debt financing of consumption is one possibility. But changing expectations could easily apply to either: if increases in consumption or health increase expectations (which we have left unmodelled here for simplicity) then this could well result in an overall decrease in satisfaction, at least for some individuals. If this were indeed the case, then it could be that the answer to the question of whether money buys happiness or simply breeds dissatisfaction is in fact *both*.

As we noted in Section 5.2.1 above, the existence of individuals whose SWB responses to the same outcome are oppositely signed — in addition to being substantively interesting in its own right — also provides evidence against the idea that variation in the mapping from objective outcomes to reported SWB could be due primarily to variation in the reporting function, rather than in their "true" effects. After all, we would expect the reporting function to be monotonically increasing in true SWB.

An alternative possibility, however, is that switches in sign are simply artefacts of the normality assumption imposed by the LMM, and would not survive abandoning it. While Verbeke and Lesaffre (1996a, 1997) have shown that estimates of both the mean (which depends on the fixed coefficients) and the variance of the coefficient distributions are robust to such misspecification, these do not (without the added normality assumption) dictate that a specific proportion of the population should have oppositely-signed coefficients. Rather than a normal distribution, the true distribution could instead be, say, a truncated normal distribution, with the same mean and variance but with no coefficients below zero, and this would be equally consistent with the evidence. This might seem less plausible in the case of reference income (where the main effect is close to zero, and the standard deviation of the coefficients is much larger), but it nonetheless remains a possibility. We address this concern in more detail in Section 5.2.3.

The structure of unobservable heterogeneity

The pattern of correlations between the random coefficients also display some interesting properties. One hypothesis here is that individual values may cluster around a "materialism" axis. If this were the case, then we would expect those who especially value consumption to also respond more negatively to reference income, but to respond less to other outcomes (i.e. less negatively to unemployment and hours worked, and less positively to health, relationships and children). As a result, we would expect to see positive correlations between reference income, health, relationships, and children; positive correlations between consumption, unemployment, and hours worked; and negative correlations between the two groups.

In fact, while this pattern does not fit exactly, and many of the correlations are insignificant, we do see something resembling these predictions. Out of 21 total coefficients, only five are incorrectly signed;[26] while of the seven that are statistically significant, all have the predicted signs. To investigate the underlying structure of the unobserved variation more formally, we conduct a factor analysis on the simulated distributions of the random portion of both the β_{ki} coefficients (i.e. the ω_{ki}s from equation (5.4)) and the

[26]Four of the exceptions involve correlations with β_{unemp} coefficient — namely β_{cons}, β_{health}, $\beta_{relationship}$, and β_{work}. The other is the correlation between β_{work} and β_{health}. In all cases, the correlations are small and insignificant.

intercept, η_i.[27] A single principal factor emerges from this analysis (using the retention criterion that the eigenvalue be greater than 1), and the factor loadings on the β_{ki} coefficients (displayed in Table 5.4) indeed map neatly onto the general concept of "materialism": the principal factor has positive loadings on consumption, unemployment, and hours of work, and negative loadings on reference income, health, relationships, and children (though the loading on health is quite small).

Table 5.4: Factor Loadings for Unobserved Heterogeneity in the Determinants of SWB

	Principal Factor
$\omega_{cons,i}$	0.221
$\omega_{ref,i}$	-0.460
$\omega_{health,i}$	-0.069
$\omega_{unemp,i}$	0.326
$\omega_{relationship,i}$	-0.190
$\omega_{kids,i}$	-0.559
$\omega_{work,i}$	0.972
η_i	-0.164

Factor loadings for the principal factor extracted from the simulated distribution of the random portion of the β_{ki} coefficients, and the random intercept, η_i. Estimated using the iterated principal factor method. Only one factor was retained with an eigenvalue greater than 1.

Finally, it is also interesting to note the correlations between the random coefficients and the random intercept. In general, greater materialism seems

[27] Intuitively, factor analysis seeks to identify common variation in a set of variables that can be attributed to one or more underlying factors, each composed of linear combinations of the original variables. As employed here, it is intended merely as a descriptive aid, and is not intended to uncover causal structure. See e.g. Gorsuch (1983) for an overview.

to be significantly related to lower (unobserved) SWB: those who value consumption more, are more negatively affected by unemployment, and are less negatively affected by working longer hours, also tend to be significantly less happy than we would expect, given their observable circumstances. This general pattern is also confirmed by the factor analysis, which loads the intercept negatively into the "materialism" factor.[28] This comports with previous findings of a negative relationship between materialism and SWB, but in contrast to those results — which are based on self-reported measures of materialism — we have identified materialism directly via estimates of individual-level SWB coefficients. For this reason, our findings cannot be explained by the idea that materialistic individuals are simply pursuing goals that are fundamentally unsatisfying (as suggested by e.g Van Boven, 2005). Instead, it may simply be that materialistic individuals' higher aspirations make them less satisfied in any given set of circumstances (see e.g. Sirgy, 1998; Sirgy et al., 1998).[29]

As we noted earlier, equity concerns may mean we are particularly interested in the factors that yield the highest impact for those who are least

[28] On the other hand, when we look at the individual correlations between the β_{ki}s and η_is, there is one notable exception to this general pattern: a significant negative correlation between the intercept and β_{health}, which suggests that those who place greater importance on health are also more difficult to satisfy.

[29] These findings are also more robust to concerns that individuals may misstate their values — concerns which seem especially relevant in light of the results of Chapter 6, which suggest that individuals frequently do a poor job of predicting what will bring them satisfaction.

satisfied with their lives. The correlations between the β_{ki} coefficients and the random intercept in the bottom panel of Table 5.2 suggest that those who are unobservably least-satisfied gain more from consumption and health (the random coefficients for which both correlate negatively with η_i), but gain less from relationships and social time ($\omega_{relationship,i}$ correlates positively with η_i, while $\omega_{workhours,i}$ correlates negatively).

However, these correlations do not take into account either the interaction effects in the top panel of Table 5.2, or the other determinants of SWB. If, for example, those who most value consumption have a low η_i, but are also rich, then they may actually have reasonably high SWB. To see what is really most important to the SWB-poor, we therefore estimate the correlations between reported SWB and Empirical Bayes predictions of the individual-level β_{ki}s.[30]

These correlations are displayed in Table 5.5, and appear to confirm not only that consumption and health have a greater impact for the SWB-poor (i.e. their coefficients have sizeable and significant negative correlations with SWB), but that the same is true, albeit to a smaller extent for unemployment (the positive correlation indicates that the SWB-rich are less negatively affected by unemployment). They also confirm that that the SWB-poor re-

[30]Such EB estimates are the standard method of predicting individual coefficients from a multilevel model, and are available even when the number of observations for each individual is too small to run separate regressions for each. Compared with maximum likelihood estimates of the individual coefficients, the EB predictions are "shrunken" towards the sample average coefficient, making them biased, but more accurate. See e.g. Rabe-Hesketh and Skrondal (2008) for details of their calculation.

spond significantly less positively to relationships, social time, and having children than do the SWB-rich. It is unclear at this stage whether it is having low SWB that results in individuals gaining little from such "social" outcomes, or instead whether gaining little from such outcomes leads to low SWB.

Table 5.5: Correlations between reported SWB and Empirical Bayes estimates of the β_{ki}s

	Correlation
$\beta_{cons,i}$	-0.445***
$\beta_{ref,i}$	-0.007
$\beta_{health,i}$	-0.445***
$\beta_{unemp,i}$	0.084***
$\beta_{relationship,i}$	0.139***
$\beta_{kids,i}$	0.048***
$\beta_{work,i}$	-0.083***

* p<0.10, ** p<0.05, *** p<0.01.

As we noted in Section 5.2.1 above, the other interesting facet of the correlations between the coefficients on different outcomes relates to the issue of heterogeneity in the reporting function. As we demonstrated there, if *all* the variation we have identified were due merely to (affine) heterogeneity in the reporting function, then (a) there should be perfect correlation (either positive or negative) between the coefficients; and (b) the effect sizes for each variable should be constant. From Table 5.2 and our discussion so far, it seems apparent that the correlations are far from perfect, and as Table 5.6

demonstrates, we can also reject the hypothesis that the absolute effect sizes are constant. Of 21 pairwise comparisons, 12 reject the equality hypothesis at (at least) the 5% level, and the overall hypothesis (that all the effect sizes are equal) is also conclusively rejected.

Table 5.6: Effect Sizes

	β_{cons}	β_{ref}	β_{health}	β_{unemp}	β_{rel}	β_{kids}	β_{work}
$\frac{\lvert\beta_{k0}\rvert}{\sigma_k}$	0.254	0.000	0.728	0.533	0.244	0.161	0.155
	(0.065)	(0.045)	(0.042)	(0.081)	(0.076)	(0.166)	(0.215)
P-values for equality tests: $\frac{\lvert\beta_{k0}\rvert}{\sigma_k} = \frac{\lvert\beta_{l0}\rvert}{\sigma_l}$							
Reference Income	0.001						
Health	0.000	0.000					
Unemployment	0.008	0.000	0.032				
Relationship	0.920	0.006	0.000	0.009			
Children	0.601	0.349	0.001	0.043	0.650		
Hours Worked	0.656	0.471	0.009	0.101	0.691	0.980	

Top panel displays absolute effect sizes (standard errors in parentheses). Bottom panel displays p-values for the tests that the corresponding pairs of absolute effect sizes are equal. 12 out a possible 21 tests are rejected at the 5% level, indicating that the effect sizes are not equal.

Heterogeneity in consumption equivalents

Nonetheless, while these tests allow us to soundly reject the hypothesis that heterogeneity in the reporting function is *entirely* responsible for our results, we still cannot rule out the idea that it may play a substantial role in our results, and there remains a legitimate question about how much true variation would remain once we took account of it. As we have already noted, one way to proceed here would be to model both "true", and reporting heterogeneity explicitly (see Appendix A.8). However, owing to the computational difficulty of such an approach, we pursue a simpler strategy here: by focusing on variation in the consumption multiplier equivalents (CMEs) of various outcomes, rather than variation in their coefficients directly. As we pointed out in Section 5.2.1, the distribution of these equivalents does not depend on the scale of individuals reporting functions, and thus, if they exhibit substantial heterogeneity, we can be confident that such heterogeneity is real.

Below, we simulate the distribution of (CMEs) for reference income, unemployment, hours of work, health, being in a relationship and having children.[31] The CME for a variable k is defined as $\text{CME}_k = \exp(\beta_k/\beta_{consumption})$, and reflects the amount by which consumption would need to be multiplied to have an equivalent effect to a one unit increase in k.[32] The possibility that some individuals may have negative coefficients on log consumption complicates this exercise somewhat, so in what follows, we focus primarily on the CMEs for those with positive values of β_{cons}.[33] A CME_k of 1 indicates that

[31] We ignore income in these calculations, as its effect is so small compared with the other variables, but incorporating it would not greatly affect the results.

[32] As we noted in Chapter 2, this seems more accurate than calculating currency-denominated money equivalents, if we take the log functional form seriously.

[33] CMEs are undefined for $\beta_{cons} = 0$, because when $\beta_{cons} = 0$ there is no change in consumption that could offset the effect of a change in another outcome.

For $\beta_{cons} < 0$, consumption equivalents still represent the level at which an individual is indifferent between the amount of consumption and a unit change in the equivalised variable. However, the direction of preference is reversed: that is, they would prefer consumption to be multiplied by *less than* the CME, rather than more. We have calculated the CMEs for $\beta_{cons} < 0$, but they do not suggest substantive conclusions different from those for $\beta_{cons} > 0$, and we omit them for simplicity.

k has no effect on SWB. For $\beta_{cons} > 0$, $\text{CME}_k < 1$ corresponds to k having a negative effect, and $\text{CME}_k > 1$ corresponds to a positive effect. The opposite is true for negative values of β_{cons}.

Figure 5.2 displays the distributions of the CMEs for each variable on a logarithmic scale. Because the densities are heavy-tailed, their means and variances approach infinity, and are not especially informative. Table 5.7 therefore presents alternative measures of central tendency and dispersion, that are based on the median, and median absolute multiplier of the median (MAM) (this is just the common median absolute deviation from the median in a logarithmic scale). Dividing the median by the MAM is then roughly analogous to subtracting one standard deviation from the mean for a normal distribution, and multiplying the median by the MAM is analogous to adding one standard deviation.[34] We also display the proportion of the population that has an oppositely signed coefficient to the median (i.e. on the opposite side of the neutral value of 1).

As we can see, there remains significant heterogeneity even once we factor out heterogeneity in the reporting function (if there were no remaining heterogeneity, then the graphs in Figure 5.2 would reduce to vertical lines at the median). Even for the lowest variation density — the CME of hours

[34]More precisely, they correspond to the median +/− 1 mean absolute deviation in a log scale.

Figure 5.2: Heterogeneity in Consumption Multiplier Equivalents

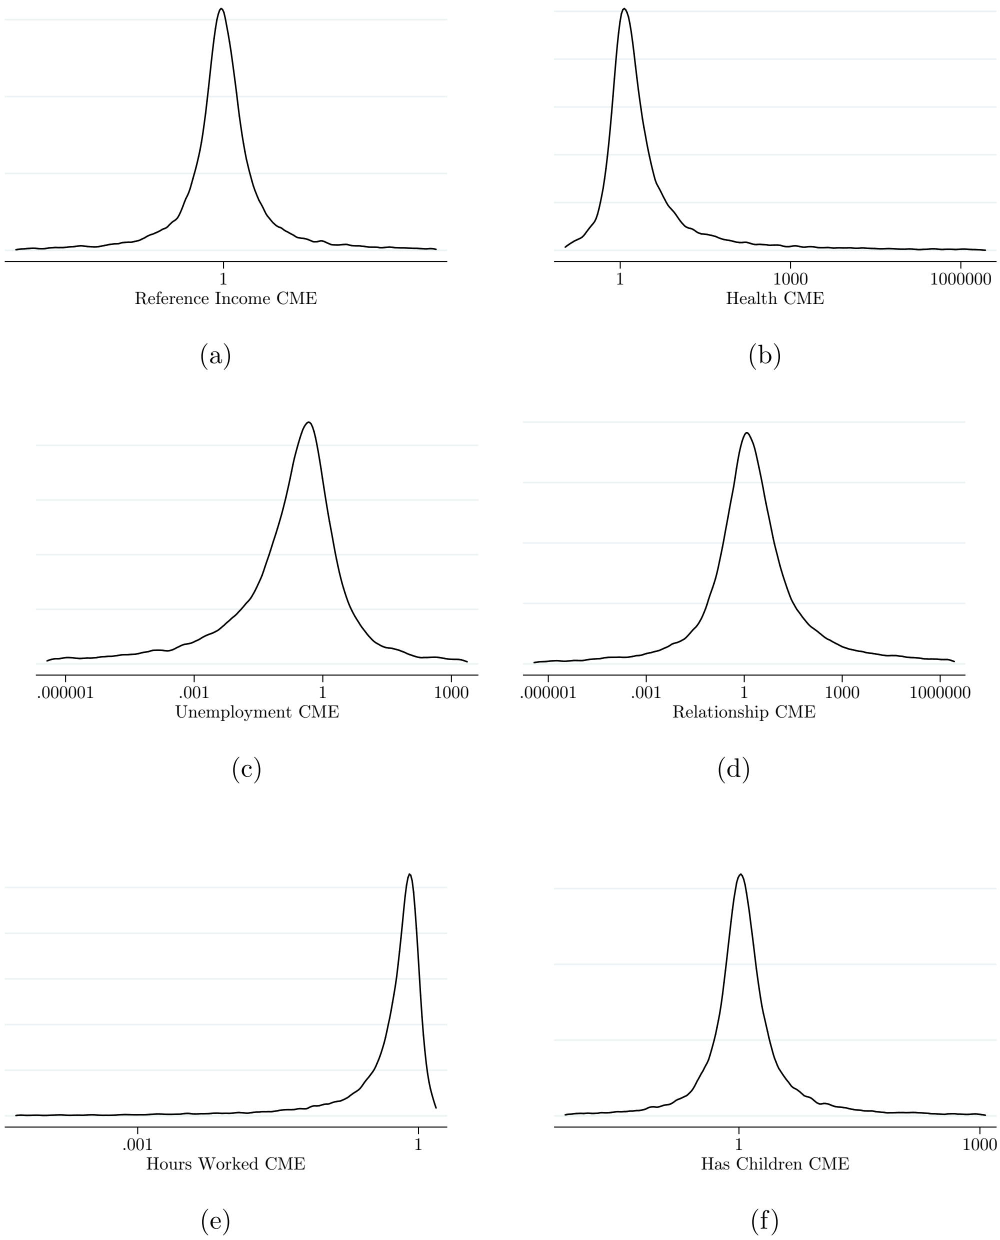

Kernel density estimates of the distribution of consumption multiplier equivalents for (a) reference income, (b) health, (c) unemployment, (d) being in a relationship, (e) hours worked, and (f) having children. The densities are truncated on the left at the maximum of the 2.5th percentile or 10^{-6}, and on the right at the minimum of the 97.5th percentile or 10^{6}, and the multipliers are displayed on a log scale.

Table 5.7: Location and Dispersion of Consumption Multiplier Equivalents

CME	Median	MAM[1]	Med/MAM[2]	Med×MAM[3]	Opposite Sign[4]
Reference Income	.963	1.668	0.577	1.606	48%
Unemployment	.260	4.745	0.054	1.233	23%
Hours Worked	.665	1.392	0.477	0.927	12%
Health	1.582	1.920	0.824	3.039	28%
Relationship	1.593	5.631	0.283	8.971	42%
Has Children	1.094	1.653	0.662	1.809	44%

[1] Median Absolute Multipliers of the the Median. This is equivalent to the Median Absolute Deviation from the Median (MAD) on a log scale. (For a log-normal distribution MAM/σ = $2/\pi$.) [2] Equivalent to a the Median minus 1-MAD, on a log scale. [3] Equivalent to a the Median plus 1-MAD, on a log scale. [4] I.e. on the other side of CME=1 Distributions for individuals with $\beta_{cons} > 0$ only.

worked — the median variation around the median CME involves a 30% or 40% change (for reductions $\approx 1 - \frac{1}{1.4}$ and increases $\approx 1.4 - 1$ respectively). Meanwhile, relationships and unemployment exhibit much higher variation, with a median multiplier in the vicinity of a factor of 5 (i.e. the multiple of consumption to required generate an an equivalent increase in SWB is 5 times more than the median). This provides evidence of significant heterogeneity in the impact of various life outcomes, even after accounting for variation in the reporting function.

Conclusion

The results of this section suggest a number of broad conclusions about the presence, scope, and structure of heterogeneity in the determinants of individual SWB. First, we find that a basic set of demographic variables are capable of explaining some differences in what brings individuals' satisfaction. While most of these variables predict variation in only one (urban residence and gender) or two coefficients (education), there are significant age-related patterns in the coefficients on reference income, unemployment, and hours worked (all of which have their greatest negative effect around the age of 50) and a suggestive relationship between age and the importance of relationships (which appear to be most important around one's mid-thirties).

Second, we show that there is nonetheless significant residual variation in the importance of different outcomes to different individuals, that cannot be accounted for by such basic demographic variables. The main axis along which these values appear to cluster broadly reflects notions of materialism vs. non-materialism, with consumption and reference income being most important at one end, and health, employment, relationships, children and social time being more important at the other. Consistent with previous research, we find that those at the more materialistic end of the spectrum tend to have lower (unobservable) SWB levels. Conversely, those who are SWB-poor tend to gain greater satisfaction from consumption, health, and employment, and less satisfaction from relationships, children and social time than those who are SWB-rich.

Finally, the both the scale and the structure of the heterogeneity we identify suggest strongly that such individual differences cannot be due merely to differences in the way individuals report SWB. Both the pattern of variations between coefficients (i.e. their correlations and relative effect sizes) and direct estimates of the distribution of consumption equivalents serve to rule out such a possibility. This conclusion is further reinforced by the presence of a substantial number of individuals who appear to respond to outcomes such as consumption and reference income, not only to a different degree, but in an entirely different *direction* to the average. However, as we have already alluded to, there is a legitimate concern that some of these results — especially the last — could be an artefact of the normality assumption. Section 5.2.3 attempts to allay such concerns.

5.2.3 Allowing for non-normality in the random effects: The Penalized Gaussian Mixture Model

While the results to this point suggest both that there is significant *real* heterogeneity in individuals' responses to objective life outcomes, and unearth some interesting patterns in such heterogeneity, they do suffer from a potential weakness, particularly relative to the latent class approach taken by Clark et al. (2005). While it has been shown that estimates of both the main model parameters (α_0, β_0, θ, Γ) and the variance components (σ^2, Ω) are consistent under non-normality (Verbeke and Lesaffre, 1996a, 1997), inference about other aspects of the random effects distribution (e.g. quantiles of the distribution, or higher-order moments characterising its shape) can be misleading in such cases.

There are in fact good reasons to think that the normality assumption may be violated in practice. As a simple example, consider that the symmetry of the normal distribution implies that there are just as many individuals who value consumption highly (specifically, for whom $\beta_{it,cons} > 2\beta_{0,cons}$) as there are for whom it has a negative effect ($\beta_{0,cons} < 0$). That is, our inference about the number of individuals *harmed* by consumption will be partly determined by the number of individuals who greatly value it (and vice versa). The assumption may turn out to be reasonable, but it is not clear that we have a particularly strong reason to believe it a priori.

This obviously has particular implications for the argument made in the previous section that a significant proportion of individuals have coefficients of the opposite sign to the average. This is interesting not only in its own right, but also because it can potentially provide the strongest evidence we have against idea that reporting heterogeneity could be responsible for our results. Unlike the other two arguments we advance, it does not depend on assuming that the reporting function is approximately affine.[35]

[35] Violation of the normality assumption could also affect the construction of the CMEs

As Verbeke and Lesaffre (1996b) show, it is difficult to test the normality assumption without explicitly fitting a more flexible model, and variety of such models have been proposed to deal with this issue. At one extreme (as we have already alluded to) we have the latent class approach of Clark et al. (2005) (which, as we let the number of latent classes increase, becomes equivalent to non-parametric maximum likelihood (NPML) estimation (Laird, 1978)). However, this approach can be computationally difficult, especially when trying to allow for variation in multiple coefficients simultaneously. Moreover, a number of authors take issue with the latent class assumption that the random coefficients are discrete (e.g. Magder and Zeger, 1996), and have therefore suggested a variety of semi-parametric approaches that both allow for continuous distributions of the coefficients, and ease the computational burden, while nonetheless avoiding the assumption of normality.[36]

pursued in the previous section. In general, the variance of (a function of) the ratio of two random variables depends on more than just their first- and second-order moments (see e.g. Goodman, 1960; Bohrnstedt and Goldberger, 1969). The normality assumption may therefore give misleading estimates of the scale of variation in individual money equivalents, even if the main variance components of the coefficients themselves are estimated consistently.

Similar concerns about possible violations of normality may also arise in any form of Bayesian updating that makes use of the estimated random effects distribution as a prior — whether this means an analyst using empirical Bayes methods to estimate individual specific coefficients for the sample population, or, as we discuss in Chapter 6, individuals attempting to integrate private information about their personal values with regression-based estimates of β coefficients.

[36]These include smoothed NPML (Magder and Zeger, 1996), predictive recursive estimation (Tao et al., 1999), the heterogeneity LMM (Verbeke and Lesaffre, 1996b), smoothing by roughening (Shen and Louis, 1999), and semi-nonparametric (SNP) estimation (Zhang and Davidian, 2001).

The approach we adopt here, which is both computationally tractable in our context, and which has been shown to perform well relative to the alternatives (see Ghidey, Lesaffre, and Verbeke, forthcoming, for a comparison), is the "penalized Gaussian mixture" (PGM) model developed by Ghidey, Lesaffre, and Eilers (2004). To our knowledge, this is the first application of this model in the economics literature, though it has been used in a number of applications in biometrics (e.g. cholesterol increase Ghidey, Lesaffre, and Eilers, 2004; the onset of AIDS-related illness Komárek, Lesaffre, and Hilton, 2005; the onset of dental problems Komárek and Lesaffre, 2006; and the emergence of permanent teeth Bogaerts and Lesaffre, 2008).[37]

The PGM approach assumes that the true random effects density can be approximated by a mixture of (multivariate) Gaussians with fixed means. In addition to the parameters of the standard Gaussian LMM, it requires us to estimate mixing probabilities for these basis densities. To avoid over-fitting, it applies a (data-determined) smoothing penalty to the log-likelihood arising from the mixture densities. The technical details of the PGM approach and our implementation of it are set out in Appendix A.7.

[37] It has recently been extended to the generalized linear mixed model (GLMM) context by Komárek and Lesaffre (2008), who also present a slightly more general exposition of the underlying Gaussian mixture model. (Ghidey, Lesaffre, and Eilers, 2004, focus on the special case with a single random coefficient and a random intercept.)

PGM Results

The PGM approach, while more tractable than alternative methods of weakening the normality assumption, nonetheless has difficulty fitting many random effects simultaneously. Consequently, we focus here on the random coefficients on three variables: consumption, reference income, and health (as well as allowing for a random intercept). It is the consumption and health coefficients that seem most likely to violate normality by having a consistent positive sign.

As detailed in Appendix A.7, our algorithm for fitting the PGM model begins with a standard LMM fit (indeed, this is one of the main reasons for its tractability). Because the main parameters of this model are consistent, even when normality is violated, the algorithm essentially assumes that they are correct, and focuses on attempting to estimate the *shape* of the random coefficient and intercept distributions under that assumption.[38] The upshot of this fitting procedure is that the main coefficients do not differ appreciably from those fit in Section 5.2.2.[39] To conserve space and attention, therefore, we include these in Appendix A.9, and focus instead on the shape of the

[38]The procedure does in fact update these main coefficients after fitting the mixing parameters, but they change very little.

[39]The main differences are that: there are now significant age-patterns in all the included coefficients; men are significantly more negatively-affected by reference income than women; and the negative correlation between the random coefficients on consumption and reference income has become significant. In general these tend to strengthen, rather than detract from the conclusions of the previous section.

resulting random coefficient densities, which are displayed in Figure 5.3.[40]

Figure 5.3: PGM estimates of unobserved heterogeneity in the determinants of SWB

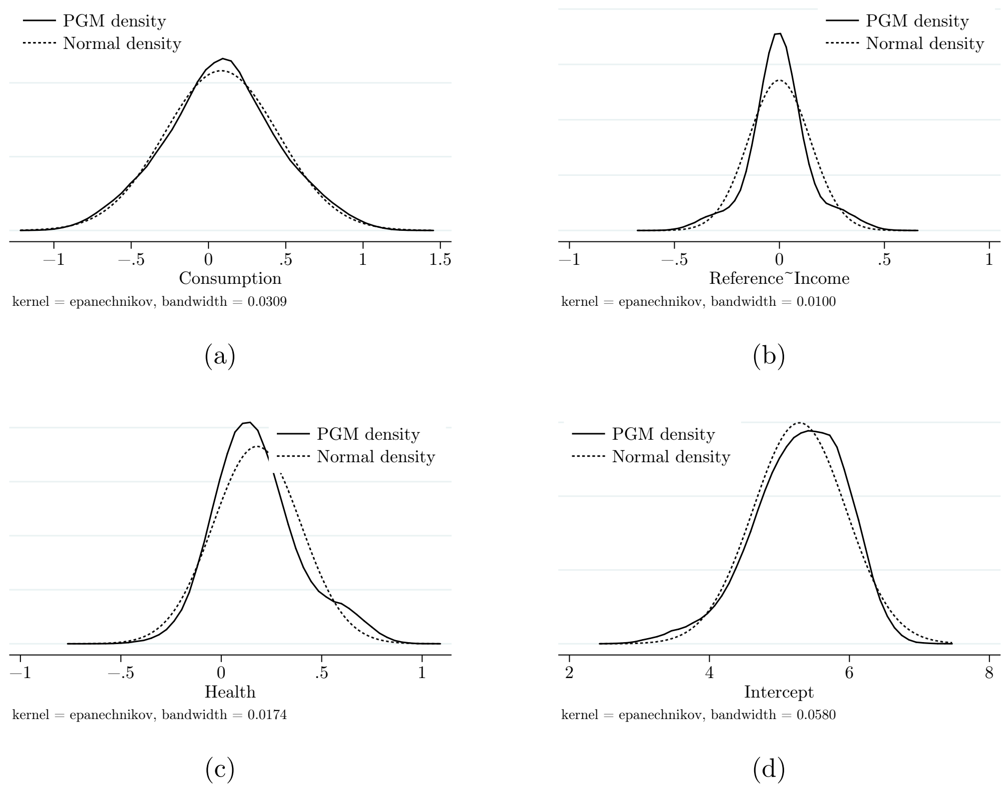

Kernel density estimates of the distribution of unobserved heterogeneity in individual coefficients on (a) log household consumption, (b) reference income, (c) health, and (d) of the distribution of the individual intercepts. β_i coefficients for each variable are on the horizontal axis, while the vertical axis represents the density. Each panel displays PGM estimates against the corresponding normal densities. While the β_{cons} is well-approximated by the normal, the others have noticeable departures from normality. The β_{refinc} density has higher kurtosis and heavier tails, while the β_{health} and η densities are skewed to the right (0.5) and left (-0.5) respectively.

Each panel compares the estimated PGM distribution of the random effects with the corresponding normal distribution. As we can see, the PGM

[40]We have also estimated the distributions of the CMEs implied by these results. The (unreported) results still support the presence of significant heterogeneity.

distribution of consumption coefficients is very well approximated by the corresponding normal distribution (skewness 0.01 and kurtosis 2.9). However, the densities of the the other coefficients depart more noticeably from normality. The β_{refinc} distribution is still roughly symmetric (skewness=0.07), but has much heavier tails (kurtosis=4.2). The β_{health} and η densities are skewed to the right (skewness=0.5) and left (skewness=-0.5) respectively, and both have very slightly heavier tails than the normal (kurtosis 3.3 and 3.2 respectively).

As to the matter of substantive interest, the first two random coefficient distributions demonstrate clearly that there are indeed individuals for whom both consumption and reference income have oppositely signed effects. Indeed, for both coefficients the proportion of individuals with such coefficients are essentially unchanged from the LMM results: an estimated 40% of individuals continue to have negative consumption coefficients, while an estimated 49% of individuals continue to have positive responses to changes in others' income. In contrast, while we still estimate a substantial portion of the population to have negative coefficients on health, this percentage has now fallen to 20% (from 24% previously).

It is possible that the PGM estimations could still overstate the proportion of oppositely-signed effects — by construction, they will always place some probability on such effects, even if it is very small. Nonetheless, the proportions for β_{cons} and β_{refinc} in particular seem difficult to explain purely as a statistical artefact, and the estimations therefore provide much more solid evidence that such effects are real. As we have already suggested, there are a number of possible explanations of for negative effects, ranging from omitted variables (such as debt) to unmodelled changes in aspirations. Determining which of these is responsible for our findings seems a potentially important topic for future research. Regardless of the explanation, however, these findings lend further support to the contention that reporting heterogeneity alone is unlikely to be responsible for the variation in the determinants of SWB that we have identified.

5.3 Conclusion

In this chapter, we have investigated the structure of individual differences in what brings them satisfaction. In line with the limited existing literature, we find substantial unobserved heterogeneity in the way individuals translate objective outcomes into reported SWB. However, in contrast to this previous work, we adopt two major innovations that allow us to gain greater insight into the structure of individual differences, and crucially, allow us to demonstrate that the differences we observe are not merely due to differences in the way individuals translate "true" SWB into a verbal life-satisfaction scale.

The first is to allow simultaneously for heterogeneity in the effects of a wider range of variables, including consumption, reference income, hours worked, health, relationships, and children (most previous studies have focused on a single outcome such as income: Clark et al. 2005; or marriage Lucas et al. 2003). Doing so, we discover that there are interesting correlations between the importance of different outcomes. In particular, there is evidence to suggest that "materialism" is the major axis of variation in individual values: individuals who value consumption tend to pay more attention to others' income, be hurt more by unemployment, and be less affected by long working hours. At the same time, they also seem to value less such non-material outcomes as health, relationships, and children. We also find that this cluster of values tends to be associated with lower overall life satisfaction (conditional on other outcomes).

Turning to focus on the outcomes that are most important to the SWB-poor (rather than just the conditionally SWB-poor), we find a slightly different pattern: such individuals do respond more to changes in consumption, but they are also especially affected by health and employment status. In contrast, they appear to gain significantly less from "social" outcomes such as relationships, social time, and children. Finally, allowing for heterogeneity in multiple outcomes also allows us to calculate money equivalents for the (non-consumption) variables. Because such money equivalents do not depend on the scale of the reporting function, showing that there is significant variability in these demonstrates that a sizeable amount of individual-level differences are real and not merely due to reporting differences.

The second innovation is to relax the normality assumption underlying the Linear Mixed Model approach to identifying individual differences (used e.g. by Lucas et al., 2003). Using a novel (to the economics literature) yet relatively tractable, semi-parametric method — the penalised Gaussian mixture model (PGM) — we are able to demonstrate, firstly, that the normality assumption seems relatively unproblematic in this context, and secondly, that there is sufficient variation in the determinants of SWB that different individuals will often not only react to changes in these consumption and reference income to different *degrees*, but in fact in entirely different *directions*.

A number of possible explanations exist for such differences. We discussed in Part II the possibility of reference income could have contrasting information and comparison effects. For consumption, it is possible that the measured negative effect is due to unobservables such as debt, or alternatively, that increases in consumption can sometimes raise expectations faster than they can be met. Determining which (if any) of these is responsible for the negative effect may be an important avenue for future research. Whatever the reason, the fact that different individuals appears to react in different directions further strengthens the case that the differences we observe in individual values cannot be attributed merely to variation in reporting.

Although this chapter has focused primarily on inter-individual differences in values, one interesting pattern that emerged was the apparent changes in the importance of different outcomes over the life-cycle, with a number of coefficients (reference income, unemployment, working hours, and relationships) having distinct age-profiles. This serves to highlight the potential importance of allowing for intra-individual variation in the importance of different outcomes. Understanding the way individual values may change over time remains an important avenue for research.

Chapter 6

Do We Know What Makes Us Happy?

The previous chapter demonstrated that there exists significant heterogeneity in individuals' SWB responses to various life outcomes. Moreover, while some of this could be explained by observable demographic factors such as age, gender, education, and residence, much of it could not. As we explain in greater detail in Section 6.1, the presence of significant *un*explained heterogeneity raises questions about the ability of SWB research to usefully contribute to both policy and individual decision-making; but its value in this sense is *relative*. If individuals' private judgments about what will bring them satisfaction are sufficiently poor, then (a) SWB research may be able to improve individual decision-making, and (b) policy rules based on population averages may perform better relative to more decentralised decision-making.

While economists have traditionally been content to assume that individuals are the best judges of what will bring them satisfaction, recent work in both psychology and behavioural economics has called this assumption into question — finding that individuals tend to overestimate the impact of events on SWB in a wide variety of domains (see e.g. Gilbert and Wilson, 2000; Wilson and Gilbert, 2003, 2005; Gilbert and Wilson, 2007, for overviews). In Section 7.3.2, we use BHPS data about individuals' stated values to conduct a number of simple — but to our knowledge novel — tests of the quality of individuals' self-knowledge.

We seek to address two distinct questions in this regard. First, in Section 6.2.2, we investigate whether individuals accurately perceive the relative value of different life outcomes (such as the importance of health relative to the importance of being in a relationship) *on average*. While existing research has found evidence of a general tendency to overestimate the SWB impact of various outcomes, it remains an interesting question whether such biases are stronger in some domains than others — especially as it is such relative biases that will introduce the strongest distortions in individuals' perceptions of the trade-offs they face between different outcomes. We find that, while there is some evidence that individuals overestimate the relative importance of having children and money, they appear to have roughly accurate perceptions of the relative value of health, friends, relationships, and jobs.

Nonetheless, even when individuals' perceptions are approximately unbiased, they may still be very noisy predictors of actual SWB responses — an issue that has received less attention in the literature. We therefore go on, in Sections 6.2.3 and 6.2.4, to investigate the *precision* of individuals' private information. Section 6.2.3 focuses on whether individuals accurately perceive whether particular outcomes are especially important *to them* (e.g. whether those who say that money is particularly important actually get more satisfaction from money than those who do not), while Section 6.2.4 investigates the general ability of individuals to correctly prioritise some outcomes over others.

Our results suggest that individuals' ability to judge what will bring them satisfaction varies both between different domains and between different individuals. Section 6.2.3 suggests that we are generally poor judges of whether particular life outcomes are especially important to us. Those who say that money, health, or children are important do not actually get greater satisfaction from these outcomes. However, those who say that relationships or a fulfilling job are especially important do seem to gain greater satisfaction from them, and there is weak evidence that the same may be true of friendships. Finally, Section 6.2.4 provides further support for the conclusion that we generally do a poor job of prioritising different outcomes, but also suggests that there is substantial variation in the quality of different individuals' information.

While there are reasons to be cautious in our interpretation of these results, they offer a fairly pessimistic view of the individuals' ability to predict what will bring them satisfaction. It is also noteworthy that the areas in which individuals appear to have the best information (friendships and employment) are those where there is the greatest variation in population responses; on the other hand, the areas where individual perceptions seem most likely to be biased (having children and money) are those that exhibit the least population variation. It is may therefore be in these latter areas that SWB research has the most to offer.

6.1 Incorporating Population Information into Individual Decisions

One possible use of SWB research is to assist individuals in making various life choices: it may provide information about such matters as whether getting married (see Stutzer and Frey, 2006) or having children (see e.g. Powdthavee, 2009) will make us happy, how we should make the trade-off between higher pay and longer working hours/less time with friends and family (Powdthavee, 2008), or how we might most enjoyably spend our time (see e.g. Kahneman and Krueger, 2006). More generally, money equivalents derived from SWB regressions could assist individuals in making trade-offs in a wide range of areas: e.g. money equivalents for health or unemployment could inform trade-offs between higher paying jobs that have a higher risk of injury or job loss, and lower paying jobs with correspondingly lower risks (or indeed, trade-offs between equally-paying jobs with different types of risk).

Despite its potential in this regard, applications of SWB research to individual decision-making are hampered by the fact that it is often unclear how relevant are estimates of population average effects to any particular individual. If I think that having a child will make me more satisfied with my life (say, by as much as a 20% increase in consumption) and I then learn that the average effect of having a child is to reduce SWB (by as much as a 20% drop in income), how am I to reconcile these conflicting pieces of information? Should I assume that I am a perfectly-informed neoclassical agent and ignore the econometric estimate entirely? Should I, as Gilbert et al. (2009) and Gilbert (2006, esp. Ch. 11) advise, recognize that my guesses about what will make me happy are so poor that I would be best to ignore them, and instead place my trust in the reactions of a hypothetical statistical construct?[1] Or should I "split the difference" and assume that having a child will leave me roughly as satisfied as I am now?

As we explain below, the answer to this question depends heavily on the amount of variation in the population: in an extreme case where there were literally *no* variation in population responses (i.e. children make everyone equally (un)happy) I should assume that I am wrong; but if there is large variation, then this leaves more scope for me to be right. In this light, the results presented in Chapter 5 — which suggested that there *is* significant individual-level variation in the determinants of SWB — do not bode well for the practical value of SWB research to individual decision-makers. On the other hand, much recent work in psychology and behavioural economics has called into question the quality of individuals' judgments about what will bring them satisfaction. If our powers of introspection are sufficiently poor, then SWB research could be the better of two imperfect alternatives — and at the very least may still be able to contribute *something* to our decision-making.

Rather than placing complete faith in either my private information (π_i) about the SWB impact of having a child, or the information contained in econometric estimates (θ), Bayes' rule suggests a middle ground that takes

[1] Of course, I could still decide to have a child, even if I knew it would probably make me unhappy — my own satisfaction need not be my only, or even my primary motivation. Nonetheless, I am presumably better off making this decision fully informed of the potential costs.

account of both. If we denote the effect of having a child (or indeed, any other life outcome) on my SWB as β_i, then my posterior distribution should satisfy (with the traditional abuse of notation):

$$\begin{aligned} f(\beta_i|\theta, \pi_i) &\propto f(\beta_i|\pi_i) f(\theta|\beta_i, \pi) \\ &= f(\beta_i|\pi_i) f(\beta_i|\theta) c[F(\theta|\beta_i), F(\pi_i|\beta_i)] \end{aligned} \tag{6.1}$$

The factorization of the likelihood in the second line follows Jouini and Clemen (1996), and shows that the posterior distribution can be thought of as depending on three distinct factors: our prior, $f(\beta_i|\pi_i)$; the population distribution of the βs, $f(\beta_i|\theta)$; and a copula $C[F(\theta|\beta_i), F(\pi_i|\beta_i)]$ describing the dependence between the two (where $c[F_1, F_2] = \frac{\partial^2 C[F_1, F_2]}{\partial F_1 \partial F_2}$, and $F(\cdot)$ is the cumulative distribution function corresponding to the density $f(\cdot)$).

Depending on the nature of the decision we face, attempting to employ Bayes' rule in this fashion may be quite complicated, and it is difficult to make fully general statements about the impact of the population information on our posterior.[2] Nonetheless, intuition suggests that the usefulness of

[2] Indeed, many of the results in the literature regarding the expected value of information prove the impossibility of making general statements that will hold for all possible priors, likelihoods, and decision contexts (see e.g. Hilton, 1981). For example, it is not true *generally* that greater uncertainty in the prior will increase the value of additional information in all decision contexts (Gould, 1974), though it does appear that this will *usually* be the case (see e.g. Marschak and Radner, 1972; Itami, 1977; Hilton, 1979, 1981, Table 3). On the other hand, it *is* generally true that the expected value of information is non-decreasing in its accuracy (at least for a certain definition of accuracy: see e.g. Blackwell, 1953; Marschak and Radner, 1972; Hilton, 1981).

econometric estimates will usually (if not universally) increase: (a) the less accurate is individuals' prior private information; (b) the less variation there is in the population under consideration; and (c) the less individuals have already incorporated similar information in their priors.

A Simple Example

As an example, consider a decision made under the following simplifying assumptions.

(1) The optimal decision depends only on the posterior mean, and not any other moments of the posterior distribution. This will be the case for a binary decision, such as whether or not to have a child;[3] as is well-known, it will also be the case for any quadratic loss function, and for "proper" loss functions more generally (see e.g Lindley, 1985, for a further characterisation of proper loss functions).

(2) The prior, $f(\beta_i|\pi_i)$, and the population distribution, $f(\theta|\beta_i)$ are both normal.

[3] Designate the choice of whether to have a child by a binary variable k. If the utility of having a child ($k=1$) is β^*, and the utility of not having a child ($k=0$) is neutral, then $u(\beta^*,k)=k\beta^*$, and maximising expected utility corresponds to choosing $k\in\{0,1\}$ to maximise:

$$E(u(\beta^*,k(f_\beta)))=\int_{-\infty}^{\infty} f_\beta(\beta^*)\beta^* k d\beta^* \tag{6.2}$$

The optimal decision rule is then:

$$k(f_\beta)=\mathbb{1}(E_\beta(\beta^*)>0) \tag{6.3}$$

(3) The copula is of the form:

$$c[F(\theta|\beta_i), F(\pi_i|\beta_i)] = \alpha + \frac{(1-\alpha)}{f(\beta_i|\theta)} \tag{6.4}$$

where $\alpha \in [0, 1]$ is a parameter indexing the "degree of independence" between the prior and the posterior. When $\alpha = 1$ the two are completely independent, and the copula drops out of (6.1) entirely; when $\alpha = 0$, the prior already incorporates all of the information in the population distribution, and the the posterior is not updated at all.[4] Intermediate cases result in a posterior that is a mixing distribution of these two extremes.

Under assumptions (2) and (3), the posterior mean ($\hat{\beta}_{\theta,\pi}$) will simply be a weighted average of the prior and population means ($\hat{\beta}_\pi$ and $\hat{\beta}_\theta$):

$$\hat{\beta}_{\theta,\pi} = \frac{\sigma_\theta^2 + (1-\alpha)\sigma_\pi^2}{\sigma_\pi^2 + \sigma_\theta^2}\hat{\beta}_\pi + \frac{\alpha\sigma_\pi^2}{\sigma_\pi^2 + \sigma_\theta^2}\hat{\beta}_\theta \tag{6.5}$$

[4] It seems rather unlikely that a given individual would know only a part of the population distribution *and nothing else*, as would be required for the information in θ to completely encompass the private information in π_i. In general, it seems reasonable to assume that it is possible that $\theta \subseteq \pi_i$, but that $\pi_i \not\subset \theta$. (Cf. Jouini and Clemen (1996), who consider the specification of some easily parameterised copulas in the context of combining expert opinions, using the alternative assumption that their arguments are exchangeable — which here would amount to treating π_i and θ symmetrically.)

From this, we can define the relative influence measure of β_θ compared with β_π:

$$\iota(\theta, \pi) \equiv \frac{\alpha\sigma_\pi^2}{(1-\alpha)\sigma_\pi^2 + \sigma_\theta^2} \tag{6.6}$$

The posterior mean will favour $\hat{\beta}_\theta$ if $\iota(\theta, \pi) > 1$ and favour $\hat{\beta}_\pi$ if $\iota(\theta, \pi) < 1$, and reduce to the "split the difference" approach if $\iota(\theta, \pi) = 1$.

The analysis of this simple case supports our earlier intuitions about the determinants of the value of econometric estimates of β. Taking $\iota(\theta, \pi)$ as a measure of usefulness, usefulness is increasing in prior variance σ_π^2, decreasing in the population variance σ_θ^2, and increasing in the degree of independence α (i.e. exactly the three factors mentioned earlier).

This underscores the potential importance of accounting for individual heterogeneity: if we have estimated only a population average parameter for β, we will have little clear evidence on which to base an estimate of the full distribution.[5] If, on the other hand, we have estimated the population distribution directly (e.g. using the sort of random coefficient model adopted in the previous section, or the sort of latent class approach adopted by Clark

[5] The standard errors on the β coefficient reflect the accuracy of the coefficient as a measure of the average population effect, rather than the degree of variation in the population *per se*. While the two are usually related, they are not identical: if we had data for the entire population, rather than just a sample, our standard errors would be zero, but that would obviously not imply that there was no individual-level variation the population. In the univariate case, the relationship between the standard deviation of the mean, s, and the standard deviation of the values themselves, σ, is approximately given by $s = \sigma\sqrt{\frac{N-n}{n(N-1)}}$. In a more general multiple regression setting, there is no simple relationship between the two.

et al. 2005) then specifying $f(\beta_i|\theta)$ is straightforward.

In this light, the results of Chapter 5 — which showed that σ_θ^2 is often large — are particularly interesting. They lend support to the view that econometric estimates may have little to contribute to individual decision-making (even if they might be relevant in policy formation). Nonetheless, how much such estimates can contribute also depends heavily on the quality of individuals' private information: if, as some have argued, we are really quite poor at predicting what will bring us satisfaction (i.e. σ_π^2 is also large), then they may be of some use after all.

One way to assess the size of σ_π^2 is simply to elicit a subjective probability distribution from a given individual (e.g. by obtaining a small number of its fractiles) and to treat this as representative of the uncertainty in their private information. However, as we detail in the next section, it is possible that such subjective distributions could be systematically biased, overconfident, or both, and we therefore set out to provide a number of novel tests of the accuracy of individuals' private information.

6.2 Do we know what makes us happy?

Gilbert (2006) presents an overview of the evidence that we frequently do a poor job of predicting what will make us happy, and suggests two broad reasons why such "miswanting" typically occurs.[6] The first is that our imagination often misrepresents the specific details of future events, for example, by overemphasising their "essential" features and neglecting "incidental" ones (see e.g. Wilson et al., 2000; Dunn, Wilson, and Gilbert, 2003; Liberman and Trope, 2008). The second is that we tend to project our current feelings and attitudes onto our future selves — and in particular tend not to realise that we will often adapt (at least to some extent) to whatever comes our way.[7] This last factor results in an "impact bias" whereby we overestimate the effect of many events on our SWB, and there is some evidence that such impact bias is stronger for negative than positive events.[8] Although we might

[6]Gilbert himself describes these variously as four separate reasons (Gilbert and Wilson, 2007) three separate reasons (Gilbert, 2006) and as merely a single reason (Gilbert et al., 2009). We follow Wilson and Gilbert (2005) in grouping the major reasons into two groups. For other overviews of this evidence see e.g. Gilbert and Wilson (2000), Wilson and Gilbert (2003), Gilbert and Wilson (2007). For further discussion of the potential divergence between decision utility and experienced utility, see e.g. Kahneman (1994), Kahneman, Wakker, and Sarin (1997).

[7]Clark, Frijters, and Shields (2008) provide a recent overview of some of the evidence that individuals adapt to at least some experiences. An earlier review is provided by Frederick and Loewenstein (1999). For specific examples of adaptation see e.g. Brickman, Coates, and Janoff-Bulman (1978) (finding that lottery winners and accident victims tend to adapt at least partially to their situation), Lucas et al. (2003) (finding adaptation to marriage, but not to widowhood) and Lucas (2005) (finding partial adaptation to divorce), Wu (2001) (finding partial adaptation to heart disease) and Oswald and Powdthavee (2008) (finding partial adaptation to disability), and Di Tella, Haisken-De New, and MacCulloch (2007) (who find evidence of adaptation to income and status).

For evidence that we tend not to anticipate changes in tastes generally see e.g. Herrnstein and Prelec (1991) Kahneman and Snell (1992), Loewenstein and Adler (1995), Loewenstein and Angner (2003), and Loewenstein, O'Donoghue, and Rabin (2003). For evidence specifically on failure to anticipate adaptation, see e.g. Loewenstein and Schkade (1999), Loewenstein and Frederick (1997), Gilbert et al. (1998), Sieff, Dawes, and Loewenstein (1999), and Ubel, Loewenstein, and Jepson (2005).

[8]Kermer et al. (2006) argue that explains the phenomenon of loss-aversion in decision contexts as an "affective forecasting error".

expect to learn to overcome such biases, Gilbert also argues that failures of memory (such as the peak-end effect, whereby we judge past experiences by their most memorable moments at the peak and the end, rather than on the basis of the full experience (see e.g. Kahneman, 1999)) often conspire to prevent us from learning from our mistakes.

In light of such tendencies towards "miswanting", Gilbert et al. (2009) suggest that *others'* affective responses might provide more accurate predictions than do our own imaginings: a strategy that they term *surrogation*. While recognising that others' responses may differ from our own, they nonetheless put forward two reasons why individual-level variation may be small enough to make surrogation worthwhile. The first is based on fundamental physiological similarity in our evolved affective responses, which suggests that for some events at least, such responses are likely to be relatively uniform. The second is based on the idea that, because we tend to associate with people with similar interests and values, their experiences may be especially informative to us. They then present evidence that surrogation — in this case, knowing the response of a single other person — does indeed outperform imagination in two experiments conducted on US college students: one involving a speed date, the other a peer evaluation.[9]

[9]They also suggest that the individuals concerned appear not to believe this result: given the choice between getting more specific details about the event they were about to experience, and being told how others reacted to the same event, they preferred the former, even though it was ultimately less useful to them.

While such results are encouraging for individuals' prospects of learning from SWB research, the extent to which they will generalize to econometric surrogation, or to domains other than speed-dating and peer-feedback is unclear. Surrogating with population average SWB responses $\hat{\beta}_{k0}$, rather than the responses of a single individual β_j seems likely to improve its value. But other differences in the relative accuracy of private information and surrogation could potentially outweigh this effect.[10]

As we noted in the previous section, greater variability of individual responses to different types of outcomes will tend to reduce the relative value of surrogation. Table 6.1 displays the standard deviations of the random effects derived from Table 5.2 (standardised so as to make them roughly comparable, and ranked from least to most important). Assuming that the quality of private information about each were roughly equal, these suggest that surrogation would be least informative about the SWB impact of unemployment, and most informative about the SWB impact of having children.

Individuals may also have better private information about their responses to some outcomes than others. One example of this has already been noted above, in the suggestion that individuals overestimate the im-

[10]It is also possible that violations of the assumptions embodied in our regression approach (e.g. the presence of non-linearity or other misspecification), could reduce the accuracy of surrogation compared with the more direct experimental approach of simply asking how happy a given individual is with a particular outcome. Unfortunately, it is difficult for us to assess the extent of such violations here.

Table 6.1: Standardised Standard Deviations of Random Coefficients from Table 5.2

	σ_β
Children	0.208
Reference Income	0.378
Consumption	0.402
Health	0.468
Hours Worked	0.484
Relationship	0.638
Unemployment	0.736

Standard Deviations of the random coefficients for continuous variables are multiplied by twice the standard deviation of the variables, so as to put them on roughly the same scale as the standard deviations for the coefficients on the dummy variables (see Gelman, 2008). Higher values reduce the value of surrogation relative to private information.

pact of negative events more than positive ones. Other possible sources of variation include the fact that we may have more opportunity for learning and feedback in some domains than others.[11] Perhaps most importantly, there may be differences in the degree to which individuals adapt to different outcomes (see e.g. Di Tella, Haisken-De New, and MacCulloch, 2007, who find evidence of differential adaptation to income and status). If, as Gilbert (2006) argues, much of our failure as affective forecasters is due to not anticipating our eventual adaptation to future events, then our predictions of what will bring us satisfaction are likely to be worse in domains that exhibit greater adaptation.

[11] Although Gilbert (2006) (esp. Ch. 10) suggests that our ability to learn what makes us happy may be limited, it is also possible that the factors that inhibit learning could vary between domains.

Differences in predictive ability between different domains are likely to be particularly important for individual decision-making. Even if individuals do consistently overestimate their affective responses, if this impact bias is fairly uniform, then it may not greatly distort their perceptions of the trade-offs they face between different outcomes (because overestimating the effect of two things by the same proportion will not affect their marginal rates of substitution). On the other hand, if we tend to overestimate the impact of some events more than others, then these variations will distort our perceptions of these trade-offs, and surrogation may become more useful as a potential corrective.

Another issue that arises here is whether some individuals have better information than others. For example, if individuals adapt to changes to different degrees (as suggested by Di Tella, Haisken-De New, and MacCulloch, 2007; Lucas et al., 2003), then those who adapt less may have more accurate perceptions of what will bring them satisfaction.

The BHPS contains a number of questions that can be considered proxies for individuals' private information; it therefore affords us a unique opportunity to examine the issues of individual heterogeneity and private information in an integrated fashion. We describe these private information proxies in Section 6.2.1, before moving on, in Sections 6.2.2–6.2.4, to examine the extent to which they reliably predict individuals' responses to changes in consumption, income, health, employment, being in a relationship, social contact, and having children (which are roughly the variables on which we estimated random coefficients in the previous chapter[12]). We also make a preliminary attempt to determine whether there is significant heterogeneity in the quality of different individuals' private information.

While there are reasons to be cautious in our interpretations, the results lend little support to the view that individuals have especially good private information about the likely impact of changes in their circumstance on SWB. Section 6.2.2 asks whether there is any evidence of bias in individuals' judgments, and finds that while the average importance of most outcomes seems roughly in line with individual perceptions, there is some evidence that individuals systematically overestimate the value of having children and of money to material well-being.

Perhaps more importantly, Section 6.2.3 presents evidence that, with the exceptions of being in a relationship and job satisfaction, individuals' private information about the importance of various outcomes does not appear to strongly predict variation in their SWB responses. Thus, despite the large variation in individual values documented in Section 5.2, estimates of average

[12]The major difference is that we look at income as well consumption. While we estimated a random coefficient on hours worked rather than social hours, the latter can simply be considered the converse of the former. (We initially tried to fit a random coefficient on social hours rather than work, but ran into numerical difficulties achieving convergence, and so switched to hours worked instead.)

coefficients may still be somewhat useful. Section 6.2.4 provides further support for this conclusion, but also suggests that there may be large differences in the quality of different individuals' information.

6.2.1 The Private Information Proxies

We employ two proxies for individuals' private information about their values. In waves 8 and 13, the BHPS questionnaire contained a series of questions asking individuals to indicate the importance of various life outcomes to them, on a scale of 1-10 (p^{IMP}). This question was phrased as follows:

> I'm going to read you a list of things that different people value. For each one I'd like you to tell me on a scale from 1 to 10 how important each one is to you, where "1" equals "Not important at all", and "10" equals "Very important".

(1) Your health

(2) Having a lot of money

(3) Having children

(4) Having a fulfilling job

(5) Being independent

(6) Owning your own home

(7) Having a good marriage or partnership

(8) Having good friends

Items 1-4 and 7-8 correspond most closely to the existing independent variables in our SWB equations, and it is these on which we focus. Because our focus is primarily on inter-individual differences, rather than changes in individuals' values over time, we begin by defining p_{ki}^{IMP} as the average importance of x_k to individual i over the sample: $p_{ki}^{IMP} = \frac{1}{T_i}\sum_{\tau} p_{ki\tau}$. We later relax this assumption, and allow p_{kit}^{IMP} to vary over time.

In waves 7 and 12, respondents were also asked to describe in their own words, what things were important to their "quality of life" (p^{QOL}). The question was worded as follows:

> [This] question asks you to think about things that are important to you. There is a lot of discussion these days about "quality of life", yet that means very different things to different people. Would you take a moment to think about what "quality of life" means to you, and tell me what things you consider are important for your own quality of life?

The responses were subsequently coded into a variety of categories,[13] which we further collapse into groupings reflecting the importance to individuals of material well-being, health, a good job, a good relationship, and social

[13]The initial coding categories are contained in the BHPS documentation.

contact. (Unfortunately, the codings do not allow us to define a p^{QOL} measure for the importance of having children.) In our initial coding, we use the time-invariant binary measure, p_{ki}^{QOL} which defines x_k as being important iff it is mentioned as being important in either of the waves in which the question was asked: $p_{ki}^{QOL} = sgn(p_{ki,7} + p_{ki,12})$.

6.2.2 Is individuals' private information accurate *on average*?

The first question we seek to answer is whether individuals' private perceptions of the importance of various life outcomes align, *on average*, with the importance assigned to those outcomes by SWB regressions. Unfortunately, however, the importance of a given outcome in terms of 10-point importance scale, or a binary judgment that "x_k is (or is not) important" do not have obvious or direct interpretations in terms of the parameters of an SWB equation. It is arguable that the three measures of importance have comparable *origins*: a p_{ki}^{IMP} rating of 1, not being mentioned as important to QOL ($p_{ki}^{QOL} = 0$), and a β_k coefficient of zero, each reflect a judgment that a particular variable has no (or at least little) impact on SWB. But they nonetheless seem to lack a common scale: it is difficult, for example, to say exactly what β_k would be implied by a p^{IMP} rating of 8 out of 10, or a judgment that "x_k

is important" ($p^{QOL} = 1$).

Rather than attempting to directly compare sample average p_k measures with regression coefficients, we instead attempt to test the average accuracy of individuals' private information by comparing the *rankings* of different k variables implied by the p_k measures, with those implied by the relevant β_k regression coefficients. If both measures rank the same variables as being more or less important, then this would provide evidence that individuals have approximately unbiased private information about their values.

In order to conduct the comparison, we first need to translate the β_k coefficients themselves into a comparable scale. While it makes intuitive sense to compare the SWB impact of binary variables like having children, a stable relationship, or a job, directly in terms of their β_k coefficients, it is less obvious how to compare the impact of continuous predictors without specifying how much of a change in health, money, or social time the coefficients reflect. In order to facilitate comparisons amongst the continuous variables, we scale the coefficients on continuous predictors (health, social time, and log consumption) by twice their standard deviation, which, as Gelman (2008) notes, serves to put them on the same scale as the more directly interpretable dummy variables (being in a relationship, having a job, and having children).[14]

[14]The intuition behind this is that the standard deviation, σ_k of a binary variable that

The standard deviations of health, social time, and log consumption are respectively 0.9, 0.53, and 0.54. Scaling the coefficients by twice their standard deviation then means that: β_{health} reflects the SWB impact of a 1.8 unit change (on a 1–5 scale); β_{social} reflects the SWB impact of a 1 hour and 6 minute change in time spent socialising each day; and β_{cons} is reflects the impact of a 1.1 log unit (an approximately 3-fold) change in consumption. While the p^{QOL} question allows for a broad notion of material well-being (including e.g. regular meals and shelter), and thus seems comparable to the standardised β_{cons} coefficient, the p^{IMP} question focuses explicitly on the importance of "having *a lot* of money" [emphasis added]. To take account of this difference, we define an additional coefficient β_{rich} that reflects the SWB impact of moving from the middle to the middle of the upper third of the consumption distribution (i.e. the 3$^{\text{rd}}$ to the 5$^{\text{th}}$ sextile). In our sample, this is a distance of 0.51 log units, and means that β_{rich} will be around half the size of β_{cons} (0.51/1.1).

We run a standard SWB regression to obtain the (standardised) coefficients on money, health, being in a relationship, time spent socialising, being employed, and having children. We also run two additional regressions to

is equally distributed between two categories is 0.5. Dividing the continuous predictors by $2\sigma_k$ gives them the same standard deviation. In contrast to the more common standardisation procedure of dividing variables by σ_k, this approach simultaneously preserves the natural interpretation of the dummy variable coefficients (as not depending on σ_k) and makes them roughly comparable to the continuous predictors.

obtain coefficients on being employed for those in the labour force (i.e. relative to the unemployed and marginally attached), and relative to being out of the labour force.

Do the private information measures agree?

Columns 1 and 2 of Table 6.2 displays the sample average responses for the $\bar{p}_k^{IMP}$ and $\bar{p}_k^{QOL}$ measures relating to health, friends, partner, job, children and money. With a couple of caveats, the rankings implied by the two proxies seem broadly consistent with each other, which gives us some faith in their validity as measures of individuals' private information. For both the $\bar{p}_k^{IMP}$ and the $\bar{p}_k^{QOL}$ proxies, health is the most important outcome, and both measures rank friends as being more important than either being in a relationship, or having a good job.

The major difference between the two rankings is in the relative importance of having a good job versus having a good relationship. Individuals rank being in a relationship as being more important on average (9.1 out of 10) than having a good job (7.8 out of 10, or 8.1 for those in the labour force); yet both in the full sample, and amongst individuals in the labour force, having a good job is more commonly mentioned as contributing to quality of life (by 18% and 25% of respondents respectively, as opposed to 12% of respondents who mention relationships). On the other hand, those who are out

Table 6.2: Importance of Various Life Outcomes

	$\bar{p}_k^{IMP*}$	$\bar{p}_k^{QOL\dagger}$	$\hat{\beta}_k{}^{\ddagger}$
Health	9.6	69	.35***
Money		63	.12***
Friends/Social Time	9.2	26	.15***
Relationship	9.1	12	.22***
Job (Full Sample)	7.8	18	.04
Job (In Labour Force only)	8.1	24	.29***
Job (Out of Labour Force only)	7.3	3	-.08
Children	7.7		.01
Being Rich	6.4		.06***

* Average out of 10. † Percentage of respondents who mention x_k as being important to their quality of life. ‡ * p<0.10, ** p<0.05, *** p<0.01.

of the labour force hardly ever mention having a good job as being important to their quality of life (3% of respondents); but they nonetheless rate it relatively highly (on average 7.3 out of 10) when explicitly asked about its importance.[15] One possibility is that (some) out-of-the labour force individuals interpret the question as being about the importance of having a *fulfilling* job, conditional on having a job at all, but we cannot be sure of this.

Consistent with our expectations of diminishing marginal returns to consumption, there is a stark difference between the average stated importance of being rich, and the number of people who mention aspects of material well-being as important contributors to quality of life: while being rich is

[15]Checks indicate that this is not simply an artefact of averaging across individuals who move into or out of the labour force between waves 8 and 13.

the least important life outcome when assessed according to the p^{IMP} metric (with an average rating of 6.4 out of 10), money is the second-most frequently mentioned contributor to quality of life (mentioned by 63% of respondents).

As formal tests of the correspondence between the p^{IMP} and p^{QOL} measures, we calculate the values of Kendall's τ (Kendall and Gibbons, 1990) for the full sample, the sub-sample in the labour force, and the sub-sample out of the labour force (excluding the responses to the money questions, on the basis that they seem to be asking about the importance of quite different things).[16] These are displayed in the first row of Table 6.3. As we can see, the correlations are all reasonably large (and indeed, are perfect for the out-of-the-labour-force sub-sample). While they are not statistically significant in the full or in-the-labour-force samples, we would not wish to place too much weight on this: the test has little power when we compare the rankings for only 4 variables, and can never reject the null of no correlation unless the two measures are perfectly aligned (and even then, can only do so at the 10% level). Overall, the τ measures seem to confirm that the measures track each other fairly well on average.

[16] Kendall's τ provides a non-parametric measure of the correlation between the two sets of rankings, with values of 1, -1, and 0 representing perfect agreement, perfect inversion, and independence, respectively. In contrast to Spearman's ρ, τ provides unbiased estimates, better small sample confidence interval coverage (see e.g. Kendall and Gibbons, 1990; Arndt, Turvey, and Andreasen, 1999). It also has a more natural interpretation:the probability of the two measures agreeing on the relative ordering of any two given variables is $(1 + \tau)/2$. (Alternatively, the odds ratio of agreement-to-disagreement is given by $(1 + \tau)/(1 - \tau)$.)

Table 6.3: Kendall's τ

Measures	Sample (by labour force status) All	In	Out
$\tau(p_k^{IMP}, p_k^{QOL})$	0.667	0.667	1.000
	(0.308)	(0.308)	(0.089)
$\tau(p_k^{IMP}, \beta_k)$	0.600	0.467	0.600
	(0.133)	(0.259)	(0.133)
$\tau(p_k^{QOL}, \beta_k)$	0.200	0.000	0.400
	(0.806)	(1.000)	(0.462)
$\tau(p_k^{QOL}, \beta_k)$	0.333	0.333	0.667
(for $k \neq money$)	(0.734)	(0.734)	(0.308)

Kendall's τ reflects the correlation between the ranks implied by pairs of measures reflecting the relative importance of different variables. p-values are displayed in parentheses, but are not expecially informative due to the low power of the test in small number of variables being compared (e.g. in the first line, their are only 4 common alternatives, and the hypothesis test can never reject the null (no correlation) unless both rankings are in perfect agreement).

Do the private information measures align with regression coefficients?

The final column of Table 6.2 displays the standardised regression coefficients derived from the SWB regressions described above. Sensitive as they are to the choice of scaling, we should be wary of over-interpreting these findings; but the rankings do seem broadly comparable to those suggested by the private information proxies. Consistent with those results, health appears as the most important factor, while being rich and having children are least important. The regression results tend to support the p^{QOL} measure over the p^{IMP} proxy in assigning little value to having a job for those out of the labour force, but otherwise seem roughly in line with both sets of results in assigning middling value to being in a relationship, having a job, and social time. There are however two major exceptions to this general pattern of agreement.

The first relates to the p^{IMP} measure of the importance of children, which seems to overestimate their SWB impact in both relative and absolute terms. In relative terms, the p^{IMP} measure places children above the importance of being rich, while the regression estimates suggest that becoming rich would actually have a greater SWB impact. It is also striking that, while the coefficient on having children is tiny and statistically insignificant, it is still rated as being reasonably important in absolute terms: not only does it obtain an average p^{IMP} rating of 7.7 out of 10 (well above the "not at all important" rating of 1); over 40% of respondents assign having children the maximum rating of 10 out of 10, making it the modal response.[17]

The second relates to the p^{QOL} measure of the importance of money. While 63% of respondents mention material well-being as important (placing it second), it rates only 4$^{\text{th}}$ or 5$^{\text{th}}$ (depending on the sample) according to the regression coefficients. This provides suggestive evidence that individuals

[17] An alternative explanation for this difference is that our fixed effects regressions may only capture the short-run impact of having children, while the stated importance of children may focus more on the long-term. This suggestion finds little support from Table 5.3, which suggests that those who spend a greater portion of the sample with children are actually less satisfied than those without — although this latter result could also be due to correlation with unobserved personality traits.

may overestimate the importance of money to SWB. On the other hand, the importance of being rich seems to be placed reasonably accurately by the p^{IMP} measure, so we would not wish to read too much into this — particularly as the size of the β_{cons} coefficient may be more sensitive than others to mis-specification of the log functional form.[18]

Turning to the Kendall's τ values in Table 6.3, we see that the correlation between the p_k^{IMP}s and βs is fairly strong in all three samples, though never attaining statistical significance. In contrast, the correlations between p_k^{QOL} and β_k, while usually positive, are somewhat weaker. The main difference (apart from the absence of a $p_{children}^{QOL}$ measure, which reduces the tests' power) is the ranking of money mentioned in the previous paragraph. If we exclude this from the rankings, the τ correlations rise to 0.333 or .667 (depending on the sample) and actually exceed those for p^{IMP} for those out of the labour force.

The results of this analysis are consistent with the idea that individuals private information provides *unbiased* estimates of the the SWB impact of most outcomes — with the caveat that the impact of having children, and possibly also of consumption are likely to be overstated. However, this provides little information about the *precision* of this private information,

[18]Though see Layard, Mayraz, and Nickell (2008) for evidence that the log form is a reasonable approximation.

which is also crucial to assessing the relative value of SWB estimates. In the following subsections, we therefore provide two more direct tests of this precision. The first is a regression-based interaction approach that assumes a level of interpersonal comparability in individuals' responses to the p^{IMP} and p^{QOL} questions, but avoids the need to compare the importance of different outcomes. The second is an alternative approach based on calculating individual-level Kendall's τ values, which resurrects the assumption of inter-outcome comparability, but allows us to abandon the assumption of interpersonal comparability, and also to identify heterogeneity in the quality of individuals' private information.

6.2.3 Does variation in private information predict variation in SWB responses?

In this section, we investigate the ability of private information to predict inter-individual variation around the population average β_{k0}s. We do so simply by including, in standard SWB regressions, a set of interactions between the proxies described in the previous section and the corresponding variables of interest (e.g. an interaction between health and the reported importance of health). A significant (positive) interaction can then be taken as evidence that individuals have at least some ability to predict where they sit relative to others in the population, and suggests that they can be reasonably confident in private estimates of βs that diverge from the population average.

Formally, we estimate a standard SWB equation of the sort used in the previous chapter:

$$u_{it} = x'_{it}\beta_{it} + \eta_i + \varepsilon_{it} \tag{6.7}$$

but with the β_{kit} coefficients modelled as depending on either the time-invariant private information proxy p_{ki} (defined in 6.2.1):

$$\beta_{ki} = \beta_{k0} + \ddot{p}_{ki}\gamma_k \tag{6.8}$$

or a time-varying version, that allows for the possibility that values might change:

$$\beta_{kit} = \beta_{k0} + \ddot{p}_{kit}\gamma_k \tag{6.9}$$

In this latter equation, to avoid dropping rounds in which the relevant questions were not asked, we impute p_{kit} as equal to the nearest available measure.[19] In both cases, we centre the private information proxies around their sample average values, so that β_{k0} will capture the importance of x_k for an average individual, and γ_k will capture the extent to which the proxy predicts variation around this average.

[19] We also experimented with interpolation; the results did not differ appreciably.

The specific interactions we include are as follows: p_{money} with each of log household consumption, log household income, and reference income (constructed as described in Chapter 5; p_{health} with health; $p_{relationship}$ with a dummy indicating that an individual is in any form of relationship (married, cohabiting, or non-cohabiting); p_{job} with dummies for each of employed or unemployed (either looking for work, or discouraged); $p_{friends}$ with hours spent socialising; and $p_{children}$ with a dummy indicating whether the individual has children. In specifications where we allow the p measures to vary over time, we also include their main effects directly (the time-invariant p variables are absorbed within the fixed effects).

The results of regressions using equations (6.8) and (6.9) with the p^{IMP} measure are displayed in Table 6.4, while those for the p^{QOL} measure are shown in Table 6.5. As the results are fairly consistent between the two sets of estimations, we discuss them together. Looking first at the controls, we see the standard result that household size (now adjusted so as to reflect non-child household members — the effect of children being captured in separate children dummies) has a negative relationship with SWB, while, just as in Chapter 5, living in an urban area does too. The coefficients on age suggest the usual u-shaped relationship, but are now statistically insignificant.

The main effects on the variables of interest also exhibit fairly standard patterns: average household income has an insignificant effect on SWB, while

Table 6.4: Do reported p^{IMP} values predict SWB responses?

	$\bar{p}_{ki}^{IMP}$		p_{kit}^{IMP}	
(Log) Household Consumption	0.109***	(0.025)	0.106***	(0.025)
Consumption x p_{money}	-0.009	(0.013)	-0.000	(0.010)
(Log) Household Income	0.002	(0.010)	-0.006	(0.030)
Income x p_{money}	0.003	(0.005)	0.001	(0.004)
Reference Income	-0.009	(0.009)	0.003	(0.026)
Reference Income x p_{money}	-0.002	(0.005)	-0.002	(0.004)
Health	0.196***	(0.009)	0.196***	(0.008)
Health x p_{health}	-0.017*	(0.010)	-0.011	(0.007)
Non-cohabiting	0.160***	(0.031)	0.158***	(0.030)
Married	0.202***	(0.045)	0.190***	(0.045)
Couple	0.262***	(0.038)	0.252***	(0.038)
Widowed	-0.123	(0.080)	-0.136*	(0.080)
Divorced	-0.074	(0.063)	-0.075	(0.063)
Separated	-0.334***	(0.068)	-0.333***	(0.068)
In Relationship x $p_{relationship}$	0.046***		0.040***	(0.011)
Employed	-0.053	(0.035)	-0.057	(0.035)
Unemployed (Seeking Work)	-0.460***	(0.060)	-0.453***	(0.112)
Unemployed (Discouraged)	-0.172***	(0.039)	-0.166*	(0.099)
Employed x p_{job}	0.010	(0.014)	0.010	(0.008)
Unemployed x p_{job}	0.001	(0.015)	-0.001	(0.013)
Hours Social Leisure	0.149***	(0.033)	0.149***	(0.033)
Hours Work	-0.000	(0.025)	0.001	(0.025)
Hours Sleep	1.067***	(0.220)	1.085***	(0.219)
Hours Sleep^2	-0.059***	(0.011)	-0.060***	(0.011)
Social Hours x $p_{friends}$	0.026*	(0.014)	0.015	(0.011)
Has Children	-0.000	(0.024)	0.003	(0.023)
2 Children	-0.015	(0.025)	-0.012	(0.025)
3 Children	-0.055	(0.043)	-0.048	(0.043)
4 Children	-0.219**	(0.102)	-0.212**	(0.102)
Children x $p_{children}$	-0.001	(0.011)	-0.008	(0.009)
p_{money}			0.007	(0.091)
p_{health}			0.067**	(0.030)
$p_{relationship}$			-0.015*	(0.008)
p_{job}			-0.004	(0.006)
$p_{friends}$			-0.017	(0.029)
$p_{children}$			0.008	(0.005)
Age	-0.008	(0.019)	-0.009	(0.019)
Age^2/10	0.001	(0.001)	0.001	(0.001)
Household Size (Adjusted)	-0.066***	(0.011)	-0.066***	(0.011)
Urban	-0.091***	(0.034)	-0.093***	(0.034)
Constant	-1.110	(1.470)	-1.113	(1.467)
Observations	54804		54804	
Groups	11348		11348	
R^2	0.077		0.073	

* p<0.10, ** p<0.05, *** p<0.01. Standard Errors (robust to clustering within pids) in parentheses. Wave dummies included but not reported

Table 6.5: Do reported p^{QOL} values predict SWB responses?

	$\bar{p}_{ki}^{QOL}$		p_{kit}^{QOL}	
(Log) Household Consumption	0.102**	(0.040)	0.104***	(0.030)
Consumption x p_{money}	0.013	(0.046)	0.012	(0.031)
(Log) Household Income	-0.000	(0.016)	0.003	(0.013)
Income x p_{money}	-0.002	(0.019)	-0.007	(0.016)
Reference Income	-0.007	(0.013)	-0.003	(0.011)
Reference Income x p_{money}	-0.003	(0.016)	-0.014	(0.014)
Health	0.214***	(0.016)	0.209***	(0.012)
Health x p_{health}	-0.024	(0.019)	-0.022	(0.015)
Non-cohabiting	0.116***	(0.030)	0.122***	(0.030)
Married	0.159***	(0.046)	0.168***	(0.045)
Couple	0.222***	(0.039)	0.229***	(0.038)
Widowed	-0.153*	(0.080)	-0.162**	(0.080)
Divorced	-0.079	(0.063)	-0.079	(0.063)
Separated	-0.335***	(0.068)	-0.333***	(0.068)
In Relationship x $p_{relationship}$	0.272***	(0.086)	0.277***	(0.100)
Employed	-0.048	(0.037)	-0.060*	(0.036)
Unemployed (Seeking Work)	-0.443***	(0.063)	-0.443***	(0.061)
Unemployed (Discouraged)	-0.163***	(0.040)	-0.170***	(0.039)
Employed x p_{job}	-0.104	(0.071)	-0.031	(0.101)
Unemployed x p_{job}	-0.159	(0.103)	-0.144	(0.145)
Hours Social Leisure	0.147***	(0.034)	0.146***	(0.033)
Hours Work	0.005	(0.025)	0.006	(0.025)
Hours Sleep	1.084***	(0.221)	1.063***	(0.221)
Hours Sleep^2	-0.060***	(0.011)	-0.058***	(0.011)
Social Hours x $p_{friends}$	0.029	(0.036)	0.038	(0.033)
Has Children	-0.001	(0.023)	0.001	(0.023)
2 Children	-0.015	(0.025)	-0.015	(0.025)
3 Children	-0.051	(0.043)	-0.054	(0.043)
4 Children	-0.225**	(0.101)	-0.223**	(0.101)
p_{money}			0.109	(0.291)
p_{health}			0.122*	(0.065)
$p_{relationship}$			-0.203**	(0.101)
p_{job}			0.047	(0.102)
$p_{friends}$			-0.048	(0.091)
Age	-0.005	(0.019)	-0.005	(0.019)
Age^2/10	0.001	(0.001)	0.001	(0.001)
Household Size (Adjusted)	-0.066***	(0.011)	-0.066***	(0.011)
Urban	-0.090***	(0.034)	-0.092***	(0.034)
Constant	-1.349	(1.477)	-1.373	(1.482)
Observations	57357		57667	
Groups	13819		13963	
R^2	0.070		0.080	

* p<0.10, ** p<0.05, *** p<0.01. Standard Errors (robust to clustering within pids) in parentheses. Wave dummies included but not reported

consumption has a sizeable and significant positive effect. The employed are slightly, though not significantly, less satisfied than those out of the labour force, while the unemployed and the marginally attached are significantly less satisfied than either (F-tests were used to confirm the difference between the employed and the marginally attached). Health has a sizeable and significant positive effect. Being in a relationship (either married, cohabiting, or non-cohabiting) has a sizeable positive effect, while being widowed, divorced, or separated has a negative effect relative to never having been married. However, the coefficients on widowhood and divorce are sometimes insignificant. Hours of social leisure time have a strong positive effect, while SWB is n-shaped in hours of sleep with those who sleep around 8.5 hours a night reporting the highest SWB. Meanwhile, hours of work are indistinguishable from hours spent in passive leisure (the omitted time-use category) in terms of SWB. Having children does not have a statistically significant effect on SWB until one has 4 or more, at which point the effect becomes significantly negative.

Turning to the interactions of interest, we see that the only private information proxy that consistently predicts differences in individuals' SWB responses is that for the importance of being in a relationship. Looking first at the p^{IMP} proxy, a one point increase in stated importance predicts around a 0.04 point stronger response to being in a relationship — somewhere in the vicinity of one quarter (0.04/0.16) to one sixth (0.046/0.26) of the SWB impact of being in a relationship, depending on what type of relationship we consider. For the p^{QOL} proxy, the SWB of those who mention a good relationship as being a major contributor to their quality of life in some cases responds *over three times more strongly* to being in a relationship than those who do not (e.g. the implied effect of a non-cohabiting relationship is $0.116 + 0.272 = 0.388$ rather than .116 in column 1 of Table 6.5).

In contrast the SWB impact of other outcomes (consumption, income, references income, health, employment, hours of social leisure, and children) in general bears no significant relationship to the corresponding private information proxies. There is a barely significant (at the 10% level) interaction for health in column 1, but it is wrongly signed — implying that those who say health is more important actually benefit *less* from it. One possible explanation for this oddity is that it is due to the subjective nature of the self-reported health variable: if those who say they value health more also tend to report larger changes in subjective health status then the same "objective" change in health would be measured as having a smaller effect.

These initial results suggest that, while individuals may have relatively good information about the importance of relationships to their SWB, their information about the importance of other factors is surprisingly poor (certainly in relative terms). Nonetheless, as we alluded to above, there are (at least) two concerns that might be raised about the robustness of these results, and which we attempt to address in the following section.

Robustness Checks

Adjusting for the quality of outcomes

The first concern relates to the fact that heterogeneity in the response of SWB to some outcomes may be driven by heterogeneity in the quality of the outcomes rather than heterogeneity in values *per se*. For example, the fact that the stated importance of having a fulfilling job does not predict individuals' SWB responses to employment could be due to many jobs not being especially fulfilling, rather than to inaccurate beliefs about what will bring satisfaction; or the failure of $p_{friendship}$ to predict the SWB returns to social time could be due to the quality of the friends, or of the time spent socialising with them. On the other hand it seems unlikely that variation in the quality consumption or health could be responsible for the irrelevance of individuals' stated priorities to their SWB reactions.

In an attempt to control for variation in the quality of jobs, social time, and intimate relationships, we rerun the regressions displayed in Tables 6.4 and 6.5, but now include additional interactions with measures of job satisfaction, satisfaction with social time, and satisfaction with one's relationship, q_{kit}. We interact these both with the main effects of the original variables, and with the existing interaction terms, so that the equations for the relevant

β coefficients are now:

$$\beta_{kit} = \beta_{k0} + \ddot{p}_{ki}\gamma_k + \ddot{q}_{kit}\theta_k + \ddot{p}_{ki}\ddot{q}_{kit}\phi_k \tag{6.10}$$

(and the equivalent version with time-varying information proxy p^*_{kit}). Note that this specification allows that private information could predict the *average* response to x_k as before (via γ_k), or the sensitivity to changes in "quality-adjusted" x_k, $q_{kit}x_{kit}$ (via ϕ_k). We centre both the proxies and the quality measures around their sample means prior to generating the interactions, so that the coefficients can always be interpreted as sample average values (e.g. $\beta_{relationship,0}$ will be the effect of an average quality relationship, for an individual who values relationships an average amount).

The results of regressions including the interactions specified in equation (6.10) are shown in Tables 6.6 and 6.7 (for p^{IMP} and p^{QOL} respectively). The main effects all have their usual signs, and the quality variables also behave as predicted: better relationships, more satisfying jobs, and higher-quality social time all contribute positively to SWB. However, in contrast to our expectations, controlling for the quality of outcomes does *not* improve the predictive ability of individuals' private information. If anything, it seems to have the opposite effect, with the previously positive $p_{relationship}$ becoming small and insignificant. This appears to be due to the fact that it is precisely those whose relationships are more satisfying who say that relationships are especially important to them (e.g. the correlation between the average 10-point importance of being in a relationship, and relationship quality is around 0.3). This suggests that the accuracy of individuals' predictions in Tables 6.4 and 6.5 may *reflect* accurate perceptions of the quality of their relationships, rather than being confounded by such concerns.

Again, the interactions with the other private information proxies are either insignificant (for income, consumption, reference income, jobs, and children) or significant, but wrongly signed. In particular, the interactions with p_{health} which before were marginally significant in only one specification, are now significantly negative at (at least) the 5% level in every column, while the interaction between $p_{friends}$ and quality-adjusted social time *is* significantly negative in the p^{QOL} specifications. Once again, it is possible to explain the negative effects as possible artefacts of subjectively reported health and quality-of-social-time measures, but in neither case does the evidence suggest that individuals' private information is especially good.

Comparability

Another possible question mark about individuals' apparently poor predictive ability relates to the comparability of the private information proxies between individuals. It seems reasonable enough to treat the binary p^{QOL}

Table 6.6: Do reported p^{IMP} values predict SWB responses? (Quality-Adjusted)

	$\bar{p}_{ki}$		p_{kit}	
(Log) Household Consumption	0.068***	(0.022)	0.067***	(0.022)
Consumption x p_{money}	-0.014	(0.012)	-0.001	(0.008)
(Log) Household Income	0.009	(0.009)	0.001	(0.028)
Income x p_{money}	0.002	(0.005)	0.001	(0.004)
Reference Income	0.000	(0.008)	-0.003	(0.023)
Reference Income x p_{money}	0.001	(0.004)	0.001	(0.004)
Health	0.137***	(0.007)	0.137***	(0.007)
Health x p_{health}	-0.023***	(0.008)	-0.013**	(0.006)
Non-cohabiting	0.278***	(0.027)	0.276***	(0.026)
Married	0.298***	(0.039)	0.297***	(0.039)
Couple	0.327***	(0.033)	0.326***	(0.033)
Widowed	0.021	(0.068)	0.013	(0.068)
Divorced	-0.055	(0.052)	-0.060	(0.052)
Separated	-0.137**	(0.058)	-0.141**	(0.058)
In Relationship x $p_{relationship}$	-0.018	(0.027)	-0.018	(0.021)
Relationship Quality	0.212***	(0.007)	0.212***	(0.007)
Relationship Quality x $p_{relationship}$	0.005	(0.005)	0.003	(0.004)
Employed	-0.067**	(0.030)	-0.067**	(0.030)
Unemployed (Seeking Work)	-0.346***	(0.053)	-0.248**	(0.097)
Unemployed (Discouraged)	-0.106***	(0.034)	-0.011	(0.086)
Employed x p_{job}	-0.030	(0.022)	-0.022	(0.017)
Unemployed x p_{job}	-0.012	(0.013)	-0.012	(0.011)
Job Quality	0.083***	(0.005)	0.083***	(0.005)
Job Quality x p_{job}	0.005	(0.003)	0.004	(0.003)
Hours Social Leisure	-0.011	(0.028)	-0.008	(0.028)
Hours Work	-0.033	(0.021)	-0.032	(0.021)
Hours Sleep	0.712***	(0.185)	0.715***	(0.185)
Hours Sleep^2	-0.042***	(0.010)	-0.042***	(0.010)
Social Hours x $p_{friends}$	-0.001	(0.012)	0.001	(0.010)
Quality-Adjusted Social Time	0.128***	(0.002)	0.128***	(0.002)
Quality-Adjusted Social Time x $p_{friends}$	-0.000	(0.001)	-0.000	(0.001)
Has Children	0.030	(0.020)	0.033*	(0.020)
2 Children	0.000	(0.021)	-0.000	(0.021)
3 Children	-0.039	(0.037)	-0.038	(0.037)
4 Children	-0.137	(0.086)	-0.135	(0.085)
Children x $p_{children}$	0.012	(0.009)	0.003	(0.007)
p_{money}			-0.009	(0.077)
p_{health}			0.070***	(0.023)
$p_{relationship}$			-0.009	(0.007)
p_{job}			-0.000	(0.005)
$p_{friends}$			0.009	(0.025)
$p_{children}$			0.006	(0.005)
Age	-0.000	(0.018)	-0.002	(0.018)
Age^2/10	0.001	(0.001)	0.001	(0.001)
Household Size (Adjusted)	-0.055***	(0.010)	-0.055***	(0.010)
Urban	-0.065**	(0.028)	-0.067**	(0.028)
Constant	1.233	(1.271)	1.268	(1.272)
Observations	54373		54373	
Groups	11345		11345	
R^2	0.448		0.481	

* p<0.10, ** p<0.05, *** p<0.01. Standard Errors (robust to clustering within pids) in parentheses. Wave dummies included but not reported

Table 6.7: Do reported p^{QOL} values predict SWB responses? (Quality-Adjusted)

	$\bar{p}_{ki}$		p_{kit}	
(Log) Household Consumption	0.078**	(0.035)	0.069***	(0.025)
Consumption x p_{money}	-0.011	(0.039)	0.001	(0.026)
(Log) Household Income	-0.000	(0.014)	0.004	(0.011)
Income x p_{money}	0.008	(0.017)	0.003	(0.014)
Reference Income	0.001	(0.011)	0.003	(0.010)
Reference Income x p_{money}	-0.002	(0.014)	-0.007	(0.013)
Health	0.162***	(0.014)	0.152***	(0.010)
Health x p_{health}	-0.034**	(0.016)	-0.027**	(0.013)
Non-cohabiting	0.270***	(0.026)	0.276***	(0.026)
Married	0.283***	(0.039)	0.292***	(0.038)
Couple	0.312***	(0.033)	0.322***	(0.032)
Widowed	0.007	(0.067)	0.000	(0.067)
Divorced	-0.059	(0.052)	-0.061	(0.052)
Separated	-0.137**	(0.058)	-0.136**	(0.058)
In Relationship x $p_{relationship}$	0.043	(0.135)	0.205	(0.162)
Relationship Quality	0.208***	(0.007)	0.213***	(0.007)
Relationship Quality x $p_{relationship}$	0.020	(0.019)	-0.015	(0.022)
Employed	-0.063**	(0.032)	-0.065**	(0.031)
Unemployed (Seeking Work)	-0.340***	(0.056)	-0.342***	(0.053)
Unemployed (Discouraged)	-0.104***	(0.035)	-0.104***	(0.034)
Employed x p_{job}	0.013	(0.083)	-0.100	(0.108)
Unemployed x p_{job}	-0.047	(0.084)	-0.044	(0.126)
Job Quality	0.087***	(0.005)	0.083***	(0.005)
Job Quality x p_{job}	-0.010	(0.011)	0.008	(0.011)
Hours Social Leisure	-0.004	(0.030)	-0.008	(0.029)
Hours Work	-0.031	(0.021)	-0.029	(0.021)
Hours Sleep	0.724***	(0.185)	0.706***	(0.185)
Hours Sleep^2	-0.042***	(0.010)	-0.041***	(0.010)
Social Hours x $p_{friends}$	-0.005	(0.031)	0.001	(0.028)
Quality-Adjusted Social Time	0.132***	(0.003)	0.131***	(0.002)
Quality-Adjusted Social Time x $p_{friends}$	-0.004***	(0.002)	-0.004**	(0.002)
Has Children	0.035*	(0.020)	0.035*	(0.019)
2 Children	0.000	(0.021)	0.001	(0.021)
3 Children	-0.037	(0.037)	-0.038	(0.037)
4 Children	-0.146*	(0.086)	-0.143*	(0.086)
p_{money}			0.045	(0.251)
p_{health}			0.137**	(0.055)
$p_{relationship}$			-0.078	(0.089)
p_{job}			0.081	(0.088)
$p_{friends}$			0.032	(0.077)
Age	0.003	(0.018)	0.003	(0.017)
Age^2/10	0.001	(0.001)	0.001	(0.001)
Household Size (Adjusted)	-0.055***	(0.010)	-0.055***	(0.010)
Urban	-0.063**	(0.028)	-0.069**	(0.028)
Constant	1.018	(1.273)	0.989	(1.274)
Observations	56883		57188	
Groups	13779		13920	
R^2	0.469		0.476	

* p<0.10, ** p<0.05, *** p<0.01. Standard Errors (robust to clustering within pids) in parentheses. Wave dummies included but not reported

measures as inter-personally comparable in (6.8) or (6.9) (in this case, x_k is either important or it is not). However, it is less clear that individual responses on the 10-point importance scales should be comparable in an obvious sense. The fact that the results (particularly those in Tables 6.4 and 6.5) are reasonably consistent between the p^{IMP} and p^{QOL} proxies, as well as the fact that the relationship proxy does appear to have reasonable predictive ability suggests that this may not be too much of a problem. However, we attempt a rudimentary robustness check below.

Conceptually, the possibility that that the information proxies may not be inter-personally comparable introduces a form of measurement error, and (for standard reasons) may lead to attenuation bias in the estimates. This could potentially provide an explanation for the insignificance of many of the interactions — though it would have more difficulty accounting for significant negative effects, such as that found for the health proxy in Tables 6.6 and 6.7.

To demonstrate the nature of the measurement error problem, begin by assuming that the reported importance of a variable x_k can be approximated by a stable, individual-specific affine transform of "true" importance $p_{kit} = a_{ki} + b_{ki}p^*_{kit}$. Rewriting this and interacting it with the the relevant x variable gives:

$$x_{kit}p_{kit} = x_{kit}p^*_{kit}\bar{b}_k + x_{kit}\bar{a}_k + [x_{kit}\tilde{b}_{ki}p^*_{kit} + x_{kit}\tilde{a}_{ki}] \qquad (6.11)$$

When this is included in the main equation, the $x_{kit}\bar{a}_k$ will be absorbed into the β_{k0} coefficient (which is also why any bias in the overall population information will not affect the interaction coefficients directly). The remaining terms in brackets have conditional expectations of 0 (by construction), and can thus be viewed as a form of classical measurement error.

Common strategies for dealing with measurement error (e.g. instrumental variables, or corrections using information about the measurement error variance) are not available to us here. Instead, we use the fact that individuals provide multiple answers to the importance questions (both for different variables with each wave, and for the same variable across waves) to construct proxies for the error. Intuitively, introducing these error proxies into our equations can provide a partial attenuation correction, by reducing the level noise in the importance estimates. A more formal account of how this works is included in Appendix A.10.

To generate proxies for the measurement error, we decompose the reported importance measures, $\bar{p}_{ki}$ (stable) and p_{kit} (time-varying) into separate components, as:

$$\bar{p}_{ki} = \bar{\bar{p}}_i + \tilde{\bar{p}}_{ki} \tag{6.12}$$

$$p_{kit} = \bar{\bar{p}}_i + \tilde{\bar{p}}_{ki} + \tilde{p}_{kit} \tag{6.13}$$

where $\bar{\bar{p}}_i = \frac{1}{T_i}\sum_\tau \frac{1}{K_{i\tau}}\sum_k p_{ki\tau}$ is an individual-specific constant, representing the average importance of all variables across all time periods, $\tilde{\bar{p}}_{ki} = \bar{p}_{ki} - \bar{\bar{p}}_i$ is the deviation from this average for a specific variable k, and $\tilde{p}_{kit} = p_{kit} - \bar{p}_{ki}$, is the time-specific deviation from these averages. Removing $\bar{\bar{p}}_i$ serves to capture any individual-specific heterogeneity in the intercept of the reporting functions, $a_i = \frac{1}{K}\sum_k a_{ki}$, and thus to help reduce attenuation bias in the results. Similarly, $\tilde{\bar{p}}_{ki}$ may capture the effect of variable-specific heterogeneity in the reporting function a_{ki} (though we expect this to be less important than a_i).

With these decompositions, we can then estimate the information content of each component separately using modified versions of equations (6.8) and (6.9):

$$\beta_{ki} = \beta_{k0} + \bar{\bar{p}}_i\gamma_{kp_1} + \tilde{\bar{p}}_{ki}\gamma_{kp_2} + \dots \tag{6.8'}$$

$$\beta_{ki} = \beta_{k0} + \bar{\bar{p}}_i\gamma_{kp_1} + \tilde{\bar{p}}_{ki}\gamma_{kp_2} + \tilde{p}_{kit}\gamma_{kp_3} + \dots \tag{6.9'}$$

The results of regressions using the modified equations (6.8') and (6.9'') are shown in Table 6.8. The coefficients of interest are the interactions between the main variables and the deviated proxies $\tilde{\bar{p}}_{ki}$ and $\tilde{p}_{kit}$. The coefficients on the interactions with $\bar{\bar{p}}_i$, while occasionally significant (for reference income and health they are weakly significant and negative) are not directly interpretable.[20] We see immediately that the $\tilde{\tilde{p}}_{relationship,i}$ and $\tilde{p}_{relationship,it}$ interactions remain highly significant in both specifications, confirming the quality of individuals' private information in this regard. Of particular interest is the fact that the interactions with $\tilde{\tilde{p}}_{friends,i}$ are now significant (though the interaction with the time-varying deviation $\tilde{p}_{friends,it}$ is not) — suggesting that individuals may also have accurate perceptions of the SWB impact of social time. In contrast, none of the other interactions is significant, which seems to reinforce our earlier conclusions about the generally poor quality of individuals' information.

6.2.4 Do individuals know the relative importance of different outcomes?

As we noted towards the end of the previous section, a potential drawback of the interaction-based approach to investigating the quality of individuals' private information is that we needed to assume that responses to the p^{IMP} and p^{QOL} questions were inter-personally comparable. Although we made some attempt to allow for this in the final robustness check, even this required us to assume some level of comparability. In contrast, this section adopts an alternative approach to assessing the quality of individuals' infor-

[20] As we show in Appendix A.10, we nonetheless need to include them to ensure the accuracy of the main interactions of interest.

Table 6.8: Do reported p^{IMP} values predict SWB responses? (Measurement Error Decompositions)

	$\bar{p}_{ki}$		p_{kit}	
(Log) Household Consumption	0.073*	(0.037)	0.071*	(0.037)
Consumption x $\bar{\bar{p}}$	0.036	(0.024)	0.036	(0.024)
Consumption x $\tilde{\bar{\bar{p}}}_{money}$	-0.021	(0.015)	-0.021	(0.015)
Consumption x $\tilde{p}_{money}$			0.009	(0.014)
(Log) Household Income	0.012	(0.013)	0.013	(0.013)
Income x $\bar{\bar{p}}$	-0.007	(0.011)	-0.007	(0.011)
Income x $\tilde{\bar{\bar{p}}}_{money}$	0.006	(0.005)	0.006	(0.005)
Income x $\tilde{p}_{money}$			-0.006	(0.008)
Reference Income	-0.005	(0.015)	-0.004	(0.015)
Reference Income x $\bar{\bar{p}}$	-0.016*	(0.009)	-0.015	(0.009)
Reference Income x $\tilde{\bar{\bar{p}}}_{money}$	0.001	(0.005)	0.001	(0.005)
Reference Income x $\tilde{p}_{money}$			-0.001	(0.007)
Health	0.213***	(0.015)	0.212***	(0.015)
Health x $\bar{\bar{p}}$	-0.021*	(0.011)	-0.021*	(0.011)
Health x $\tilde{\bar{\bar{p}}}_{health}$	-0.013	(0.010)	-0.013	(0.010)
Health x $\tilde{p}_{health}$			-0.004	(0.012)
Non-cohabiting	0.115***	(0.032)	0.115***	(0.032)
Married	0.159***	(0.047)	0.148***	(0.048)
Couple	0.220***	(0.040)	0.210***	(0.040)
Widowed	-0.124	(0.080)	-0.135*	(0.080)
Divorced	-0.076	(0.063)	-0.079	(0.063)
Separated	-0.336***	(0.068)	-0.335***	(0.068)
In Relationship x $\bar{\bar{p}}$	0.033	(0.028)	0.035	(0.028)
In Relationship x $\tilde{\bar{\bar{p}}}_{relationship}$	0.059***	(0.018)	0.061***	(0.019)
In Relationship x $\tilde{p}_{relationship}$			0.034**	(0.014)
Employed	-0.039	(0.037)	-0.042	(0.037)
Unemployed (Seeking Work)	-0.448***	(0.062)	-0.448***	(0.062)
Unemployed (Discouraged)	-0.174***	(0.043)	-0.176***	(0.043)
Employed x $\bar{\bar{p}}$	-0.027	(0.033)	-0.027	(0.033)
Employed x $\tilde{\bar{\bar{p}}}_{job}$	0.019	(0.017)	0.020	(0.017)
Employed x $\tilde{p}_{job}$			0.010	(0.010)
Unemployed x $\bar{\bar{p}}$	0.028	(0.039)	0.030	(0.038)
Unemployed x $\tilde{\bar{\bar{p}}}_{job}$	-0.007	(0.019)	-0.006	(0.019)
Unemployed x $\tilde{p}_{job}$			-0.008	(0.023)
Hours Social Leisure	0.117***	(0.035)	0.116***	(0.035)
Hours Work	-0.002	(0.025)	-0.002	(0.025)
Hours Sleep	1.049***	(0.220)	1.067***	(0.219)
Hours Sleep^2	-0.059***	(0.011)	-0.059***	(0.011)
Social Hours x $\bar{\bar{p}}$	-0.014	(0.024)	-0.013	(0.024)
Social Hours x $\tilde{\bar{\bar{p}}}_{friends}$	0.038**	(0.017)	0.039**	(0.017)
Social Hours x $\tilde{p}_{friends}$			0.002	(0.017)
Has Children	-0.001	(0.023)	-0.001	(0.023)
2 Children	-0.016	(0.025)	-0.014	(0.025)
3 Children	-0.055	(0.043)	-0.049	(0.043)
4 Children	-0.217**	(0.102)	-0.209**	(0.102)
Children x $\bar{\bar{p}}$	-0.009	(0.028)	-0.009	(0.028)
Children x $\tilde{\bar{\bar{p}}}_{children}$	-0.005	(0.014)	-0.005	(0.014)
Children x $\tilde{p}_{children}$			-0.016	(0.013)
$\tilde{p}_{money}$			-0.026	(0.126)
$\tilde{p}_{health}$			0.044	(0.047)
$\tilde{p}_{relationship}$			-0.011	(0.009)
$\tilde{p}_{job}$			-0.003	(0.007)
$\tilde{p}_{friends}$			0.015	(0.042)
$\tilde{p}_{children}$			0.010*	(0.006)
Age	-0.008	(0.019)	-0.008	(0.019)
Age^2/10	0.001	(0.001)	0.001	(0.001)
Household Size (Adjusted)	-0.065***	(0.011)	-0.065***	(0.011)
Urban	-0.091***	(0.034)	-0.093***	(0.034)
Constant	-1.091	(1.470)	-1.154	(1.466)
Observations	54804		54804	
Groups	11348		11348	
R^2	0.076		0.076	

* p<0.10, ** p<0.05, *** p<0.01. Standard Errors (robust to clustering within pids) in parentheses. Wave dummies included but not reported

mation, that does not rely on inter-*personal* comparability at all, and instead relies on inter-*outcome* comparability to attempt to assess whether individuals accurately perceive the relative importance of different outcomes to their SWB.

Similar to the analysis in Section 6.2.2, the approach we adopt here also seeks to compare the rankings of outcomes implied by the private information proxies with those obtained from SWB regressions — but this time seeks to do so at the individual-level, rather than for the overall sample. As in Section 6.2.2, this requires us to scale the β_k coefficients so that they are comparable, and we do this in the same way as we did there: multiplying the coefficients on the continuous variables (health, social time, and log consumption) by twice their standard deviation, and constructing a β_{rich} coefficient by multiplying the unstandardised β_{cons} by 0.51 (the log distance from the 3$^{\text{rd}}$ to the 5$^{\text{th}}$ sextile of the consumption distribution).

In contrast to 6.2.2, however, we now require individual-level coefficients on each of the relevant independent variables, rather than merely their sample averages. To obtain estimates of these coefficients, we use the results of the random coefficient model estimated in Section 5.2.2 to generate Empirical Bayes (EB) predictions for the individual level coefficients.[21]

Figure 6.1 displays the distribution of the individual-level Kendall's τ_b

[21] See e.g. footnote 30 in Chapter 5.

correlations between the ranking of outcomes implied by p_k^{IMP} and β_k (solid line), and p_k^{QOL} and β_k (dashed line). As we can see, the (p_k^{IMP}- and β_k-based rankings are positively correlated on average, but only weakly so, with a mean of 0.15 (roughly corresponding to a 58% chance of correctly ranking any two pairs of outcomes).[22] In contrast, the (p_k^{QOL}- and β_k-based rankings are barely correlated at all on average, with a mean of 0.01. This seems to confirm the results of the previous section, suggesting that individuals' private information is in general fairly poor, though the information content of the 10-point p^{IMP} measure is — perhaps unsurprisingly — greater than that of the binary p^{QOL} variable.[23]

Despite this relatively pessimistic conclusion, there does appear to be a fair amount of variation in the quality of information between individuals. The standard deviation of the $\tau_i(p_k^{IMP}, \beta_k)$ is 0.32 (suggesting that those one standard deviation above the mean in accuracy have a 16% higher chance of correctly ranking two outcomes than the average); while the standard deviation for the $\tau_i(p_k^{QOL}, \beta_k)$ is 0.44 (corresponding to a 22% higher chance

[22] This would be the exact interpretation for Kendall's τ_a. The correction for ties in the τ_b measure complicates this interpretation somewhat, but it nonetheless give a reasonable intuitive sense of the magnitude of the correlation.

[23] One possible concern with this analysis is that the low correlation between the stated and estimated importance of a these outcomes is that it could be due to imprecision in the individual-level coefficient estimates, rather than the quality of private information. One way to test this possibility would be to see whether the individual-level standard errors on the empirical Bayes predictions correlate with the individual-level Kendall's τ_b values. A negative correlation between the two would suggest that errors are indeed likely to be driven in part by inaccuracy in the $\hat{\beta}_i$s. Unfortunately, Stata does not currently provide a facility for calculating such standard errors.

of correct ranking). Attempting to find predictors of such variation seems a useful task for further research.

Figure 6.1: The distribution of individual-level Kendall's τ_b correlations

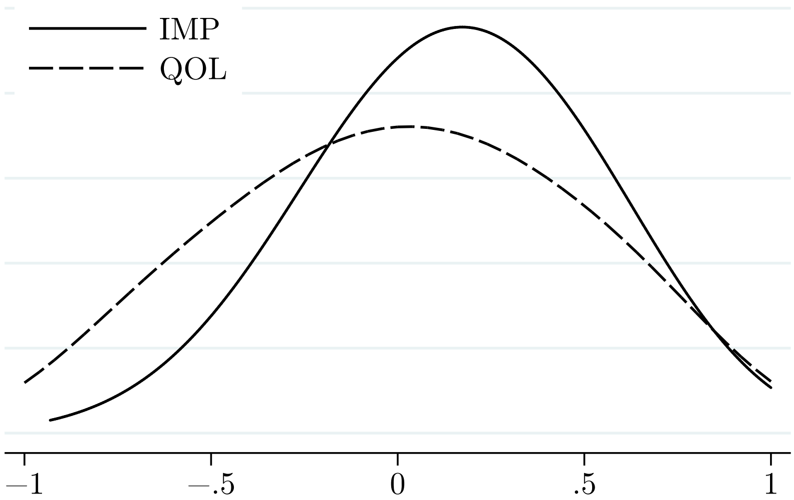

Kernel density estimates of the distribution of individual-level τ_b values: $\tau_{bi}(p_k^{IMP}, \beta_k)$ (solid line); and $\tau_{bi}(p_k^{QOL}, \beta_k)$ (dashed line). τ values of 1, -1, and 0 represent complete agreement in the rankings, complete disagreement in the rankings, and independence of the rankings respectively.

6.3 Conclusion

This chapter has examined one potential implication of finding that individuals differ significantly in what is important to them. We showed, in Section 6.1, that such variation diminishes the usefulness of population-average coefficient estimates to individuals in attempting to predict what will make them happy. However, the value of such estimates will typically also depend on the accuracy of individuals' private judgments.[24]

In Section 7.3.2, we then went on to provide three separate — and to our knowledge novel — tests of the quality of such judgments, using data from the BHPS. In contrast to much of psychology literature, which has tended to focus on biases that apply generally across many domains (and especially on impact bias), we focus on two slightly different concerns. First, in Section 6.2.2, we examine the possibility that individual judgments might be more biased in some domains than others. It is such *relative* biases that seem most likely to lead to poor decision-making, because they will distort the trade-offs that individuals perceive between different outcomes. Our results are consistent with the notion that people have roughly accurate perceptions of the relative value of health, relationships, friends, and jobs, but suggest that they may over-estimate the relative value of having children and money.

Second, even when people's judgments are relatively unbiased, they may still be very noisy predictors of their SWB responses. The issue of precision seems to have received less attention in the literature than bias, though it may be just as important in determining the value of the sorts of "surrogation" strategies advocated by Gilbert et al. (2009). Sections 6.2.3 and 6.2.4 therefore seek to assess the precision of individual judgments in two different

[24]Although we have not examined the issue in depth here, a similar set of considerations would seem to arise in certain policy contexts, when considering the extent to which decision-making should be centralised.

ways.

The first assumes that the BHPS measures of the stated importance of different outcomes are inter-personally comparable, and examines whether the SWB of those who say that a given outcome is especially important responds more strongly to changes in that outcome. We find that stated values generally have surprisingly poor predictive ability in this regard. In only one case — the importance of relationships — do we find a consistent positive relationship between stated values and estimated SWB responses. When we attempt to control for differences in the way individuals report the importance of outcomes, we also find that those who especially value jobs respond more negatively to unemployment. However, for income, consumption, reference income, health, social time, and having children, we find little evidence that individuals' stated values have useful information content.

The second test abandons the assumption of inter-personal comparability and instead examines the extent to which individuals' stated rankings of different outcomes agree with those implied by empirical Bayes estimates of their (individual-level) coefficients (derived from the LMM model we fit in Chapter 5). While the rankings of any two outcomes are more likely to agree than not, the correlation between the rankings is relatively weak — supporting the general tenor of our earlier conclusion about surprisingly poor predictive ability of individuals' stated values. Nonetheless, there does appear to be a fair amount of variation in the quality of different individuals' information, and the best are rather more accurate than the average. Whether those who are in general more accurate have any common observable characteristics seems a useful avenue for further research.

There are, as we have already noted, reasons to be cautious in the interpretation of these results. There are legitimate concerns about the extent to which individuals' stated values reflect their true beliefs (although the correspondence between the two alternative value proxies, p^{IMP} and p^{QOL}, suggests they are capturing *something*, as does the fact that *some* of the proxies perform fairly well). Moreover, we could also question the degree to which these measures are comparable either between individuals or between outcomes (although we have also adopted alternative empirical strategies that reduce the need to rely on such assumptions). Even with these concerns, however, the extent to which the private information measures generally fail to predict variation in individuals' actual responses is quite striking.

Such results lend support to existing evidence from the psychological and behavioural literature that throws doubt on the traditional neoclassical assumption that individuals are good judges of what will bring them satisfaction. They also extend these existing results in a two interesting ways, suggesting both that some outcomes (children and money) are overvalued more than others, and that, despite our generally poor predictive ability, we also have less noisy estimates of the importance of some outcomes (relationships and jobs) to SWB. Though we have suggested a range of possible explanations for such differences — including the possibility that we may adapt more to some outcomes than others — we have not been able to investigate this possibility here, and this remains an interesting avenue for further research. Finally, it is interesting to note that it is precisely those areas that exhibit the greatest population-level variation in responses (relationships and jobs) about which individuals appear to have the best private information, and those that exhibit the least population-level variation (children and money) where individual judgments appear most likely to be biased. This suggests that population-level surrogation is likely to be most useful for the latter, and least useful for the former.

Chapter 7

Conclusion

Recent years have seen an explosion in interest amongst economists in the determinants of SWB. This thesis has used data from Russian and British household surveys to examine three major issues in this burgeoning literature: the importance of distinguishing between consumption, income, and wealth as separate influences on SWB; the estimation of reference groups; and the scale, structure, and importance of individual differences in the determinants of SWB. Although the substance of the three Parts of the thesis share some common concerns, the issues they address are sufficiently distinct that we have chosen below to recapitulate the major and conclusions of each separately.

7.1 Consumption, Income, & Wealth

Part I of the thesis used Russian and British household panel data to explore the apparent disjunction between the traditional specification of the utility function — which holds that it is current consumption that contributes to current utility — and the standard approach to empirical SWB work, which (primarily for reasons of data availability) has focused instead on estimating the effect of income.

In methodological terms, our results strengthen the findings of previous research, which has suggested that the empirical focus on income tends to underestimate the impact of money on SWB. However, unlike earlier results, our findings are more robust to the presences of unobservable heterogeneity and possible biases arising from the use of financial-satisfaction as a dependent variable, rather than the more global measure of life-satisfaction that we employ. We have also attempted to use consumption measures that more closely align with best practice — in particular, via the incorporation of rental values for housing and durable consumption.

This result itself is entirely consistent with the implications of the traditional, consumption-based utility function: if income were simply a noisy measure of consumption, then we would expect income-based estimates of the importance of money to understate its impact. Indeed, it is in this light that that the Headey and Wooden (2004) results have been interpreted by subsequent authors such as Clark, Frijters, and Shields (2008). However, a major contribution of Part I is to provide evidence that this interpretation — and the exclusively–consumption-based model on which it is based — are incomplete in important ways. The standard theory predicts that, once we control for current consumption, income and wealth should have no further effect on current SWB. In contrast, we provide evidence that both variables influence SWB, and that they do so in a way that is difficult to attribute merely to their capturing measurement error in consumption.

Instead, our results suggest that income and wealth can both affect SWB independently of consumption. The wealth measures available in the RLMS are somewhat limited; however, the more extensive measures available in the BHPS allow us to conduct a fairly extensive model selection exercise to determine exactly *how* wealth affects SWB. In this respect, we advance the existing literature in four ways. First, we are able to resolve a tension between the approach taken by Headey and Wooden (2004) and Headey, Muffels, and Wooden (2008), and that employed by Brown, Taylor, and Wheatley Price (2005), as to whether it is individuals' overall net position that influences SWB, or whether they instead operate distinct mental accounts for assets and debts, and for different types of wealth. Our results strongly support the latter view. Second, we provide a novel test of the "hedonic editing" hypothesis put forward by Thaler (1985, 1999) and Thaler and Johnson (1990), which suggests that individuals may attempt to frame their mental accounts so as to maximise their SWB. In contrast to experimental results finding partial support for this view in other contexts, we find no evidence that individuals frame their wealth holdings in this way. Third, we examine whether individuals respond more to debt than to assets, and (again in contrast to existing results) find no evidence for debt-aversion. Finally, our results suggest that some types of wealth have a larger effect on SWB than others. In particular, long-term wealth holdings (such as investments) appear to have a smaller impact than either housing wealth, or shorter-term savings and debt. We also find that savings for specific items have a greater effect on current SWB.

Collectively, these findings suggest that at least part of the reason for the importance of wealth is rooted in the anticipation of future consumption, and we find some further support in the RLMS data for the idea that both income and wealth can influence SWB via the sense security they provide about the future. In contrast to the standard discounted utility model, which treats inter-temporal decisions as depending on a detached weighing of future against current consumption, these results lend support to an older theory of utility and inter-temporal consumption-savings behaviour. Championed by W.S. Jevons (1888), and even more so by his son, H.S. Jevons (1905), the idea that such decisions may depend significantly on the impact of future consumption on *current*, rather than just future utility has been recently resurrected in a number of behavioural models put forward by e.g. Caplin and Leahy (2001) and Köszegi (2005).[1] Some further implications of this broader view of utility can also be found in Frederick, Loewenstein, and O'Donoghue (2002).

In addition to the importance of future consumption, we also find evidence that the effect of income and wealth on current utility is — especially in Russia — partly related to the sense of status and/or self-esteem that they provide. As with anticipatory consumption, the inclusion of such factors in the utility function also recalls an older, more inclusive tradition in utility theory. Much of e.g. Duesenberry's (1949) relative income theory is heavily rooted in the importance of such status concerns — as of course is much of the more recent research into the effect of others' economic outcomes on SWB. One particularly interesting implication of incorporating such concerns is the finding that, in Russia, individuals appear to gain more SWB from income that they are somehow responsible for, than income they receive as a transfer.

Altogether, these findings point towards a rather broader conception of the ways in which economic outcomes can contribute to SWB than that

[1]Although, as we noted above, there are certain respects in which our finding that mental accounts moderate the SWB-impact suggests some glosses on the way in which these models incorporate future consumption into current utility.

embodied in the standard consumption-based model. As we might expect, the inclusion of such concerns may have a variety of implications for welfare analysis, of which we have sketched the outline of only a few here. These include e.g. the suggestion that consumption may be a less preferable basis for welfare and poverty analysis than is often claimed, and that the welfare cost of income volatility may be greater than the standard model would suggest. While we should be cautious about assuming that such findings will necessarily translate from the realm of experienced-utility to decision-utility, they may also have important implications for behaviour.

Another theme that emerges in these results is the importance of context in determining what matters to SWB. In at least one respect, our quantitative results differ quite markedly between Russia and Britain: while income is roughly as important to SWB as consumption in Russia, it is much less so in Britain (and in some specifications, fails to attain significance at all). Two major explanations for this difference appear to find some support in the data. The first is that income may be more important to status in Russia than in Britain. Consistent with this hypothesis, we find that earned income has a greater effect than unearned income in Russia, but not in Britain. The second is that income is especially important to individuals' sense of security about the future in volatile economic environments such as post-transition Russia, but less so in more stable circumstances. Consistent with this idea, income is relatively unimportant to SWB over most of the BHPS sample, but exerts a greater influence around the time of the stock market crash of 2000.

7.2 Reference Groups

In Part II of the thesis, we made two major contributions to the literature on the estimation of social effects — i.e. the effect of others' outcomes — on individual SWB. The first, building on the analysis of Part I, is to demonstrate the potential importance of distinguishing between the effects of others' income (which, again, has been the focus of most empirical work) and others' consumption.

Until recently, it has been a fairly standard finding in that others' income exerts a negative, comparison effect on individual SWB (see e.g. Clark, Frijters, and Shields, 2008, for an overview). However, in two recent papers, Senik (2004, 2008) has shown that this effect appears to reverse in Eastern Europe and the US, where there is substantial economic volatility, and/or perceptions of high mobility. In such circumstances, she argues, individuals take others' income as a signal of their future prospects, rather than a reference point against which to measure themselves.

We revisit this finding in the context of the RLMS, and, while confirming

Senik's result that others' income has a positive effect on SWB, find that others' *consumption* — which is likely to have less value as a signal, and more influence as a point of comparison — has a countervailing negative effect. We show further that, depending on how we define the relevant reference groups for income and consumption, we can obtain different conclusions about whether the information or the comparison effect dominates.

The sensitivity of our conclusions to the specification of the reference groups is what motivates the second main innovation of Part II. Instead of imposing a preconceived notion of the appropriate groups, we develop a novel estimation strategy that lets the data determine the reference groups and reference effects jointly. Our estimator generalises both the major existing approaches to specifying the reference groups, and can be thought of as either: a "fuzzy" cell-mean approach, which allows individuals (or households) outside a given (geographic, gender, age-group etc.) cell to affect the reference level to some extent; or as a non-linear kernel predictor of the relevant reference outcome.

Importantly, it offers the promise of identifying appropriate reference groups without the need to survey individuals on this point directly. In fact, it can potentially provide a data-based check on such survey responses, even when they are available. Indeed, in line with the results of existing survey responses, we find that reference groups are primarily local rather than global. However, we find that this is true only for some dimensions of social distance: some variables that predict income or consumption are nonetheless irrelevant in defining the reference group. In fact, the reference groups differ for income and consumption, with geography and household size the most relevant dimensions in defining the consumption referent, but productive characteristics more relevant to the referent in income (as we would expect if income acts primarily as a signal).

Estimating the reference effects jointly with the reference groups also allows us to gain a better idea of the net effect of others' economic outcomes on SWB. Regardless of whether we choose to focus on individual or household income, the negative effect of others' consumption consistently outweighs the positive effect of others' income identified by Senik. While information dominates comparison in respect of income, comparison dominates information overall.

As with any approach, the estimation of reference groups is limited somewhat by data availability. If, for example, reference groups are defined in terms of social networks, but we lack data on these, then we will be unable to identify this. The approach taken in Part II could also be extended in a number of ways to allow for even more flexibility in the specification of reference groups — for example by allowing for interactions between different dimensions of social distance (equivalent to non-Euclidean distance measures), or by allowing for asymmetry in upward and downward comparisons. Nonetheless, even with these data constraints, and the simplifying assumptions imposed, our estimator is still more flexible that existing alternatives, and allows us to tackle questions which existing approaches have difficulty addressing.

7.3 Heterogeneity in the Determinants of SWB

In Part III we turned to examine the scope, structure, and significance of individual differences in the determinants of SWB, with a specific focus on Britain. While there is a body of existing work that attempts to identify *observable* variation in what brings different individuals satisfaction, this can only provide a lower bound on the total amount of variation in the population. Moreover, while there have been some attempts to identify the full scale of both observable and unobservable differences, they have so far focused on only a small number of outcomes, such as the effect of income on financial satisfaction (Clark et al., 2005) or adaptation to marriage and divorce (Lucas et al., 2003). As such, they have been unable to address the possibility of correlations between the importance of different outcomes; and have also had difficulty distinguishing "true" variation in the determinants of SWB from variation in how individuals *report* SWB.

7.3.1 Identifying Heterogeneity

In Chapter 5, we use random coefficient models to simultaneously estimate not only the amount of unobservable variation in seven distinct determinants of SWB (consumption, reference income, health, unemployment, hours of work, relationships, and children), but also the correlations between different values. This allows us to advance the existing literature in a number of ways.

First, we find that there is indeed a large amount of variation in the way different individuals report responding to each of the outcomes we examine. In some cases, such as the effect of reference income or of having children, taking account of such individual differences seems especially important: despite the fact that the average effect of these variables is insignificantly different from zero, there appear to be both individuals who respond significantly negatively, and those who respond positively to changes in them. In respect of other outcomes too, we find that there is often variation, not only in the *degree* to which different individuals respond, but in the *direction* they respond. For example, while people respond positively to both consumption and relationships *on average*, there seem to be significant numbers for whom consumption breeds dissatisfaction, or whose relationships take a heavy emotional toll on them.

Second, and related to this latter point, we are able to go beyond previous work (e.g. Lucas et al., 2003) to verify that the appearance of oppositely-signed effects is not merely an artefact of the normality assumption relied on by the standard Linear Mixed Model. To do so, we employ a semi-parametric method, new to the SWB literature, to relax the normality assumption. An interesting task for future research would be to compare the performance of alternative parametric, semi-parametric and non-parametric (latent-class) approaches to estimating variation in coefficients.

Third, we find that there is indeed some underlying structure in the way different individuals respond to different outcomes. Some of this is related to observable demographic characteristics (especially age) but most of it is not. In the remaining unobserved variation, we find evidence that individuals' values tend to cluster to some degree around a "materialism" axis, with those who most value consumption and respond most negatively to others' income also tending to place less value on other outcomes such as health, employment, social time (relative to time spent working), relationships, and having children. In line with previous research, we also find that those who are more materialistic tend to have lower SWB than their observable characteristics would otherwise predict.

Fourth, we find that those who are least satisfied with their lives overall (whether for observable or unobservable reasons) tend to respond more strongly to changes in consumption, health and employment, and less to more "social" factors such as social time, relationships and children. They do not, however, appear to respond differently to others' income. As it stands, we do not know why these "SWB-poor" individuals respond to outcomes in the way they do, and in particular whether it is being SWB-poor that, say, leads relationships to be generally unfulfilling, or instead that these individuals are SWB-poor in part *because* they gain little from relationships. This is a potentially interesting avenue for further research.

These results also feed into out final conclusion: in contrast to earlier work, we are able to show that the variation we have identified in individual SWB-responses to different variables cannot be due simply to differences in the way they use reporting scales. We provide three pieces of evidence for this conclusion. First, as already discussed, we have found that some individuals respond in different directions to the same variables — a result that is difficult to explain with reporting heterogeneity alone. Second, if all heterogeneity were due merely to reporting differences, then this would imply both perfect correlations between the effects of different variables, and that the amount of variation in each coefficient should be proportional to its average effect. Both of these restrictions are conclusively rejected by the data. Finally, we show that there are significant inter-individual differences in the "consumption multiplier equivalents" of various outcomes. These measures, which define the trade-offs each individual faces between consumption and a given variable, are invariant to (affine) transforms of the reporting function, and variation in these is therefore unlikely to be significantly contaminated by reporting differences.

7.3.2 Do we know what makes us happy?

In Chapter 6 we go on to examine one potential implication of the results of Chapter 5: that large variations in individual SWB-responses to a given variable undermine the usefulness of information about population-average coefficients, making them less informative to both individuals and policy-makers. In both cases, however, the value of population-average information depends in part on the quality of the information that individuals themselves have about what is important to them. If people are poor judges of what will bring them satisfaction, then information about population-average responses will be more useful to them in attempting to predict their response to a change in circumstances. Moreover, while greater variation in individual responses increases the costs of centralised decision-making (e.g. the provision of a single level of health-care to all), it may nonetheless outperform more decentralised decision-making if individuals' private information is sufficiently poor.

The quality of individuals' predictions about what will make them happy has been a topic of interest to both psychologists and behavioural economists, who have accumulated a variety of evidence that individuals suffer from a number of biases in this regard. This chapter makes a novel contribution to this literature by examining whether individuals' stated values match their SWB-responses to changes in their situation. Although our results should be interpreted with some caution, they suggest that individuals have surprisingly poor information about what contributes to their satisfaction. Interestingly however, there appears to be a fair degree of variation in the quality of individual information, both between different domains and between different individuals.

Our results suggest that, on average, individuals appear to have roughly accurate senses of the relative importance of health, employment, relationships, and social time to SWB. We nonetheless find evidence they may systematically overestimate the SWB impact of having children, and — more tentatively — money. In contrast to existing results, which tend to focus on general biases, such *relative* biases are arguably more interesting — as they will tend to distort individuals' perceptions of the trade-offs they face between different outcomes.

Even if individual information is roughly accurate on average (i.e. unbiased) this in itself says little about the *precision* of individual judgments. We therefore go on to investigate this precision in two ways. First, we use an interaction strategy to see whether inter-personal variation in the stated importance of an outcome predicts variation in how individuals respond to changes in the corresponding variable. In general, we find little evidence that those who say that a given outcome is especially important to them respond especially strongly to such changes. There are however two major exceptions to this general pattern: (a) those who say that they especially value relationships tend to have relationships that are more satisfying, and to gain greater SWB from them; and (b) those who say that friends are especially important may respond more to changes in the amount of time they spend on social leisure.

In contrast to this first strategy, which relies on some level of inter-personal comparability in individuals' stated values, our second strategy instead focuses on the degree to which each individuals' stated ranking of outcomes matches the (empirical Bayes) predicted individual-level coefficients derived from the estimations conducted in Chapter 5. While the rankings of any two outcomes are more likely to agree than not (58% vs. 42%, based on the best performing measure of individual values), the effect is relatively small. Nonetheless, there is a reasonable amount of variation in the quality of different individuals' information, with those one standard deviation above the mean having around a 75% chance of correctly ranking any two outcomes.

Overall, these results — while somewhat discouraging from the perspective of individual decision-makers — do at least suggest that SWB research can usefully contribute to their knowledge. Together with the findings of earlier chapters, they provide further confirmation of the important role that such research can play in helping both individuals, and the governments that serve them, to advance the cause of human well-being.

Appendices

A.1 Preferences Over Income Streams

The inclusion of income and wealth in the instantaneous utility function generates preferences over income streams that do not exist in the standard, consumption-based model. The inclusion of income creates a preference for income-smoothing, while the inclusion of wealth generates a (competing) preference for "front-loaded" income streams, which will allow the benefits of wealth to be enjoyed in more periods.

This can be seen most easily in a stripped down consumption-savings model, where we assume the agent's problem is to maximise the (undiscounted[2]) sum of a generalised instantaneous utility function (additively separable in consumption, income and wealth) over a finite time horizon of length T. In order to focus attention on the relevant issues, we assume a

[2]Not discounting seems preferable from the perspective of welfare analysis, and serves to focus attention on the effects of interest.

zero interest rate, and fix both hours worked each period, and total income constant.

$$\max_{c_1,\ldots,c_T,y_1,\ldots,y_T} U = \sum_{t=1}^{T} \{f(c_t) + g(y_t) + h(A_t)\} \tag{A.1}$$

$$\text{s.t.} \qquad A_T \geq 0 \tag{A.2}$$

$$\sum_{t=1}^{T} y_t = Y \tag{A.3}$$

where c_t is consumption, y_t is income, and $A_t = \sum_{\tau=1}^{t}(y_\tau - c_\tau)$ represents the agent's asset holdings. To illustrate the key results, we consider a number of possible assumptions on the functions $f(\cdot)$, $g(\cdot)$, and $h(\cdot)$:

$$f'(x) > 0, \quad f''(x) < 0, \quad \lim_{x\to 0} f'(x) \to \infty \tag{F}$$

$$g(x) = 0 \tag{G0}$$

$$g'(x) > 0, \quad g''(x) < 0 \tag{G1}$$

$$h(x) = 0 \tag{H0}$$

$$h'(x) > 0, \quad h''(x) < 0 \text{ if } x > 0 \tag{H1}$$

Considering the four possible combinations of these in turn yields the following propositions.

Proposition 1. *If (F), (G0), and (H0) hold (i.e. consumption is the only* *argument in the utility function) then: (a) the optimal solution to the problem in (A.1)-(A.3) sets consumption constant at $c_t = Y/T$, for all $t \in T$; and (b) the agent is equally satisfied with any income stream that sets total income equal to Y.*

Proposition 2. *If (F), (G1), and (H0) hold (i.e. both consumption and income, but not wealth matter) then: (a) Proposition 1(a) continues to hold; and (b) the optimal solution to (A.1) now sets income constant at $y_t = Y/T$, for all $t \in T$.*

Proposition 3. *If (F), (G0), and (H1) hold (i.e. consumption and wealth, but not income matter) then: (a) the optimal consumption profile is increasing over time; and (b)(i) the agent prefers any income profile y' to another $y = \{y_1, \ldots, y_T\}$ if it is more "front-loaded" (i.e. if $y' = \{y_1, \ldots, y_s + \delta, \ldots, y_t - \delta, \ldots, y_T\}$ for any $s < t$, $\delta > 0$); and (ii) the optimal income profile is "fully front-loaded", $y^* = \{Y, 0, \ldots, 0\}$.*

Proposition 4. *If (F), (G1), and (H1) hold (i.e. consumption, income, and wealth all matter) then: (a) The optimal consumption profile is increasing over time; and (b)(i) the optimal income profile is strictly declining over time, unless and until it reaches 0 (after which point it is constant); but (ii) the optimal income profile is not, in general, fully front-loaded.*

Propositions 2(b), 3(b) and 4(b) show that, in all three cases where income and or wealth enter into the utility function, this generates preferences over income streams that do not exist in the standard consumption-only model (Proposition 1(b)). Moreover, Propositions 3(a) and 4(a) show that the inclusion of wealth in the utility function generates an incentive towards increasing consumption streams not otherwise present.[3]

With the addition of non-negativity constraints on the y_t, the Lagrangian for the problem in (A.1)- (A.3) becomes:[4]

$$\mathcal{L} = \sum_{t=1}^{T} f(c_t) + \sum_{t=1}^{T} g(y_t) + \sum_{t=1}^{T} h\left(\sum_{\tau=1}^{t}(y_\tau - c_\tau)\right) + \lambda \sum_{t=1}^{T}(y_t - c_t)$$
$$- \gamma \left(\sum_{t=1}^{T} y_t - Y\right) + \theta_1 y_1 + \ldots + \theta_T y_T \tag{A.4}$$

with FOCs:

$$c_t: \quad f'(c_t) - \textstyle\sum_{\tau=t}^{T} h'(A_\tau) - \lambda = 0, \qquad \forall\, t \in T \tag{A.5}$$

$$y_t: \quad g'(y_t) + \textstyle\sum_{\tau=t}^{T} h'(A_\tau) + \lambda - \gamma + \theta_t = 0, \quad \forall\, t \in T \tag{A.6}$$

Proof of Proposition 1. The result is well-known and the proof straightforward.[5]

[3] Relaxing the assumption of zero interest creates similar incentives, but the point here still stands.

[4] The assumptions we place on $f(\cdot)$ in (F) above mean we can ignore issue of non-negativity constraints on c_t.

[5] Relaxing the assumptions of zero interest and fixed hours does generate preferences over income streams. But the point of the analysis — that including income, and wealth in the utility function generates additional preferences that do not exist in these models

Proof of Proposition 2. The proof is again straightforward, and follows in the same manner as the consumption result in Proposition 1(a).

Proof of Proposition 3. To show Proposition 3(a), subtract any FOC c_t from c_s, where $s < t$:

$$c_s - c_t : \quad f'(c_s) - f'(c_t) = \textstyle\sum_s^{t-1} h'(A_t) > 0, \quad \forall\, s < t \tag{A.7}$$

$f''(\cdot) < 0$ immediately implies that $c_s < c_t$. Proposition 3(b)(i) follows from the fact that the marginal utility of income (through the wealth term in the utility function) is always greater in an earlier period than in a later period:

$$\frac{dU}{dy_s} - \frac{dU}{dy_t} = \sum_s^{t-1} h'(A_t) > 0, \quad \forall\, s < t \tag{A.8}$$

Proposition 3(b)(ii) then follows immediately with the addition of non-negativity constraints on y_t. Subtracting any FOC y_s from y_t, where $s < t$, gives:

$$y_t - y_s : \quad \theta_t - \theta_s = \textstyle\sum_s^{t-1} h'(A_t) > 0 \quad \forall\, s < t \tag{A.9}$$

which implies that $\theta_t > 0$, $\forall\, t \neq 1$, and thus, that income will be zero in every period except the first (where it will be equal to Y).

— is not affected.

Proof of Proposition 4. Proposition 4(a) follows in exactly the same manner as Proposition 3(a), with the only change being that, if income is not fully front-loaded, A_t will now be smaller in each period, and the RHS, $\sum_s^{t-1} h'(A_t)$will be greater, implying a fall in initial consumption, and a concomitant increase in the rate of consumption growth. Proposition 4(b)(i) can be seen as follows. We first show that, if income is positive in any period s, then $y_s > y_t$ for all $t > s$. If $y_t = 0$, then this is obvious. If not, then subtracting the FOC on y_s from the FOC on y_t gives:

$$y_t - y_s: \quad g'(y_t) - g'(y_s) = \textstyle\sum_s^{t-1} h'(A_t), \quad \forall\, s < t \tag{A.10}$$

and hence, $y_t < y_s$ (with $g''(\cdot) < 0$ and $h'(\cdot) > 0$). Second, we show that if $y_s = 0$ then $y_t = 0$, for all $t > s$. Allowing the constraints to bind, (A.10) becomes:

$$y_t - y_s: \quad \theta_t - \theta_s = \textstyle\sum_s^{t-1} h'(A_t) + g'(0) - g'(y_t) > 0 \quad \forall\, s < t \tag{A.11}$$

The final inequality holds because $g''(x) < 0$ implies $g'(0) \geq g'(y_t)\, \forall\, y_t$, and immediately implies that $\theta_t > \theta_s > 0$, which is what we require. Proposition 4(b)(ii) will hold provided that there exist circumstances under which (A.10) holds for $s = 1$ and $t = 2$. While it is not difficult to generate examples when this will be the case, more generally, precisely when any of the constraints will bind, and thus, whether earning zero income in some periods will be optimal, depends in a rather complicated way on the nature of all three sub-utility functions. In general, it can be seen from (A.10) that the steepness of the income decline (and hence the probability of the constraints binding) will be greater:

(1) the greater is the marginal utility of wealth, $h'(A_t)$;

(2) the lower the agent's level of risk aversion, $-\frac{g''(y_t)}{g'(y_t)}$ in income; and

(3) the greater is the agent's level of risk aversion in consumption, $-\frac{f''(c_t)}{f'(c_t)}$.

The first two follow from inspection. (3) follows because increased consumption smoothing will decrease A_t in each period, and thus increase $h'(A_t)$ (because $h''(A_t) < 0$). Intuitively, the more risk averse an agent is in consumption, the more she will prefer to increase wealth by front-loading income rather than "rear-loading" consumption.

A.2 RLMS Variable Construction

RLMS Individual Income

RLMS Individual Income is the sum of individual main cash salary, individual non-cash salary, additional salary, additional non-cash salary, other paid work salary, amount of pension income, and unemployment benefits. Where individuals report receiving income from a given source, but do not report the amount, the amount is imputed using the median value conditional on settlement type, gender and age category (if possible). In general, less than 1% of any particular variable is imputed, and it never materially affects the characteristics of the variable.

RLMS Household Income

RLMS Household Income is the sum Individual Income for each individual in the household, plus the value of home production sold, consumed or given away. Where households report home production of a given item, but do not report the amount, it is imputed using the median conditional on settlement type and family size (if possible).

RLMS Household Expenditure

RLMS Household Expenditure is the sum of expenditure on food, clothing, durables, luxury durables, fuel, services, rent, bonds, insurance, alimony and loan repayments, plus savings, and home-produced goods consumed or given away. As above, where households report expenditure on a given item, but do not report the amount, it is imputed using the median conditional on settlement type and family size (if possible). For those data on which substitution counts are reported, the total imputation never exceeds 1% of the data.

Household Expenditure

Our Household Expenditure measure is the RLMS Household expenditure measure minus reported savings, expenditure on bonds and loan repayments.

A.2.1 Household Consumption

Construction of our household consumption measure follows the guidelines set out in Deaton and Zaidi (2002). It consists of the RLMS Household expenditure measure less reported savings, expenditure on bonds, loan repayments, durables and luxury durables (which are technically additions to wealth and not consumption), plus estimates of (a) the rental value of durables owned, and (b) the rental value of housing consumed.

Constructing the Rental Value of Durables

Unfortunately, while the RLMS has data on the ownership and age of a variety of consumer durables, it does not have information on the value of individual items owned, except in Round 5.[6] In the absence of such data, we use the Round 5 data to impute predicted values for durables with particular

[6] Moreover, because the reported expenditure categories are wider than the ownership categories – e.g. "Cultural Goods" and "Appliances" rather than "TV", "Video" or "Computer" and "Refrigerator", "Freezer", or "Washer" – we cannot reliably use expenditure data to impute such values.

characteristics. This can only give us values for those items asked about in round 5, and therefore excludes Trucks, Hair-Dryers, and Computers, which were only added to the survey later. We also drop tractors and other apartments, because we consider them more likely to be investment than consumption.[7]

Following Deaton and Zaidi (2002), we assume a constant depreciation rate for each of the remaining durable items:

$$V_a = V_0(1 - \delta)^a$$

where V_a is the value of an item of age a, and V_0 is it's initial value when new. Taking logs gives:

$$\ln V_a = \ln V_0 + a \ln(1 - \delta)$$

If we further specify that:

$$\ln V_0 = \beta_0 + \beta_1 \ln c_i + \gamma z_i + \varepsilon_i$$

where c_i is household non-durable expenditure (intended to capture the general effect of household wealth on durable value) and z_i is a vector of region

[7] In any event, the regressions outlined do a poor job of predicting these values.

and settlement-type dummies,[8] this yields:

$$\ln V_a = \beta_0 + \beta_1 \ln c_i + \gamma z_i + \theta a + \varepsilon_i$$

The results of running this regression on the Round 5 data (shown in A.1 and A.1[9]) gives us two pieces of information. First, the estimated coefficient on a, $\hat{\theta}$ is an estimate of $\ln(1-\delta)$. Second, it allows us to predict values for the stock of durable good g, owned by family i, in period t as:[10]

$$\hat{V}_{g,a(t)} = \exp\left(\hat{\beta}_0 + \hat{\beta}_1 \ln c_i + \hat{\gamma} z_i + a(t)\hat{\theta}\right)$$

As Deaton and Zaidi (2002) show, from this information, coupled with Bank of Russia monthly data on the nominal interest rate, we can calculate the user cost of good g in period t as:

$$\begin{aligned} D_{gt} &= \hat{V}_{g,a(t)}(r_t + \delta_g) \\ &= \hat{V}_{g,a(t)}(r_t + (1 - \exp(\hat{\theta}))) \end{aligned}$$

[8] The latter are excluded from the the regressions for freezers and cottages, as there are not sufficient observations here to be useful.

[9] Unfortunately, the R^2s for the regressions are in general low. However, we have few alternative available to us.

[10] In fact, as Goldberger (1968) shows, the log transformation means that $\hat{V}_{g,a(t)}$ is not an unbiased estimate of $V_{g,a(t)}$. However, because we take logs of both rental and stock values before using them in our regressions, we do not "correct" for this here.

Table A.1: Durable Values Regressions (a)

	Refrigerator		Freezer		Washer		TV(B&W)		TV(Colour)	
(Log) Household Consumption	0.128**	(5.622)	-0.028	(-0.491)	0.147**	(5.308)	-0.017	(-0.568)	0.086**	(3.744)
Age	-0.047**	(-21.598)	-0.064**	(-6.770)	-0.032**	(-12.404)	-0.031**	(-8.921)	-0.061**	(-17.082)
North & Northwest	0.298**	(3.712)	0.096	(0.459)	0.143	(1.424)	0.468**	(3.721)	0.323**	(4.580)
Central	-0.067	(-1.014)	-0.301	(-1.599)	-0.215*	(-2.545)	0.116	(1.071)	-0.012	(-0.215)
Volga	0.082	(1.189)	-0.546**	(-3.487)	-0.287**	(-3.274)	0.352**	(3.146)	0.069	(1.136)
North Caucasus	0.050	(0.686)	-0.117	(-0.442)	-0.325**	(-3.502)	0.117	(0.985)	-0.015	(-0.229)
Urals	0.202**	(2.949)	-0.337	(-1.912)	-0.121	(-1.407)	0.431**	(3.869)	0.240**	(3.889)
East Siberia	0.068	(0.902)	-0.120	(-0.685)	0.034	(0.364)	0.164	(1.363)	0.247**	(3.590)
West Siberia	-0.020	(-0.252)	-0.354*	(-2.377)	-0.300**	(-3.154)	0.119	(0.979)	0.241**	(3.520)
PGT	0.140	(1.884)			0.131	(1.470)	0.151	(1.344)	0.093	(1.363)
Rural	-0.041	(-0.969)			0.056	(1.126)	0.285**	(5.198)	0.030	(0.666)
Constant	8.440**	(37.101)	10.279**	(17.716)	7.444**	(26.758)	8.701**	(28.392)	9.058**	(39.736)
Observations	2781		180		2356		1516		2048	
R^2	0.185		0.293		0.105		0.088		0.165	

* p<0.05, ** p<0.01.
t-statistics in parentheses.

Table A.2: Durable Values Regressions (b)

	VCR		Car		Motorcycle		Dacha		Cottage	
(Log) Household Consumption	0.048	(1.316)	0.208**	(3.997)	0.024	(0.328)	0.269**	(2.616)	0.289**	(3.222)
Age	-0.049**	(-3.618)	-0.051**	(-10.470)	-0.008	(-1.363)	0.001	(0.270)	-0.008	(-1.397)
North & Northwest	0.458**	(4.129)	0.226	(1.370)	1.063**	(2.969)	-0.851**	(-3.189)	-1.618**	(-5.032)
Central	-0.064	(-0.923)	-0.125	(-0.871)	0.650	(1.903)	-0.982**	(-3.979)	-0.831**	(-3.570)
Volga	0.100	(1.104)	0.106	(0.677)	0.643	(1.889)	-1.353**	(-5.227)	-0.969**	(-3.752)
North Caucasus	0.349**	(3.417)	0.002	(0.016)	1.182**	(3.458)	-0.845**	(-2.972)	-0.580	(-1.608)
Urals	0.202	(1.926)	0.239	(1.599)	0.990**	(2.983)	-1.416**	(-4.935)	-0.967**	(-4.264)
East Siberia	0.274**	(2.744)	0.104	(0.680)	0.964**	(2.733)	-1.380**	(-4.336)	-1.166**	(-4.466)
West Siberia	0.404**	(4.209)	0.100	(0.652)	0.980**	(2.911)	-0.909**	(-3.081)	-1.120**	(-4.460)
PGT	-0.068	(-0.566)	0.351*	(1.997)	0.618**	(2.821)	-1.161**	(-4.633)		
Rural	-0.025	(-0.331)	0.075	(0.869)	0.123	(1.207)	0.462	(1.355)		
Constant	9.335**	(25.763)	10.354**	(19.504)	9.103**	(11.629)	10.204**	(9.852)	9.458**	(10.180)
Observations	462		683		310		335		356	
R^2	0.133		0.186		0.066		0.218		0.104	

* $p<0.05$, ** $p<0.01$.
t-statistics in parentheses.

Constructing the Rental Value of Housing

In each round, the RLMS includes questions regarding the amount each household spent on rent and municipal services in the past month. However, for a number of reasons, the answers to these questions are unlikely to give us useful measures of housing consumption. This is most obvious for privatized residences where no rent is due, but even for non-privatized residences, the raw data is likely to be problematic.

(1) The rental owed by most households does not reflect the market value of their accommodation. Most accommodation is subsidised by local governments, who typically require rental payments well below cost recovery. While there has been a general movement towards fuller cost recovery over the course of the sample (with reported payments towards rent and utilities rising from 3.1% of total expenditure at the end of 1994 to 8.7% by the end of 2004[11]) this has been sporadic and subject to significant variation between regions (see e.g. Lykova et al., 2004).

(2) Even then, the rental *reported* by many households may not reflect the amount actually owed. A significant portion of households have rent outstanding (26% across the entire sample), and in contributing to,

[11] This compares with a housing share of around 20% in most Western countries. Though Clements, Wu, and Zhang (2006) note that the housing share tends to decrease with GDP per capita, the RLMS reported values are low, even by developing country standards. Moreover, general rents are consistently below those reported by market renters (which rise from a 4.1% to an 11.1% share over the same period).

or attempting to reduce such debts, households may have paid either more or less than the amount due for the relevant period. (Evidence of underpayment can be seen most clearly in the case of the 27% of households living in non-privatized homes who report paying no rent at all in the previous month.)

(3) Due to the construction of the sample (surveyors revisited the same *dwellings* every wave, rather than attempting to track households that moved), there is likely to be little true variation in housing consumption. What variation there is is likely to consist mainly of changes in consumption of utilities such as power, heating, and water (which are not recorded separately except in Round 5).

The collective effect of these factors is to introduce both a sizeable amount of spurious variation, and a significant downward bias into the reported housing consumption measures, both of which are likely to result in a bias in coefficient estimates relying on them.[12] In an attempt to overcome these problems, we replace reported rental measures for *all* households with values

[12] In the case of low housing consumption values, biased coefficients will result unless the measurement error reduces our measures of overall consumption by a constant proportion (in which case taking mean deviations for the fixed effects will remove the bias). For other patterns of bias the sign is difficult to predict formally, though simulations suggest that, for a simple OLS or fixed effects model with a constant intercept, a constant linear downward bias in the untransformed independent variable typically results in a downward bias in its coefficient (and an upward bias in the intercept).

imputed from hedonic housing regressions, of the following form:

$$y_{it} = x'_{it}\beta + \varepsilon_{it}$$

where y_{it} is the log of rent paid, and x_{it} is a vector of household characteristics, including logged floor space, central heating, cold water, hot water, gas or electric stove, central sewerage, telephone connection, and housing type dummies.[13] Given the structure of the data set, we expect the resultant loss of information on true variation to be outweighed by the removal of noise due to changes in subsidies and under or over-payment.

In running these regressions, we take a number of steps in order to minimise the potential for downward bias in our estimates.

(1) First, to minimise bias due to public subsidies, we estimate hedonic housing regressions only on the subset of households paying market rents. These constitute only 5.6% of the sample (and yield only 1324 usable observations). However, the results appear to outperform the

[13] We impute missing values in housing characteristics using the following rules (based on those adopted by Romanik, 1998).

(1) Where households do not provide information on a particular question in one round, but do answer it in others, we impute values using the following algorithm:
    (a) replace missing values with lagged values from the previous round;
    (b) replace remaining missing values with lead values from the subsequent round;
    (c) repeat (1) and (2) until all possible missing values have been imputed.
(2) Where household living space, but not total floor space is reported, we impute missing values of the latter using the median living space to floor space ratio for the relevant housing type.

alternative strategies of using either the full sample, or only Round 13 data (the latter on the theory that these rents will in general be closer to full cost recovery levels).[14] In general, they predict rent levels nearer to those obtained by Tesliuc and Ovtcharova (2006), who, using better quality NOBUS data (with both a more extensive set of housing characteristics, and predicted market rents for non-market households) estimate an under-reporting of rent of around 62%. The corresponding under-reporting estimates for the market rent, Round 13, and full sample models, are 55%, 47%, and 24% respectively. The signs on all coefficients are as expected, and while the R^2 of 0.35 is somewhat low, it is higher than that for the full sample (0.34) and only slightly lower than that for the Round 13 model (0.42). Moreover, to the extent that the latter does explain more of the variation in the reported rents, it is likely to be due partially to regional differences in subsidies.

(2) Second, to minimise the effect of underpayment, we exclude households that report paying zero rent.

(3) Finally, and also to avoid the potential for bias due to underpayment, we exclude (log) household income from the independent variable set.

[14] Moreover, as Deaton and Zaidi (2002) suggest, the hedonic imputation procedure is also likely to outperform an approach based on imputation back from market housing values. The latter is highly sensitive to assumptions about the user cost (which we have little basis for making) and in any event, we have estimated market values for only a few rounds.

While household income is typically included in hedonic housing regressions, on the theory that it is likely to proxy for unobserved household quality, two factors decrease its usefulness in this context. First, to the extent that housing allocation under the Soviet system was less dependent on a family's means than under a market system (and Russia's still relatively undeveloped housing market means that Soviet era allocations have remained reasonably stable), income may be only weakly correlated with true housing consumption. Second, and more importantly, swings in income are also likely to be highly correlated with under or overpayment of rental obligations. Consistent with this idea, household income strongly influences both the probability of a household owing money for rent, and the amount owed. Moreover, hedonic housing regressions that include household income as an independent variable consistently yield lower imputed rents than those which exclude it, suggesting that, conditional on observed characteristics, the latter correlation may be higher than the former.

From the results of this regression (presented in Table A.3), we generate predicted log rental consumption, $\hat{y_{it}}$. Taking $\exp(\hat{y}_{it})$ then yields the imputed values for housing consumption.[15]

[15] As noted in footnote 10, Goldberger (1968) shows that $\exp(\hat{y}_{it})$ is not an unbiased estimate of $\exp(y_{it})$. However, because we take logs of consumption as the explanatory variable in our regression, we do not "correct" for this here.

Table A.3: Hedonic Housing Regression

	MarketRental	
Log Floor Space	0.088	(1.213)
Heating	0.530**	(4.689)
Cold Water	0.149	(1.293)
Hot Water	0.249**	(2.587)
Electric Stove	0.294**	(3.420)
Central Sewerage	0.131	(1.008)
Phone	0.134	(1.839)
North & Northwest	-0.254	(-1.462)
Central	-0.447**	(-3.123)
Volga	-0.331*	(-2.218)
North Caucasus	0.140	(0.892)
Urals	-0.372**	(-2.605)
East Siberia	-0.307	(-1.916)
West Siberia	-0.259	(-1.778)
PGT	-0.197	(-1.376)
Rural	-0.643**	(-7.215)
Constant	5.570**	(18.244)
Observations	1324	
R^2	0.350	

* p<0.05, ** p<0.01.
t-statistics in parentheses.

A.3 Attrition In The RLMS

This appendix describes the inverse probability weighting scheme I apply to correct for possible attrition bias in the results. Section A.3.1 gives an overview of the method. Section A.3.2 describes our implementation.

A.3.1 An Overview of Inverse Probability Weighting

Fitzgerald, Gottschalk, and Moffitt (1998), Wooldridge (2002) and others have shown that it is possible to correct for this by using inverse probability weighted estimators if we can find a vector z_{it} that satisfies the following assumption:

$$\mathrm{P}[s_{it} = 1 | y_{it}, x_{it}, z_{it}] = \mathrm{P}[s_{it} = 1 | z_{it}] \tag{A.12}$$

where z_{it} is observed for all s_{it}. Because z_{it} is always observed, we can estimate the selection probabilities π_{it} via a standard logit or probit of s_{it} on z_{it}, and find consistent parameter estimates by minimising the function:

$$\min M = \sum \left(\frac{s_{it}}{\hat{\pi}_{it}} \right) q_t(w_{it}, \theta)$$

where $q_t(w_{it}, \theta)$ is the general function to be minimised under a random sampling assumption. Under assumption (A.12), and standard regularity conditions, M has the probability limit:

$$\begin{aligned} \text{plim} M &= \mathrm{E}\left\{ \mathrm{E}\left[(s_{it}/\pi_{it}) \cdot q_t(w_{it}, \theta) \right] | w_{it}, z_{it} \right\} \\ &= \mathrm{E}\left\{ \left[\mathrm{E}(s_{it} | w_{it}, z_{it}) / \pi_{it} \right] \cdot q_t(w_{it}, \theta) \right\} \\ &= \mathrm{E}\left[q_t(w_{it}, \theta) \right] \end{aligned}$$

If these assumptions hold, the IPW estimator will be consistent regardless of the presence of attrition bias. However, it suffers from the weakness that it requires us to restrict our analysis to the artificial monotone sample (i.e. the sample constructed by imposing $s_{(i,t)} = 0$ if $s_{(i,s)} = 0$ for any $s < t$, where s is a selection (non-attrition) indicator) which means that we lose approximately $\frac{2}{3}$ of our observations (and even more in the case of the wealth regressions that primarily use data from later survey rounds).[16]

A.3.2 Generating Inverse Probability Weights

We generate the predicted inverse probability weights by running separate attrition probits for each round. This gives predicted response probabilities $\hat{\lambda}_{it} = \mathrm{P}[r_{it} = 1 | r_{i,t-1}, z_{it}]$, which can be used to calculate the $\hat{\pi}_{it}$ as $\hat{\pi}_{it} = \prod_{\tau=1}^{t} \hat{\lambda}_{i\tau}$. For the estimation, we include in our vector z_{it}, a progressively expanding set of non-response predictors. In all rounds we include lags of household head category dummies, family size, the proportions of children and elderly family members, settlement type dummies, region dummies, marital status dummies, age, age squared, education, and gender. We also include individual-level means (for the previous $t-1$ periods) of life satisfac-

[16] As discussed in e.g. Robins and Rotnitzky (1995) and Robins, Rotnitzky, and Zhao (1995), the case of non-monotone attrition, where attriters are allowed to (re)enter in subsequent periods is rather more involved, requiring additional assumptions, and a rather more complicated estimation procedure.

tion, log household income, log household consumption, health, employment status, and a number of surveyor-assessed attitude variables relating to both the individual, and household survey respondents (Mu, 2003, has previously found these to be useful predictors in this context). From Round 7 onwards we introduce mean deviations of the same variables, and from Round 8 onwards, we also include their sample standard deviations. The results of these regressions are displayed in Table A.4 and Table A.4 below. As Fitzgerald, Gottschalk, and Moffitt (1998) note, the low R^2s suggest that the resulting corrections are likely to be fairly small — a prediction that largely seems to be borne out in practice.

We then estimate a probit selection equation on the full panel in each round to generate predicted probabilities for s_{it} for each individual, and use the inverse of these to weight the fixed effects regressions (implemented using Stata's areg command) reported in the text.

Table A.4: Response Probits

	Round 6		Round 7		Round 8		Round 9	
Mean (Log) Household Income	-0.071**	(-2.162)	0.026	(0.564)	-0.070	(-1.164)	-0.026	(-0.335)
Mean Dev. (Log) Household Income			0.028	(0.445)	-0.043	(-0.887)	-0.058	(-0.991)
S.d. (Log) Household Income					-0.044	(-0.464)	-0.106	(-0.971)
Mean (Log) Household Consumption	0.087**	(2.349)	-0.024	(-0.458)	-0.001	(-0.021)	0.002	(0.023)
Mean Dev. (Log) Household Consumption			0.010	(0.114)	-0.128	(-1.572)	-0.042	(-0.411)
S.d. (Log) Household Consumption					-0.470***	(-3.067)	-0.411**	(-1.980)
Adult Female Head	0.009	(0.102)	0.221**	(2.231)	0.150	(1.352)	0.011	(0.079)
Elderly Male Head	0.151	(0.999)	0.080	(0.477)	0.087	(0.492)	0.070	(0.313)
Elderly Female Head	0.253	(1.563)	0.292	(1.600)	0.147	(0.755)	-0.052	(-0.222)
Young Head	-0.028	(-0.076)						
Household Head Changed			-0.012	(-0.105)	-0.218*	(-1.764)	-0.026	(-0.244)
Family Size	-0.003	(-0.102)	-0.005	(-0.170)	0.019	(0.638)	0.011	(0.266)
% Children	0.395**	(2.488)	0.223	(1.159)	0.076	(0.360)	0.314	(1.205)
% Elderly	0.045	(0.285)	0.039	(0.212)	0.150	(0.744)	0.167	(0.674)
Mean HH Respondent's Attitude	-0.161***	(-3.069)	-0.299***	(-3.672)	-0.211	(-1.563)	0.022	(0.099)
Mean Dev. HH Respondent's Attitude			0.178*	(1.930)	-0.133	(-1.371)	-0.072	(-0.656)
S.d. HH Respondent's Attitude					0.181	(1.057)	0.003	(0.013)
Mean HH Respondent's Behaviour	-0.099	(-1.615)	-0.176*	(-1.883)	0.162	(0.924)	0.076	(0.312)
Mean Dev. HH Respondent's Behaviour			0.249**	(2.420)	-0.064	(-0.624)	-0.054	(-0.427)
S.d. HH Respondent's Behaviour					0.023	(0.124)	-0.119	(-0.497)
PGT	0.038	(0.355)	0.051	(0.394)	0.072	(0.507)	-0.169	(-1.177)
Rural	0.356***	(6.014)	0.141**	(2.235)	0.321***	(4.448)	0.220**	(2.553)
North & Northwest	0.419***	(3.769)	0.190	(1.417)	0.255	(1.624)	0.654***	(3.569)
Central	0.522***	(5.813)	0.550***	(4.711)	0.421***	(3.218)	0.862***	(5.673)
Volga	0.697***	(7.267)	0.451***	(3.770)	0.649***	(4.816)	0.880***	(5.882)
North Caucasus	0.361***	(3.586)	0.376***	(2.934)	0.276*	(1.884)	0.844***	(4.849)
Urals	0.588***	(6.389)	0.529***	(4.432)	0.479***	(3.533)	0.917***	(5.846)
East Siberia	0.380***	(3.685)	0.453***	(3.383)	0.398***	(2.688)	0.729***	(4.342)
West Siberia	0.455***	(4.413)	0.174	(1.362)	0.141	(0.973)	0.542***	(3.136)
Mean Life Satisfaction	-0.020	(-1.103)	0.047*	(1.704)	0.020	(0.526)	0.034	(0.683)
Mean Dev. Life Satisfaction			-0.033	(-0.980)	0.007	(0.194)	0.054	(1.464)
S.d. Life Satisfaction					-0.089	(-1.355)	-0.113	(-1.326)
Mean Employed	0.109***	(2.700)	0.231***	(4.157)	0.099	(1.416)	-0.007	(-0.083)
Mean Dev. Employed			0.140	(1.121)	0.009	(0.074)	0.068	(0.626)
S.d. Employed					0.005	(0.036)	0.078	(0.506)
Mean Health	-0.002	(-0.079)	0.006	(0.159)	0.092*	(1.859)	0.048	(0.791)
Mean Dev. Health			0.079	(1.317)	0.118**	(2.111)	0.039	(0.606)
S.d. Health					-0.142	(-1.597)	-0.115	(-1.026)
Married	-0.032	(-0.469)	0.150*	(1.779)	0.026	(0.268)	-0.068	(-0.568)
Divorced	-0.146*	(-1.774)	-0.049	(-0.472)	-0.320***	(-2.634)	-0.034	(-0.225)
Widow(er)	-0.055	(-0.571)	-0.108	(-0.970)	-0.064	(-0.496)	0.028	(0.186)
Age	0.032***	(4.887)	0.020***	(2.610)	0.060***	(6.402)	0.069***	(6.108)
Age^2	-0.000***	(-3.767)	-0.000**	(-2.064)	-0.001***	(-6.376)	-0.001***	(-6.453)
Education	0.008	(0.775)	0.004	(0.298)	-0.002	(-0.113)	-0.012	(-0.713)
Gender	-0.076***	(-3.246)	-0.140***	(-4.590)	-0.142***	(-3.852)	-0.138***	(-3.025)
Mean Respondent's Attitude	-0.091**	(-2.270)	-0.100	(-1.535)	0.019	(0.179)	-0.384***	(-2.612)
Mean Dev. Respondent's Attitude			-0.166**	(-2.234)	-0.099	(-1.338)	-0.205**	(-2.406)
S.d. Respondent's Attitude					-0.073	(-0.610)	0.029	(0.192)
Mean Respondent's Behaviour	0.031	(0.728)	0.110	(1.545)	0.029	(0.214)	0.039	(0.200)
Mean Dev. Respondent's Behaviour			-0.036	(-0.453)	0.133*	(1.710)	0.031	(0.324)
S.d. Respondent's Behaviour					-0.040	(-0.272)	0.185	(0.994)
Constant	-0.470	(-1.087)	0.196	(0.334)	-0.779	(-0.991)	-0.756	(-0.704)
Observations	7665		5755		4586		3715	
R^2	0.056		0.050		0.072		0.079	

* p<0.10, ** p<0.05, *** p<0.01.
t-statistics (robust to clustering within hpanelids) in parentheses.

Table A.4: Response Probits

	Round 10		Round 11		Round 12		Round 13	
Mean (Log) Household Income	-0.098	(-0.979)	-0.210*	(-1.732)	-0.247*	(-1.947)	-0.206	(-1.553)
Mean Dev. (Log) Household Income	0.059	(0.834)	-0.011	(-0.141)	0.111	(1.380)	0.089	(1.135)
S.d. (Log) Household Income	-0.051	(-0.389)	-0.259*	(-1.664)	0.092	(0.505)	-0.081	(-0.492)
Mean (Log) Household Consumption	-0.089	(-0.842)	0.098	(0.755)	-0.002	(-0.014)	0.165	(1.269)
Mean Dev. (Log) Household Consumption	-0.142	(-1.322)	-0.055	(-0.478)	-0.256**	(-2.384)	-0.025	(-0.235)
S.d. (Log) Household Consumption	0.376	(1.570)	-0.190	(-0.691)	-0.354	(-1.346)	-0.244	(-0.842)
Adult Female Head	0.164	(1.115)	-0.056	(-0.326)	0.267	(1.395)	-0.105	(-0.654)
Elderly Male Head	-0.159	(-0.588)	-0.612**	(-2.029)	-0.314	(-1.218)	0.103	(0.390)
Elderly Female Head	-0.073	(-0.245)	-0.396	(-1.252)	-0.123	(-0.437)	0.252	(0.905)
Household Head Changed	-0.082	(-0.617)	-0.345**	(-2.143)	0.159	(0.937)	-0.195	(-1.181)
Family Size	0.135***	(3.036)	-0.042	(-0.870)	0.054	(1.114)	0.012	(0.252)
% Children	0.904***	(2.921)	0.862**	(2.446)	-0.137	(-0.383)	0.016	(0.046)
% Elderly	0.574**	(1.999)	0.694**	(2.121)	0.350	(1.148)	-0.087	(-0.293)
Mean HH Respondent's Attitude	0.078	(0.289)	0.002	(0.005)	-0.140	(-0.340)	0.098	(0.220)
Mean Dev. HH Respondent's Attitude	-0.099	(-0.756)	-0.139	(-0.872)	-0.030	(-0.193)	-0.038	(-0.230)
S.d. HH Respondent's Attitude	-0.394	(-1.428)	-0.692*	(-1.953)	-0.223	(-0.584)	-0.615	(-1.501)
Mean HH Respondent's Behaviour	0.137	(0.433)	0.456	(0.895)	-0.911**	(-2.018)	-0.292	(-0.556)
Mean Dev. HH Respondent's Behaviour	-0.251*	(-1.674)	-0.038	(-0.225)	-0.044	(-0.262)	-0.055	(-0.297)
S.d. HH Respondent's Behaviour	0.210	(0.720)	0.713*	(1.724)	-0.550	(-1.561)	-0.187	(-0.456)
PGT	0.150	(0.873)	0.039	(0.200)	0.240	(1.298)	0.075	(0.398)
Rural	0.080	(0.827)	0.176	(1.550)	0.273**	(2.450)	0.212**	(1.991)
North & Northwest	0.081	(0.354)	0.623**	(2.139)	-0.264	(-0.843)	-0.129	(-0.479)
Central	0.357*	(1.729)	0.509**	(2.147)	-0.048	(-0.176)	0.370	(1.552)
Volga	0.393*	(1.877)	0.700***	(2.902)	0.039	(0.141)	0.548**	(2.183)
North Caucasus	0.189	(0.843)	0.371	(1.411)	-0.171	(-0.570)	0.240	(0.904)
Urals	0.522**	(2.363)	0.315	(1.354)	0.165	(0.590)	0.474*	(1.926)
East Siberia	0.097	(0.439)	0.220	(0.863)	-0.248	(-0.841)	0.326	(1.200)
West Siberia	0.179	(0.738)	0.752***	(2.857)	0.133	(0.427)	0.385	(1.437)
Mean Life Satisfaction	0.020	(0.327)	-0.037	(-0.537)	-0.003	(-0.037)	0.019	(0.260)
Mean Dev. Life Satisfaction	-0.034	(-0.816)	-0.000	(-0.008)	0.111**	(2.259)	0.046	(0.987)
S.d. Life Satisfaction	0.112	(0.971)	0.189	(1.289)	-0.214	(-1.366)	-0.002	(-0.014)
Mean Employed	0.155	(1.367)	0.211*	(1.818)	0.028	(0.211)	-0.097	(-0.694)
Mean Dev. Employed	0.259**	(2.062)	-0.122	(-0.814)	0.212	(1.396)	-0.027	(-0.181)
S.d. Employed	0.101	(0.546)	0.063	(0.322)	0.291	(1.266)	0.149	(0.750)
Mean Health	-0.003	(-0.037)	0.198**	(2.273)	0.221**	(2.265)	0.130	(1.376)
Mean Dev. Health	-0.016	(-0.229)	0.101	(1.180)	0.050	(0.639)	0.121	(1.501)
S.d. Health	-0.334**	(-2.261)	-0.223	(-1.346)	-0.108	(-0.613)	-0.110	(-0.577)
Married	0.162	(1.094)	-0.158	(-0.826)	0.369**	(1.969)	-0.023	(-0.124)
Divorced	0.157	(0.880)	-0.300	(-1.300)	0.014	(0.066)	-0.151	(-0.747)
Widow(er)	0.061	(0.343)	-0.329	(-1.433)	0.178	(0.819)	-0.066	(-0.334)
Age	0.035**	(2.262)	0.040**	(2.391)	0.030	(1.602)	0.062***	(3.266)
Age^2	-0.000**	(-2.183)	-0.000**	(-2.010)	-0.000*	(-1.924)	-0.001***	(-3.465)
Education	-0.005	(-0.236)	0.048**	(2.110)	-0.004	(-0.158)	-0.005	(-0.207)
Gender	-0.060	(-1.116)	-0.241***	(-3.920)	-0.237***	(-3.495)	-0.221***	(-3.052)
Mean Respondent's Attitude	-0.236	(-1.141)	0.053	(0.175)	0.056	(0.201)	-0.303	(-0.871)
Mean Dev. Respondent's Attitude	-0.151	(-1.460)	-0.063	(-0.475)	-0.082	(-0.597)	-0.161	(-1.204)
S.d. Respondent's Attitude	-0.039	(-0.185)	0.073	(0.263)	-0.069	(-0.266)	0.107	(0.332)
Mean Respondent's Behaviour	-0.343	(-1.163)	-0.209	(-0.564)	0.628*	(1.806)	-0.688	(-1.422)
Mean Dev. Respondent's Behaviour	0.078	(0.702)	-0.180	(-1.454)	0.078	(0.627)	-0.114	(-0.815)
S.d. Respondent's Behaviour	-0.240	(-0.960)	-0.104	(-0.354)	0.250	(0.875)	-0.201	(-0.591)
Constant	1.682	(1.422)	-0.716	(-0.460)	3.084*	(1.955)	2.897	(1.575)
Observations	3116		2811		2591		2350	
R^2	0.070		0.081		0.085		0.067	

* $p<0.10$, ** $p<0.05$, *** $p<0.01$.
t-statistics (robust to clustering within hpanelids) in parentheses.

A.4 BHPS Variable Construction

A.4.1 Consumption

Following Deaton and Zaidi (2002), we calculate consumption as the sum of four consumption sub-aggregates: food expenditure, non-food expenditure, the user cost of durables, and the rental value of housing; defining each sub-aggregate over the six month period prior to the conduct of the survey. The following subsections set out the construction of each of the sub-aggregates in turn.

Food Consumption

The BHPS contains a question asking respondents to estimate their approximate weekly food and grocery expenditure, by indicating which of 12 expenditure categories they fall into (£10 bands up to £60, and £20 bands thereafter, with a final band of £160 and upwards). We convert responses in each bounded category to their midpoint, and recode the last as £170, before multiplying by 26 to generate average half-yearly values.

Non-Food Consumption

The BHPS contains three main sources of non-food consumption data: monthly expenditure on eating out, monthly expenditure on leisure, and annual expenditure on heating. The first two are reported in the same band structure as weekly food expenditure, and we convert their values accordingly. We then multiply by 6, or divide by 2, as appropriate, to obtain half-yearly averages.

Durable Consumption

Following Deaton and Zaidi (2002) we can calculate the user cost, $D_{g,t}$ of durable good g at time t as:

$$D_{g,t} = V_{g,a(t)} \left(1 - \frac{(1+\pi_t)(1-\delta_g)}{(1+r_t)} \right) \tag{A.13}$$

where $V_{g,a(t)}$ is the market value of good g of age a six months prior to the survey date, t, δ_g is its rate of depreciation, r_t is the nominal interest rate, and π_t is the rate of inflation over the period ending at t.[17] As nominal

[17] This formula differs from that set out in Deaton and Zaidi (2002) in two ways. The first, and most obvious, is that it employs the exact formula based on rearrangement of their Equation (2.11), rather than the common approximation set out in their Equation (2.12):

$$D_{g,t} = V_{g,a(t)}(r_t - \pi_t + \delta_g) \tag{A.14}$$

There is in general, little reason to use the approximate version, but it would be particularly inappropriate here when we have reasonably high depreciation (and in some cases interest) rates.

The second, more subtle difference lies in our variable definitions. Deaton and Zaidi (2002) calculate the user cost at time t as the cost of holding the good from time t to time $t+1$, whereas (to remain consistent with the other consumption sub-aggregates) we calculate it as the cost of holding the good from six months prior to the survey until the survey date (i.e. from $t - \frac{1}{2}$ to t). For this purpose, we require $V_{g,a(t)}$ to be the value at $t - \frac{1}{2}$, and r_t, π_t, and δ_g to reflect rates of change between $t - \frac{1}{2}$ and t. (Although, in practice, we take a three year moving average of interest rates, to allow for the fact that it may be difficult for households to quickly alter their durable holdings to take advantage of interest rate changes.)

interest rates, we use a 3-year moving average (calculated over the 2 years prior, and 1 year subsequent to the survey date) of the "Monthly Average of the Official Bank Rate" from the Bank of England (series IUMABEDR), and use monthly CPI data from the Office of National Statistics as the inflation rate. Estimates of depreciation rates for consumer durables are difficult to come by for the UK, so we follow a number of other studies in using values for the US calculated by Jorgenson and Stiroh (2000), and set out in Table A.5.[18] Because we define consumption over a six-month time horizon, r_t, π_t, and δ_g are all semi-annual rates, calculated from the annual rates r_t^*, π_t^*, and δ_g^* as, e.g.:

$$r_t = (1 + r_t^*)^{\frac{1}{2}} - 1 \tag{A.15}$$

Table A.5: Consumer Durable Depreciation Rates

Durable Category	**Goods Included**	δ_g^*
Automobiles	Vehicles	25.5%
Kitchen Appliances	Freezers, Washing Machines, Dryers, Dishwashers, Microwaves	15.0%
Computers and Software	Computers	31.5%
Video, Audio	TVs, VCRs, CD Players, Satellite Dishes, Cable TV	18.33%
Other Durable Goods	Telephones	16.5%

The BHPS contains a variety of information on the durable goods possessed by households, including: vehicles, TVs, VCRs, freezers, washing ma-

[18] See e.g. Venturini (2006), Timmer and van Ark (2005), Draca, Sadun, and van Reenen (2006), and Keeney (2007).

chines, dryers, dishwashers, microwaves, computers, CD players, satellite dishes, cable TV, and telephones. Unfortunately, we have explicit data on the value of these goods only in somewhat limited circumstances: for vehicles only in waves 3-5, 7, and 10-15; and for the remainder only for waves 7-15, and then only if they were purchased in the previous year. As a result, we need to substitute estimates of $V_{g,a(t)}$. In doing so, our main goals are to capture as accurately as possible: (a) the year-on-year changes in durable consumption, the primary drivers of which are (i) depreciation, and (ii) new durable purchases; and (b) the percentage of overall consumption that consists of durables.[19] In constructing these measures, a number of somewhat arbitrary (and in some cases obviously incorrect) choices need to be made. Nonetheless, most of these should not significantly affect the key drivers of the measures (depreciation and new purchases) and the undoubtedly imperfect measures that result are likely to be far more useful for our purposes than the alternative of excluding durable consumption entirely, which is often adopted in other contexts.

[19]The latter would not matter if we thought that SWB were linear in consumption, but with a log functional form, misestimating this percentage will result in misestimates of variation in the logs of the consumption variable, which will bias our estimates. See Footnote 12

General Approach

In general (i.e. for all goods except vehicles), where we lack the relevant data, we substitute an estimate of $V_{g,a(t)}$, assuming a constant geometric depreciation rate:[20]

$$V_{g,a(t)} = V_{g,0}(1 - \delta_g^*)^a \tag{A.16}$$

We assume that goods purchased in a given period, t_0, are purchased new, at a point half way through the period. Consequently, where we observe the purchase of a good within the sample period, we know that its age is $a(t) = t - t_0$. If the purchase occurs in waves 7-15 we also observe $V_{g,0}$ directly. In other cases, one or both of these quantities must be estimated.

Estimating a

If we observe a household with access to a particular durable good f (of type g), but do not observe its acquisition, then we estimate the lifetime of the good, L_{if}, as the greater of two quantities: (a) the expected age at which the household would typically replace the good, conditional on its income and region; and (b) the length of time, t_{if}, that the household is observed as having access to the good.[21] (Assuming that goods have zero

[20] In contrast to (A.13), the depreciation rate here is annual, and the age of the good is expressed in years.

[21] Where use data on access is missing for some waves, we perform imputation based on data for the last and next waves for which it is available, according to the following rules.

value once replaced by a new good of the same type, and that new goods are acquired half way through the relevant period, this quantity is equal to the number of full periods that the household has the use of the good prior to replacement, plus an extra half period if the good is replaced during the sample.) Formally:[22]

$$\hat{L}_{if} = \max\left\{(\hat{\beta}_g \ln y_i + \hat{\gamma}_g z_i), t_{if}\right\} \tag{A.17}$$

where y_i is mean household real income over the sample period, z_i is the household's modal region of residence, and $\hat{\beta}_g$ and $\hat{\gamma}_g$ are the estimated coefficients from the the following regression (run on the sample of households

(1) If the household reports having access to a good in both the last and next waves, assume it has access to the good in the intervening waves.
(2) If the household reports not having access to a good in both the last and next waves, assume it does not have access to the good in the intervening waves.
(3) If the household reports access to the good in the last wave, but not in the next, assume it has access to the good in the former half of the intervening waves (rounding up).
(4) If the household reports not having access to a good in the last wave, but having access to it in the next:
    (a) assume that it does not have access to the good in the intervening waves if we know the good was purchased in the next wave for which data is available; but
    (b) otherwise assume that it had access to the good in the latter half of the intervening waves (rounding up).
(5) If either there is no prior data, or no subsequent data on the use of a particular type of good, make no imputation.

[22]Ideally, we would like an estimate of $\mathrm{E}[L_{if}|y_i, z_i, t_{if}]$, which is not quite what Equation (A.17) gives us (for example, $\mathrm{E}[L_{if}|y_i, z_i, t_{if}]$ will almost certainly be greater than $\hat{L}_{if}$ in the event that $\hat{L}_{if} = t_{if}$). Nonetheless, the formula set out in (A.17) seems a reasonable approximation for our purposes.

who have access to a given durable good type prior to their first observed purchase of such a good):

$$L_{ig} = \beta_g \ln y_i + \gamma_g z_i + \varepsilon_i \tag{A.18}$$

where L_{ig} is the implied average lifetime of goods of type g for possessed by household i, defined as:

$$L_{ig} = \begin{cases} \frac{w_{ig}}{n_{ig}} & \text{if } n_{ig} > 0 \\ w_{ig} + 5 & \text{if } n_{ig} = 0 \end{cases} \tag{A.19}$$

where w_{ig} is the number of waves in which household i is observed in possession of any good of type g, and n_{ig} is the number of acquisitions of new durable goods of type g by the household over the course of the sample (which together imply an expected lifetime of $\frac{w_{ig}}{n_{ig}}$ if $n_{ig} > 0$). If we do not observe any acquisitions by the household within the sample period (i.e. $n_{ig} = 0$), we assume, somewhat arbitrarily, that the good has a lifetime of $w_{ig} + 5$ years.

From this, we can then estimate the age of the good by assuming that it is $\hat{L}_{if}$ years old when replaced (or retired). This is straightforward if we observe the good being replaced (or retired) within the sample period — in which case it will be $\hat{L}_{if} - 1$ (or $\hat{L}_{if} - \frac{1}{2}$) years old at the beginning of the period prior to replacement (or retirement). If we do not, we instead take the good's age at the beginning of the final period observed to be:

$$a(T) = \max\left\{\frac{1}{2}(\hat{L}_{if} + w_{ig}), w_{ig}\right\} - \frac{1}{2} \tag{A.20}$$

(where the first term in curly brackets is the good's expected age at the end of period t if $\hat{L}_{if} \geq w_{ig}$ and we place a uniform prior over all possible replacement dates that allow the good to be $\hat{L}_{if}$ years old when replaced).

Estimating $V_{g,0}$

To impute values for $V_{g,0}$ we estimate the following equation on the observed sample:

$$V_{g,0,i,t} = \beta_g \ln y_{i,t} + \gamma_g z_i + \varepsilon_{g,i,t} \tag{A.21}$$

where $y_{i,t}$ is household income, and z_i is a vector of region dummies. We then impute the missing $V_{g,0}$s as:

$$\hat{V}_{g,0,i,t} = \hat{\beta}_g \ln y_{i,t} + \hat{\gamma}_g z_{i,t} \tag{A.22}$$

(substituting the earliest observed $y_{i,t}$ and $z_{i,t}$ if the good was purchased prior to the beginning of the sample).

Vehicles

Due to greater data availability, we follow a slightly different procedure to generate the user cost of vehicles. Having data on the market value of vehicles for most rounds allows us to simply interpolate the missing values in a linear fashion. (This is strictly inconsistent with the assumption of a geometric depreciation rate, but seems to perform adequately for our purposes.) Because $V_{g,a(t)}$ is the good's value six months prior to the survey date, we then project this reported (or imputed) value at the survey date, $V_{g,a(t+\frac{1}{2})}$, backwards, as:

$$V_{g,a(t)} = \frac{V_{g,a(t+\frac{1}{2})}}{(1-\delta_g)} \tag{A.23}$$

before applying the formula in Equation (A.13).

Housing Consumption

The gold standard for measures of housing consumption is typically taken to be the market rent that is (or would be) paid for a household's dwelling. Depending on the type of housing, we use three alternative measures of this notional private rental.

(1) For the approximately 7% of households renting on the private market we use the rent actually paid.

(2) For the approximately 21% of households renting from local councils,

housing associations, and other providers of social housing, rents paid are likely to be below market rents. On the other hand, the alternative of using estimates based on hedonic housing regressions is likely to overstate housing consumption, because council and other social housing is likely to differ from private rentals in neigbourhood unobservables that would reduce its market price. For simplicity, we use the reported rents.

(3) For owner-occupied housing (including those who own their homes outright, and those who have outstanding mortgages), we base the estimate of imputed rent on the value of the home reported in the BHPS, multiplied by the estimated annual rental yield for the UK (and then halved, to give a half-yearly figure). The rental yield values are calculated as set out in Section A.4.1 below, and displayed in Table A.6.[23]

Calculating Rental Yields for the UK

The method we adopt for calculating rental yields for the UK is based on that set out in Carver (2006), with a few modifications. Yields for 2001-2006 are drawn directly from the Investment Property Databank (IPD) UK

[23]We attempted to calculate rental yields at the regional, rather than the national level, but were unable to generate reasonable values, as only some elements of the data used to construct the series were available at the appropriate level of disaggregation.

Table A.6: UK Rental Yields

Year	Rental Yield
1987	7.3%
1988	6.3%
1989	5.7%
1990	6.4%
1991	7.5%
1992	8.5%
1993	9.2%
1994	9.5%
1995	9.8%
1996	9.8%
1997	9.1%
1998	8.3%
1999	7.7%
2000	7.1%
2001	7.3%
2002	6.9%
2003	6.1%
2004	5.9%
2005	5.6%
2006	5.1%

Residential Investment Index.[24] For years prior to 2001, the yield is given as an estimated UK average mix-adjusted rent, divided by an estimated UK average mix-adjusted house price:

$$\text{UK Average Rental Yield}_t = \frac{\text{UK Average Mix-Adjusted Rental}_t}{\text{UK Average Mix-Adjusted House Price}_t} \tag{A.24}$$

From 1993 onwards, the Average UK Mix-Adjusted House Prices are drawn directly from the Department of Communities and Local Government (DCLG) Table 507. We then use the UK Mix-Adjusted Price Index (Table 593) to project these values back for years prior to this.

[24]Specifically, the publicly available IPD reports for 2004, 2005, and 2006. Because the series have been updated over time, for each year, we use the most recent available estimate if multiple estimates are given.

The method used to obtain the Mix-Adjusted Rentals is somewhat more involved. To begin with, we calculate the implied rental in 2001 as the DCLG Mix-Adjusted House Price, multiplied by the IPD Rental Yield. We then project this back, using an index of UK rental prices, constructed from rent data drawn from three sources: (a) weekly market rentals in England for 1990 onwards from the York Review of Housing (Table 72); (b) monthly market rentals in England for 1995 onwards from the DCLG (Table 734, based on data from the Survey of English Housing); and (c) The rent component (DOBPAU series) of the UK Retail Price Index (RPI) for 1987 onwards. These are used to construct a Composite Rent Index based on a subjectively-weighted average of the yearly percentage increases derived from each series. The weights for this calculation, closely based on those adopted by Carver (2006) are set out in Table A.7.[25] They assume that the DCLG is the best rent index, when available (as it best approximates a replacement price for owner occupied housing), followed by the York Index and then the RPI.

Table A.7: CRI Weights

Year	**DCLG**	**York**	**RPI**
1991-1995		0.7	0.3
1996-2001	0.6	0.25	0.15

Once we have the annual figures set out in Table A.6, we calculate the

[25]The divergences from Carver (2006) are minimal, and essentially due to slightly different data availability.

relevant annual yield for each household as a weighted average depending on the the number of months of the last six that fall into a given calendar year. So, for example, if a household is interviewed in the latter half of February 1999, its rental yield will be $\frac{2}{6}$ times the yield for 1999 (for January and February), and $\frac{4}{6}$ times the yield for 1998 (for September, October, November and December).

A.4.2 Health Assessment

As noted in the text, the subjective health assessment variable we use as a control is not available in Wave 9. We impute values for the missing round based on a linear fixed effects regression, using all data for Waves 6-15 for which life satisfaction data are available. The variables included in the RHS of the regression are the standard control variables used throughout the British analysis, coupled with the an extra set of dummies for a range of reported health problems (which are available in all rounds).

Specifically, the former are comprised of: log household income; log household consumption; age; age squared; marital status dummies; an employment dummy; household size; and a dummy indicating whether the observation occurred in any of waves 8, 9, or 10 (this allows for the possibility of common shocks, assuming for predictive purposes that the unobserved shock in wave

9 is the average of the observed shocks in waves 8 and 10).

The health dummies indicate whether the respondent experiences problems with: arms, legs, hands etc.; sight; hearing; skin conditions/allergies; chest/breathing; heart/blood pressure; stomach or digestion; diabetes; anxiety, depression etc.; alcohol or drugs; epilepsy; migraines; or other problems (with no reported problems as the omitted category).

A.5 Predicted Income Estimations

See Table A.8.

A.6 Reference Income Regressions

See Tables A.9 and A.10

A.7 The Penalized Gaussian Mixture Model

As noted above, the PGM proceeds by assuming that the true distribution of the random effects can be approximated by a mixture of (multivariate) basis Gaussian densities with fixed means. Specifically, it assumes that each of the K random effects to be estimated extends over a (standardized) interval

Table A.8: Predicted Income Equations (Rounds 5-9)

	Round 5	Round 6	Round 7	Round 8	Round 9
Log Individual Income:					
Education	-0.004	0.016**	0.017***	0.024***	0.027***
	(0.005)	(0.006)	(0.004)	(0.004)	(0.003)
Female	-0.368***	-0.212***	-0.391***	-0.255***	-0.252***
	(0.024)	(0.055)	(0.029)	(0.026)	(0.025)
Experience	-0.000	-0.003***	-0.001***	-0.000	-0.001*
	(0.000)	(0.001)	(0.000)	(0.000)	(0.000)
Experience^2	0.000	0.000	0.000*	-0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Army	1.308***	1.360***	1.498***	1.151***	1.102***
	(0.111)	(0.148)	(0.165)	(0.091)	(0.066)
Legislators, Senior Managers	0.918***	0.967***	0.786***	0.703***	0.717***
	(0.043)	(0.091)	(0.051)	(0.051)	(0.050)
Professionals	0.771***	1.049***	0.738***	0.579***	0.743***
	(0.044)	(0.081)	(0.050)	(0.048)	(0.047)
Technicians	0.568***	1.062***	0.681***	0.486***	0.519***
	(0.060)	(0.111)	(0.065)	(0.068)	(0.068)
Clerks	0.859***	1.236***	0.644***	0.609***	0.569***
	(0.055)	(0.094)	(0.058)	(0.057)	(0.051)
Service Workers	0.707***	-0.527	-0.155	0.325	0.420**
	(0.198)	(0.436)	(0.206)	(0.237)	(0.180)
Skilled Agr. & Fish.	0.627***	0.994***	0.599***	0.500***	0.739***
	(0.043)	(0.082)	(0.050)	(0.051)	(0.047)
Craft & Related	0.665***	0.981***	0.652***	0.587***	0.718***
	(0.044)	(0.082)	(0.051)	(0.048)	(0.045)
Operators & Assemblers	0.237***	0.587***	0.320***	0.195***	0.126***
	(0.047)	(0.083)	(0.051)	(0.050)	(0.048)
Elementary Unskilled	1.248***	1.171***	0.734***	0.779***	1.225***
	(0.155)	(0.218)	(0.135)	(0.138)	(0.157)
North & Northwest	0.026	-0.071	-0.158***	-0.128**	-0.122*
	(0.051)	(0.101)	(0.058)	(0.058)	(0.064)
Central & Black Earth	-0.378***	-0.402***	-0.392***	-0.337***	-0.297***
	(0.039)	(0.078)	(0.045)	(0.046)	(0.055)
Volga	-0.560***	-0.784***	-0.640***	-0.448***	-0.461***
	(0.040)	(0.079)	(0.047)	(0.048)	(0.055)
North Caucasus	-0.482***	-0.392***	-0.451***	-0.411***	-0.403***
	(0.043)	(0.086)	(0.052)	(0.051)	(0.058)
Urals	-0.307***	-0.470***	-0.398***	-0.322***	-0.326***
	(0.041)	(0.081)	(0.048)	(0.049)	(0.057)
Western Siberia	-0.087*	-0.123	-0.050	-0.099*	-0.253***
	(0.046)	(0.092)	(0.055)	(0.055)	(0.060)
Eastern Siberia & Far East	-0.037	-0.289***	-0.288***	-0.240***	-0.256***
	(0.047)	(0.090)	(0.055)	(0.054)	(0.060)
Constant	8.612***	8.373***	8.653***	7.940***	7.882***
	(0.054)	(0.107)	(0.064)	(0.060)	(0.065)
Selection Equation:					
Female	0.092***	-0.031	0.080***	0.044	-0.006
	(0.031)	(0.027)	(0.028)	(0.029)	(0.031)
Age	0.124***	0.100***	0.109***	0.103***	0.123***
	(0.003)	(0.002)	(0.003)	(0.003)	(0.003)
Age^2	-0.001***	-0.001***	-0.001***	-0.001***	-0.001***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
# Childen <7	-0.030	-0.049***	-0.028	-0.102***	0.036
	(0.024)	(0.017)	(0.023)	(0.024)	(0.026)
Constant	-2.494***	-1.778***	-2.385***	-2.235***	-2.434***
	(0.061)	(0.049)	(0.061)	(0.060)	(0.062)
ρ	-0.527***	-0.946***	-0.641***	-0.603***	-0.552***
	(0.027)	(0.004)	(0.030)	(0.027)	(0.023)
σ	0.888***	2.249***	0.955***	0.931***	0.945***
	(0.009)	(0.022)	(0.016)	(0.012)	(0.010)
Observations	10703	10183	10000	10153	10452

* p<0.10, ** p<0.05, *** p<0.01. Standard Errors in parentheses.

Table A.8: Predicted Income Equations (Rounds 10-13)

	Round			
	10	11	12	13
Log Individual Income:				
Education	-0.008	-0.005	-0.005	-0.002
	(0.006)	(0.006)	(0.006)	(0.006)
Female	-0.387***	-0.366***	-0.357***	-0.411***
	(0.024)	(0.023)	(0.022)	(0.022)
Experience	-0.001**	-0.000	-0.000	-0.001***
	(0.000)	(0.000)	(0.000)	(0.000)
Experience^2	0.000	0.000	0.000	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Army	1.211***	1.335***	1.492***	1.470***
	(0.058)	(0.056)	(0.063)	(0.060)
Legislators, Senior Managers	0.891***	0.988***	1.061***	1.171***
	(0.044)	(0.042)	(0.040)	(0.040)
Professionals	0.818***	0.940***	0.951***	1.027***
	(0.044)	(0.041)	(0.040)	(0.039)
Technicians	0.763***	0.863***	0.851***	1.016***
	(0.066)	(0.059)	(0.055)	(0.058)
Clerks	0.643***	0.805***	0.763***	0.853***
	(0.047)	(0.046)	(0.042)	(0.042)
Service Workers	0.280	0.364*	0.642***	0.827***
	(0.196)	(0.194)	(0.180)	(0.189)
Skilled Agr. & Fish.	0.870***	0.828***	0.900***	0.968***
	(0.046)	(0.043)	(0.042)	(0.041)
Craft & Related	0.766***	0.815***	0.873***	0.911***
	(0.043)	(0.041)	(0.040)	(0.039)
Operators & Assemblers	0.367***	0.406***	0.482***	0.571***
	(0.045)	(0.043)	(0.040)	(0.041)
Elementary Unskilled	1.130***	1.243***	1.317***	1.134***
	(0.156)	(0.159)	(0.153)	(0.145)
North & Northwest	-0.020	-0.175***	-0.165***	-0.106**
	(0.050)	(0.047)	(0.049)	(0.048)
Central & Black Earth	-0.386***	-0.432***	-0.495***	-0.443***
	(0.038)	(0.035)	(0.035)	(0.035)
Volga	-0.546***	-0.625***	-0.670***	-0.648***
	(0.039)	(0.035)	(0.035)	(0.035)
North Caucasus	-0.495***	-0.556***	-0.609***	-0.526***
	(0.041)	(0.038)	(0.038)	(0.038)
Urals	-0.432***	-0.496***	-0.633***	-0.545***
	(0.040)	(0.037)	(0.037)	(0.037)
Western Siberia	-0.363***	-0.525***	-0.614***	-0.585***
	(0.047)	(0.044)	(0.042)	(0.041)
Eastern Siberia & Far East	-0.354***	-0.413***	-0.535***	-0.494***
	(0.045)	(0.042)	(0.041)	(0.041)
Constant	8.553***	8.641***	8.752***	8.746***
	(0.059)	(0.056)	(0.056)	(0.057)
Selection Equation:				
Female	0.085***	0.024	0.061**	0.068**
	(0.030)	(0.030)	(0.030)	(0.031)
Age	0.124***	0.122***	0.127***	0.138***
	(0.003)	(0.003)	(0.003)	(0.003)
Age^2	-0.001***	-0.001***	-0.001***	-0.001***
	(0.000)	(0.000)	(0.000)	(0.000)
# Childen <7	-0.014	0.035	0.034	0.031
	(0.026)	(0.026)	(0.027)	(0.027)
Constant	-2.384***	-2.299***	-2.350***	-2.486***
	(0.060)	(0.058)	(0.058)	(0.061)
ρ	-0.562***	-0.581***	-0.612***	-0.646***
	(0.023)	(0.022)	(0.023)	(0.022)
σ	0.951***	0.934***	0.925***	0.935***
	(0.009)	(0.009)	(0.009)	(0.009)
Observations	11360	11806	11979	11886

* p<0.10, ** p<0.05, *** p<0.01. Standard Errors in parentheses.

Table A.9: Predicted Income Equations (Females: Waves 7-10)

	Wave							
	7		8		9		10	
Main Equation:								
Age	-0.060***	(0.010)	-0.039***	(0.009)	-0.019**	(0.007)	-0.080***	(0.008)
$Age^2/10$	0.008***	(0.001)	0.006***	(0.001)	0.003***	(0.001)	0.011***	(0.001)
Incentive Pay	0.072*	(0.042)	0.109***	(0.039)	0.099***	(0.033)	0.055	(0.036)
Full Time	0.544***	(0.041)	0.544***	(0.037)	0.490***	(0.031)	0.465***	(0.032)
Union Member	0.101**	(0.047)	0.086**	(0.041)	0.059*	(0.036)	0.090**	(0.037)
Pension	0.237***	(0.043)	0.200***	(0.039)	0.200***	(0.034)	0.201***	(0.036)
Tenure	-0.009	(0.006)	-0.008	(0.006)	0.001	(0.005)	0.009	(0.006)
$Tenure^2$	0.000***	(0.000)	0.001**	(0.000)	0.000	(0.000)	-0.000	(0.000)
Constant	9.561***	(0.457)	9.166***	(0.391)	8.993***	(0.361)	10.210***	(0.443)
Selection Equation:								
Log Non-labour Income	-0.084***	(0.011)	-0.136***	(0.013)	-0.116***	(0.012)	-0.084***	(0.008)
Children Aged 0-2	-0.274***	(0.093)	-0.048	(0.095)	0.051	(0.094)	0.138**	(0.070)
Children Aged 3-4	-0.179**	(0.091)	-0.055	(0.094)	-0.060	(0.096)	0.129*	(0.071)
Children Aged 5-11	-0.053	(0.083)	0.126	(0.082)	0.162**	(0.082)	0.176***	(0.060)
Children Aged 12-15	0.053	(0.084)	0.260***	(0.083)	0.285***	(0.085)	0.234***	(0.062)
Total Children	-0.054	(0.077)	-0.271***	(0.077)	-0.345***	(0.077)	-0.287***	(0.057)
In Relationship	0.132***	(0.037)	0.190***	(0.041)	0.146***	(0.041)	0.097***	(0.031)
Age	0.142***	(0.008)	0.158***	(0.008)	0.155***	(0.008)	0.159***	(0.007)
$Age^2/10$	-0.018***	(0.001)	-0.020***	(0.001)	-0.020***	(0.001)	-0.020***	(0.001)
Constant	-1.907***	(0.213)	-1.863***	(0.218)	-1.796***	(0.225)	-2.324***	(0.192)
ρ	-1.883***	(0.065)	-1.389***	(0.056)	-1.344***	(0.058)	-1.634***	(0.046)
σ	0.180***	(0.017)	-0.012	(0.017)	-0.172***	(0.017)	0.141***	(0.014)
Observations	4589		5364		5319		7590	

* p<0.10, ** p<0.05, *** p<0.01. Standard Errors in parentheses. Included, but not reported are dummies for Health (5) Region (18), Occupation (8), Industry (9), Education (6), When Work (10), Organisation Type (8), Term of Contract (3), Mangerial Responsibility (3) in the main equation, and dummies for Health (5) Region (18), and Education (6) in the selection equation.

Table A.9: Predicted Income Equations (Females: Waves 12-15)

	Wave							
	12		13		14		15	
Main Equation:								
Age	0.054***	(0.011)	-0.058***	(0.009)	-0.077***	(0.011)	-0.092***	(0.012)
$Age^2/10$	-0.006***	(0.001)	0.008***	(0.001)	0.010***	(0.001)	0.012***	(0.001)
Incentive Pay	0.083*	(0.044)	0.081**	(0.041)	0.105**	(0.050)	0.089*	(0.052)
Full Time	0.550***	(0.041)	0.500***	(0.034)	0.342***	(0.042)	0.497***	(0.043)
Union Member	0.106**	(0.044)	0.098**	(0.041)	0.091*	(0.050)	0.091*	(0.052)
Pension	0.186***	(0.044)	0.131***	(0.041)	0.158***	(0.050)	0.117**	(0.052)
Tenure	0.004	(0.007)	0.021***	(0.006)	0.021***	(0.008)	0.020**	(0.008)
$Tenure^2$	0.000	(0.000)	-0.001***	(0.000)	-0.000	(0.000)	-0.000	(0.000)
Constant	6.712***	(0.485)	9.913***	(0.408)	10.907***	(0.621)	11.719***	(0.604)
Selection Equation:								
Log Non-labour Income	-0.178***	(0.010)	-0.093***	(0.008)	-0.083***	(0.009)	-0.094***	(0.009)
Children Aged 0-2	-0.094	(0.116)	-0.215**	(0.094)	-0.104	(0.089)	0.160*	(0.090)
Children Aged 3-4	-0.037	(0.116)	-0.189**	(0.091)	-0.201**	(0.090)	0.039	(0.092)
Children Aged 5-11	0.140	(0.098)	-0.014	(0.078)	0.070	(0.075)	0.213***	(0.076)
Children Aged 12-15	0.289***	(0.097)	0.030	(0.078)	0.070	(0.075)	0.256***	(0.078)
Total Children	-0.351***	(0.090)	-0.102	(0.073)	-0.171**	(0.072)	-0.279***	(0.072)
In Relationship	0.132***	(0.049)	0.034	(0.037)	0.067*	(0.037)	0.086**	(0.037)
Age	0.187***	(0.009)	0.180***	(0.008)	0.168***	(0.008)	0.172***	(0.008)
$Age^2/10$	-0.024***	(0.001)	-0.023***	(0.001)	-0.021***	(0.001)	-0.021***	(0.001)
Constant	-3.576***	(0.205)	-4.001***	(0.184)	-3.715***	(0.189)	-3.916***	(0.190)
ρ	-0.057	(0.074)	-1.476***	(0.050)	-1.637***	(0.061)	-1.628***	(0.066)
σ	-0.082***	(0.012)	0.104***	(0.015)	0.277***	(0.016)	0.319***	(0.017)
Observations	6950		6793		6218		6241	

* $p<0.10$, ** $p<0.05$, *** $p<0.01$. Standard Errors in parentheses. Included, but not reported are dummies for Health (5) Region (18), Occupation (8), Industry (9), Education (6), When Work (10), Organisation Type (8), Term of Contract (3), Mangerial Responsibility (3) in the main equation, and dummies for Health (5) Region (18), and Education (6) in the selection equation.

Table A.10: Predicted Income Equations (Males: Waves 7-10)

	Wave							
	7		8		9		10	
Main Equation:								
Age	-0.043***	(0.010)	-0.002	(0.009)	0.037***	(0.008)	-0.030***	(0.010)
$Age^2/10$	0.006***	(0.001)	0.001	(0.001)	-0.004***	(0.001)	0.005***	(0.001)
Incentive Pay	-0.007	(0.045)	0.033	(0.037)	0.032	(0.029)	0.060	(0.040)
Full Time	0.409***	(0.075)	0.679***	(0.068)	0.430***	(0.055)	0.397***	(0.067)
Union Member	0.015	(0.054)	0.032	(0.046)	0.029	(0.036)	0.021	(0.048)
Pension	0.202***	(0.051)	0.142***	(0.042)	0.091***	(0.034)	0.115**	(0.046)
Tenure	0.002	(0.006)	0.004	(0.006)	0.013***	(0.005)	0.014***	(0.005)
$Tenure^2$	0.000**	(0.000)	0.000	(0.000)	-0.000	(0.000)	-0.000*	(0.000)
Constant	9.020***	(0.317)	7.825***	(0.291)	7.721***	(0.306)	9.112***	(0.349)
Selection Equation:								
Log Non-labour Income	-0.137***	(0.015)	-0.214***	(0.017)	-0.241***	(0.017)	-0.166***	(0.011)
Children Aged 0-2	-0.389**	(0.160)	0.196	(0.153)	-0.251	(0.189)	-0.070	(0.129)
Children Aged 3-4	-0.239	(0.165)	0.224	(0.147)	-0.206	(0.190)	-0.004	(0.131)
Children Aged 5-11	-0.284**	(0.138)	0.242*	(0.126)	-0.277*	(0.163)	-0.034	(0.112)
Children Aged 12-15	-0.269*	(0.139)	0.222*	(0.127)	-0.231	(0.164)	-0.026	(0.112)
Total Children	0.240*	(0.133)	-0.198*	(0.118)	0.214	(0.154)	0.057	(0.106)
In Relationship	0.335***	(0.054)	0.428***	(0.060)	0.553***	(0.069)	0.378***	(0.049)
Age	0.153***	(0.009)	0.155***	(0.009)	0.181***	(0.010)	0.163***	(0.008)
$Age^2/10$	-0.019***	(0.001)	-0.020***	(0.001)	-0.023***	(0.001)	-0.020***	(0.001)
Constant	-1.945***	(0.242)	-1.372***	(0.249)	-1.358***	(0.271)	-1.842***	(0.220)
ρ	-1.754***	(0.068)	-1.199***	(0.062)	-0.703***	(0.085)	-1.346***	(0.052)
σ	0.162***	(0.015)	-0.061***	(0.015)	-0.377***	(0.017)	0.201***	(0.013)
Observations	3961		4612		4475		6339	

* $p<0.10$, ** $p<0.05$, *** $p<0.01$. Standard Errors in parentheses. Included, but not reported are dummies for Health (5) Region (18), Occupation (8), Industry (9), Education (6), When Work (10), Organisation Type (8), Term of Contract (3), Mangerial Responsibility (3) in the main equation, and dummies for Health (5) Region (18), and Education (6) in the selection equation.

Table A.10: Predicted Income Equations (Males: Waves 12-15)

	Wave							
	12		13		14		15	
Main Equation:								
Age	0.091***	(0.013)	0.113***	(0.013)	0.139***	(0.017)	-0.107***	(0.013)
$Age^2/10$	-0.011***	(0.002)	-0.014***	(0.002)	-0.016***	(0.002)	0.013***	(0.001)
Incentive Pay	0.058	(0.051)	0.033	(0.053)	0.025	(0.065)	0.064	(0.059)
Full Time	0.544***	(0.090)	0.479***	(0.091)	0.535***	(0.111)	0.493***	(0.078)
Union Member	-0.003	(0.059)	-0.022	(0.063)	-0.082	(0.081)	0.020	(0.074)
Pension	0.081	(0.055)	0.087	(0.059)	0.077	(0.072)	0.059	(0.065)
Tenure	0.013*	(0.007)	0.029***	(0.007)	0.030***	(0.009)	0.037***	(0.007)
$Tenure^2$	0.000	(0.000)	-0.001***	(0.000)	-0.000	(0.000)	-0.001***	(0.000)
Constant	5.908***	(0.431)	5.607***	(0.438)	4.937***	(0.587)	12.090***	(0.499)
Selection Equation:								
Log Non-labour Income	-0.231***	(0.012)	-0.176***	(0.010)	-0.240***	(0.013)	-0.089***	(0.010)
Children Aged 0-2	0.134	(0.178)	-0.220	(0.210)	0.168	(0.196)	0.053	(0.149)
Children Aged 3-4	0.106	(0.180)	-0.300	(0.209)	0.053	(0.199)	-0.121	(0.147)
Children Aged 5-11	0.017	(0.149)	-0.272	(0.183)	-0.007	(0.165)	-0.056	(0.122)
Children Aged 12-15	0.048	(0.146)	-0.325*	(0.178)	0.004	(0.160)	-0.110	(0.123)
Total Children	0.027	(0.139)	0.360**	(0.173)	0.069	(0.153)	0.108	(0.116)
In Relationship	0.412***	(0.066)	0.496***	(0.065)	0.429***	(0.071)	0.225***	(0.049)
Age	0.171***	(0.009)	0.190***	(0.009)	0.189***	(0.010)	0.174***	(0.008)
$Age^2/10$	-0.022***	(0.001)	-0.024***	(0.001)	-0.023***	(0.001)	-0.021***	(0.001)
Constant	-2.913***	(0.220)	-3.631***	(0.218)	-3.296***	(0.246)	-3.744***	(0.202)
ρ	0.017	(0.056)	0.063	(0.052)	0.019	(0.065)	-1.915***	(0.073)
σ	0.185***	(0.012)	0.227***	(0.012)	0.334***	(0.013)	0.472***	(0.015)
Observations	5766		5611		4863		4910	

* p<0.10, ** p<0.05, *** p<0.01. Standard Errors in parentheses. Included, but not reported are dummies for Health (5) Region (18), Occupation (8), Industry (9), Education (6), When Work (10), Organisation Type (8), Term of Contract (3), Mangerial Responsibility (3) in the main equation, and dummies for Health (5) Region (18), and Education (6) in the selection equation.

$[m_k^-, m_k^+]$ but vanishes (for practical purposes) outside this.[26] Taking J_k equally-spaced points on each interval gives us the vector $m_k = (m_k^-, m_k^- + \delta_k, \ldots, m_k^+ - \delta_k, m_k^+)$ (where δ_k is the step size in dimension k, and is equal to $(m_k^+ - m_k^- + 1)/J_k$), and allows us to construct a K-dimensional grid of knots, which form the (standardized) means of the basis densities. Each mean $\mu_{\bar{\jmath}}$ is a K-vector with its kth element equal to the j_kth element of m_k. Defining the matrix R such that $RR' = \Omega$, the unstandardised means of the basis densities (i.e. in the original scale of the random effects) are then given by $R\mu_{\bar{\jmath}}$; while the variance of each basis density is RDR', where D is a diagonal matrix with non-zero elements $\frac{4}{9}\delta_k^2$.[27] The density of the random effects can then be specified as:

$$f(\omega_i, \eta_i) = \sum_{\bar{\jmath}} c_{\bar{\jmath}} N(R\mu_{\bar{\jmath}}, RDR') \tag{A.25}$$

where $c_{\bar{\jmath}} \geq 0$ are the mixing weights for the corresponding basis densities, and $\sum_{\bar{\jmath}} c_{\bar{\jmath}} = 1$. To avoid dealing with the constraints directly, we estimate the transformed parameters $a_{\bar{\jmath}}$, defined such that $c_{\bar{\jmath}} = \exp(a_{\bar{\jmath}}) / \sum_{\bar{\jmath}} \exp(a_{\bar{\jmath}})$

To prevent over-fitting the mixing parameters, a smoothing penalty is

[26] Because the component Gaussians extend over the entire real line, any estimated random effects distribution will also do so — hence, the "for practical purposes" proviso. In fact, it would be possible for a distribution estimated by a sufficiently coarse Gaussian mixture to place significant mass beyond the extremes of $[m_k^-, m_k^+]$.

[27] See Ghidey, Lesaffre, and Eilers (2004) for the justification behind this choice.

applied and the model is estimated by maximising the resulting penalised log-likelihood function $\ell_p = \ell - \pi(\lambda)$. Recasting equation (5.6) above at the individual level, and compressing the notation somewhat, we have:

$$u_i = Q_i\psi + W_i\xi_i + \varepsilon_i \tag{A.26}$$

where $Q_i\psi$ is a stacked reformulated version of the fixed components of (5.6), W_i is a stacked matrix of the variables for which we estimate random effects (including the intercept) and ξ_i is (ω_i, η_i). The log-likelihood and penalty are then given by:

$$\ell = \sum_{i=1}^{N} \ln\left\{\sum_{\bar{\jmath}} c_{\bar{\jmath}} N(Q_i\psi + W_i R\mu_{\bar{\jmath}}, Z_i RDR'Z_i' + \sigma^2 I_{T_i})\right\} \tag{A.27}$$

$$\pi = \frac{1}{2}\sum_{l=1}^{K}\left\{\lambda_l \sum_{\bar{\jmath}} (\Delta_k^3 a_{\bar{\jmath}})^2\right\} \tag{A.28}$$

where $\Delta_k^3 a_{\bar{\jmath}}$ is the 3rd-order forward difference function with respect to the kth element of $\bar{\jmath}$,[28] and $\lambda = (\lambda_1, \ldots, \lambda_K)$ is a K-vector of penalty parameters.

Because maximizing ℓ_p with respect to all the parameters can be computationally intensive, and to avoid having to deal directly with complicated constraints on the variance components, Ghidey, Lesaffre, and Eilers (2004) suggest a simplified fitting procedure that takes advantage of the Verbeke

[28] I.e. $\Delta_k^3 a_{\bar{\jmath}} = -a_{\bar{\jmath}} + 3a_{j_1,\ldots,j_k+1,\ldots,j_K} - 3a_{j_1,\ldots,j_k+2,\ldots,j_K} + a_{j_1,\ldots,j_k+3,\ldots,j_K}$

and Lesaffre (1997) result that the fixed parameters and the variance components are consistently estimated from the Gaussian linear mixed model, even when the normality assumption is violated. It is a slightly modified version of this algorithm that we use here.

(1) Fit a standard Gaussian linear mixed model (e.g. using Stata's **xtmixed** command).

(2) Holding the parameters from the estimated Gaussian model fixed, maximize ℓ_p with respect to the mixing parameters a.[29]

(3) Update $\hat{R}$ and $\hat{\psi}_*$ (where $\hat{\psi}_*$ denotes the subset of the $\hat{\psi}$ for which we estimate random coefficients) from their Gaussian LMM values, $\hat{R}^0$ and $\hat{\psi}_*^0$, as follows:

$$\hat{R} = \text{chol}\left(\hat{R}^0 \left\{\sum_{\bar{\jmath}} \hat{c}_{\bar{\jmath}}\mu_{\bar{\jmath}}\mu_{\bar{\jmath}}' + D + \left(\sum_{\bar{\jmath}} \hat{c}_{\bar{\jmath}}\mu_{\bar{\jmath}}\right)\left(\sum_{\bar{\jmath}} \hat{c}_{\bar{\jmath}}\mu_{\bar{\jmath}}\right)'\right\} \hat{R}^{0\prime}\right) \tag{A.29}$$

$$\hat{\psi}_* = \hat{\psi}_*^0 + \sum_{\bar{\jmath}} \hat{c}_{\bar{\jmath}}\hat{R}^0\mu_{\bar{\jmath}} \tag{A.30}$$

[29] The approach set out in Ghidey, Lesaffre, and Eilers (2004) supplements this with an additional step where $\hat{a}$ is held fixed and ℓ_p is maximized with respect to the β and $\ln\sigma$ parameters, alternating between maximisation with respect to each of a and $(\beta, \ln\sigma)$ until convergence. However, the number of parameters in our model makes the latter step very computationally intensive. (The Gaussian mixture means that the likelihood is no longer linear in the β parameters, so we have to calculate derivatives with respect to each parameter individually. Moreover, unlike maximisation with respect to a (where fixing β and $\ln\sigma$ means we need to evaluate each mixture density only once), maximizing with respect to β and $\ln\sigma$ requires us to explicitly recalculate the mixture density for every function evaluation.)

(4) Repeat (2) and (3) over a (K-dimensional) grid of λ values, selecting the optimum λ as that which minimises $AIC = -2\ell_p + 2\ \mathrm{trace}[H^{-1}(a)I(a)]$, where the latter term is the "effective degrees of freedom", and $H(x)$ and $I(x)$ are the negative Hessian matrices with respect to x from the penalised and unpenalised likelihoods, respectively.[30]

(5) Estimate the covariance matrix as $\hat{V} = H^{-1}(\psi, \mathrm{vec}(\Omega), \ln\sigma)$ (following e.g. Gray, 1992; Verweij and Van Houwelingen, 1994).[31]

The only real choices facing us in fitting the model are to fix the bounds m^- and m^+, and the number of basis points J_k in each dimension. With four random coefficients, the algorithm's memory requirements limit us to $J_k = 5$ for $k = 1, \ldots, 4$. With the number of basis points fixed, there is a trade-off between accurate fitting towards the edge of the distribution, and accurate fitting towards its centre. After some experimentation, we opted to set $m^- = -1.5$ and $m^+ = 1.5$, which seemed to provide the highest resolution grid without distorting the edges of the distribution.

We have written a Stata .ado command **pgm** to implement steps (2)-(5)

[30] Gray (1992) and Ghidey, Lesaffre, and Eilers (2004) suggest taking the negative Hessians with respect to the full parameter vector $(\psi, \mathrm{vec}(\Omega), \ln\sigma, a)$. We simplify this for computational reasons, and because any variance in the effective degrees of freedom arises primarily from the a parameters.

[31] The Gray (1992) variance estimator is usually defined as $\hat{V} = H^{-1}(x)I(x)H^{-1}(x)$, but this reduces to the Verweij and Van Houwelingen (1994) estimator $H^{-1}(x)$ here, because the penalty only operates with respect to the a parameters, which we treat as fixed in this calculation for computational reasons. Ghidey, Lesaffre, and Eilers (2004) provide evidence that the latter is typically close to the asymptotic covariance matrix.

of this algorithm following an initial fit using **xtmixed**, making extensive use of Stata's internal matrix programming language Mata, and in particular its **optimize()** routine.

A.8 Modelling heterogeneity in the reporting function

Although we do not pursue it further here, the modelling strategy employed in Chapter 5 could, with the imposition of some additional assumptions, also be employed to separately identify reporting heterogeneity and "true" heterogeneity. We briefly set sketch the nature of this extension here.

In Section 5.2 we suggested that we could model heterogeneity in the determinants of SWB and heterogeneity in the reporting function via the following equations:[32]

$$\begin{aligned} u^*_{it} &= q'_{it}\psi^*_{it} + \alpha^*_0 + \alpha^*_i + \epsilon^*_{it} \\ &= q'_{it}\psi^*_0 + \ddot{x}'_{kit}\omega^*_i + \alpha^*_0 + \alpha^*_i + \epsilon^*_{it} \end{aligned} \tag{A.31}$$

$$\begin{aligned} u_{it} &= u^*_{it}\phi_i + \alpha_0 + \alpha_i + \epsilon_{it} \\ &= q'_{it}\psi_{it} + \eta_i + \varepsilon_{it} \end{aligned} \tag{A.32}$$

[32]These equations are derived from a combination of the general equations (5.1) and (5.2), and the more specific form set out in equation (5.6), with some compression of notation.

In what followed, we made assumptions directly on the shape of the distribution of the random parameters in the latter equation ψ_{it}, and this left us unable to distinguish directly between ϕ_i and ψ^*_{it}.

An alternative approach would be to treat the ψ^*_{it}, α^*_i, ϕ_i, and α_i as separate parameters, each with their own parametrically specified distribution. The simplest assumption would treat be treat the ψ^*_i, α^*_i, and α_i as normal, and the ϕ_i as, say log-normal (to ensure that the reporting function is monotonic).

Defining $\tilde{\phi}_i = \phi_i - 1$ and substituting (A.31) into (A.32) then yields:

$$u_{it} = q'_{it}\psi_0 + q'_{it}\delta_i + \ddot{x}'_{kit}\theta_i + \alpha^*_0 + \lambda_i + \epsilon_{it} \tag{A.33}$$

with $\delta_i = \psi_0\tilde{\phi}_i$, $\theta_i = \omega^*_i + \omega^*_i\phi_i$, $\lambda_i = \alpha^*_i\tilde{\phi}_i + \alpha^*_i\tilde{\phi}_i + \alpha_i$, and $\varepsilon_{it} = \epsilon^*_{it} + \epsilon^*_{it}\tilde{\phi}_i + \epsilon_{it}$.

Unfortunately, estimating latent interaction models such as that set out in equation (A.33) is somewhat difficult (dealing with the normal-product, and lognormal-normal–product distributions resulting from the interactions between the random effects requires special techniques/software). And it is unclear what effect violations of (log-)normality would have on the estimates. We leave such the possibility of such investigations for future work.

A.9 PGM Tables

See Tables A.12 and A.12. The qualitative results are similar to those obtained from the LMM model with seven random coefficients. In some cases, interactions that were previously insignificant have become significant. There are now distinct age-patterns in the importance of consumption, reference income, and health; and there is also now a significant difference between the genders in the reference income coefficient, with men being more negatively affected than women. The negative correlation between the coefficients on consumption and reference income has also become statistically significant. The other coefficients are similar to their LMM estimates.

Table A.11: PGM Estimates of Heterogeneity in the Determinants of SWB in the BHPS (Random Coefficients)

	Consumption		Reference Income		Health	
Fixed Part						
Main effect	0.085***	(0.025)	-0.003	(0.011)	0.185***	(0.009)
Interactions w/:						
Employed			0.028	(0.023)		
Sex	0.014	(0.026)	0.030**	(0.015)	-0.010	(0.014)
Age	-0.008**	(0.004)	-0.005**	(0.002)	-0.004*	(0.002)
Age^2/10	0.000	(0.000)	0.001***	(0.000)	0.000**	(0.000)
Urban	0.091***	(0.027)	0.004	(0.015)	0.011	(0.015)
Secondary	-0.021	(0.033)	0.017	(0.018)	-0.025	(0.017)
Post-Secondary	-0.011	(0.040)	0.025	(0.022)	-0.048**	(0.022)
Random Part						
Standard Deviation	0.400***	(0.019)	0.177***	(0.010)	0.244***	(0.010)
Correlations w/:						
Reference Income	-0.208**	(0.083)				
Health	-0.165***	(0.060)	-0.018	(0.078)		
Intercept	-0.384***	(0.029)	0.054	(0.042)	-0.533***	(0.034)
Observations	55794					
Groups	11360					
Penalised Log-Likelihood	-77971.104					

* p<0.10, ** p<0.05, *** p<0.01. Standard errors in parentheses.

Table A.12: PGM Estimates of Heterogeneity in the Determinants of SWB in the BHPS (Main Effects and Intercept)

	Main		Intercept†	
(Log) Household Income	0.001	(0.008)	0.039**	(0.016)
Employed	-0.057*	(0.031)	-0.161***	(0.052)
Unemployed	-0.445***	(0.043)	-0.301***	(0.102)
Unemployed (Discouraged)	0.282***	(0.046)	-0.433***	(0.117)
Age	-0.024***	(0.003)		
$\text{Age}^2/10$	0.003***	(0.000)		
In Relationship	0.195***	(0.035)	0.165***	(0.052)
Married	-0.057*	(0.031)	0.140***	(0.044)
Widowed	-0.198***	(0.050)	0.280***	(0.067)
Divorced	-0.065	(0.049)	-0.071	(0.066)
Separated	-0.319***	(0.053)	-0.112	(0.098)
Household Size	-0.068***	(0.011)	0.024	(0.016)
Sex			0.043	(0.030)
Urban	-0.103***	(0.035)	0.060	(0.039)
Hours Social Leisure	0.139***	(0.029)	0.075*	(0.044)
Hours Worked	0.004	(0.025)	0.060	(0.038)
Hours Slept	1.245***	(0.201)	-0.324	(0.315)
Hours Slept^2	-0.069***	(0.010)	0.029*	(0.015)
Secondary	0.148	(0.134)	-0.319**	(0.135)
Post-Secondary	0.027	(0.146)	-0.279*	(0.148)
Has Children	0.060**	(0.025)	-0.085**	(0.042)
2 Children	0.054**	(0.027)	-0.044	(0.047)
3 Children	0.074	(0.052)	0.026	(0.078)
4 Children	-0.029	(0.094)	0.152	(0.153)
(Mean) (Log) Consumption			-0.061*	(0.031)
(Mean) Reference Income			-0.065***	(0.017)
(Mean) Health			0.376***	(0.014)
Constant			5.286***	(0.009)
Observations	55794			
Groups	7052			
Penalised Log-Likelihood	-77971.104			

* p<0.10, ** p<0.05, *** p<0.01. Standard Errors in parentheses. †Coefficients in the intercept column are for time-invariant individual-level means, $\bar{\bar{x}}_i$. Wave dummies included but not reported.

A.10 Measurement Error Correction

This appendix provides a simple illustration of how the inclusion of an imperfect proxy for measurement error in a regression model can reduce attenuation bias. The derivation of results applicable to the specific situation at hand is more involved (largely due to the more complicated expressions for the matrix inverses when more regressors are included). We nonetheless expect the result to generalize, and have confirmed this with a set of simulations.

Consider a simple measurement error problem, where the true model is

$$y_i = x_i^* \beta + \varepsilon_i \tag{A.34}$$

but we observe only a noisy measure of x_i^*:

$$x_i = x_i^* + v_i \tag{A.35}$$

As is well known, estimating Equation (A.34) using x_i in place of x_i^* results in an attenuated estimate of $\hat{\beta}$ that converges to only $\beta\sigma^2_{x^*}/(\sigma^2_{x^*} + \sigma^2_v) < \beta$.

Suppose that we have a noisy measure of v_i, which may also be related to the true measure x_i^* as follows:

$$w_i = v_i \gamma + x_i \theta + \omega_i \tag{A.36}$$

where the ω_i is independent of v_i, x_i^* and ε_i. Including this in the estimating equation as follows:

$$y_i = (x_i - w_i)\tilde{\beta} + w_i\hat{\phi} + e_i \tag{A.37}$$

will then reduce the bias in $\tilde{\beta}$. This is most obvious when w_i is a perfect measure of v_i: then $x_i - w_i = x_i^*$, $\tilde{\beta} \to \beta$, $\theta = 0$, and $\hat{\phi} \to 0$. More generally, when γ can take values other than 1, θ can take values other than 0, and σ_ω^2 may be greater than 0 we have (the derivation is provided in the next subsection):

$$\tilde{\beta} \to \frac{\sigma_{x^*}^2(\gamma(\gamma-\theta)\sigma_v^2 + \sigma_\omega^2)}{\sigma_{x^*}^2(\gamma(\gamma-\theta)\sigma_v^2 + \sigma_\omega^2) + \sigma_v^2(\sigma_\omega^2 - \theta(\gamma-\theta))}\beta \tag{A.38}$$

This will be attenuated if $\sigma_\omega^2 > \theta(\gamma-\theta)\sigma_{x^*}^2$ (which will always be true if $\theta > \gamma$), and amplified if $\sigma_\omega^2 < \theta(\gamma-\theta)\sigma_{x^*}^2$. More importantly, it is less attenuated than $\hat{\beta}$ if $\gamma(\gamma-\theta)\sigma_v^2 > \theta(\theta-\gamma)\sigma_{x^*}^2$. If γ, $\theta \geq 0$ then this requires simply that $\gamma > \theta$. Intuitively, provided that w_i contains more noise than signal, estimating (A.37) will reduce attenuation bias.[33]

[33] As it turns out, this is also the condition required for $\tilde{\beta}$ to outperform $\hat{\phi}$ as an estimate of β. This can be seen directly from the probability limit of $\hat{\phi}$:

$$\hat{\phi} \to \frac{\sigma_{x^*}^2((\gamma-1)(\gamma-\theta)\sigma_v^2 + \sigma_\omega^2)}{\sigma_{x^*}^2(\gamma(\gamma-\theta)\sigma_v^2 + \sigma_\omega^2) + \sigma_v^2(\sigma_\omega^2 - \theta(\gamma-\theta))}\beta \tag{A.39}$$

References

Akaike, Hirotugu. 1974. "A New Look at the Statistical Model Identification." *IEEE Transactions on Automatic Control* 19 (6):716–723.

Akerlof, George A. 1997. "Social Distance and Social Decisions." *Econometrica* 65 (5):1005–1028.

Amador, Francisco Javier, Rosa Marina González, and Juan de Dios Ortúzar. 2005. "Preference heterogeneity and willingness to pay for travel time savings." *Transportation* 32 (6):627–647.

Ameriks, John, Andrew Caplin, and John Leahy. 2003. "Wealth Accumulation and the Propensity to Plan." *Quarterly Journal of Economics* 118 (3):1007–1047.

Argyle, Michael. 1999. "Causes and Correlates of Happiness." In *Well-Being: The Foundations of Hedonic Psychology*, edited by Daniel Kahneman, Ed Diener, and Norbert Schwarz. New York: Russell Sage.

Arndt, Stephan, Carolyn Turvey, and Nancy C. Andreasen. 1999. "Correlating and Predicting Psychiatric Symptom Ratings: Spearmans r versus Kendalls tau Correlation." *Journal of Psychiatric Research* 33 (2):97–104.

Atkeson, Christopher G., Andrew W. Moore, and Stefan Schaal. 1997. "Locally Weighted Learning." *Artificial Intelligence Review* 11 (1):11–73.

Bafumi, Joseph and Andrew Gelman. 2006. "Fitting Multilevel Models When Predictors and Group Effects Correlate." Paper presented at the 2006 Annual Meeting of the Midwest Political Science Association.

Basu, Anirban and David Meltzer. 2007. "Value of Information on Preference Heterogeneity and Individualized Care." *Medical Decision Making* 27 (2):112.

Bauermeister, Roy F., Jennifer D. Campbell, Joachim I. Krueger, and Kathleen D. Vohs. 2003. "Does High Self-Esteem Cause Better Performance, Interpersonal Success, Happiness, or Healthier Lifestyles?" *Psychological Science in the Public Interest* 4 (1):1–44.

Belk, Russell W. 1985. "Materialism: Trait Aspects of Living in the Material World." *Journal of Consumer Research* :265–280.

Bertrand, Marianne and Sendhil Mullainathan. 2001. "Do People Mean What They Say? Implications for Subjective Survey Data." *American Economic*

Review 91 (2):67–72. (Papers and Proceedings of the Hundred Thirteenth Annual Meeting of the American Economic Association).

Blackwell, David. 1953. "Equivalent Comparisons of Experiments." *The Annals of Mathematical Statistics* :265–272.

Blanchflower, David G. and Andrew J. Oswald. 2004. "Well-Being Over Time in Britain and the USA." *Journal of Public Economics* 88:1359–1386.

———. 2008. "Is Well-being U-shaped over the Life Cycle?" *Social Science & Medicine* 66 (8):1733–1749.

Bogaerts, Kris and Emmanuel Lesaffre. 2008. "Estimating Local and Global Measures of Association for Bivariate Interval Censored Data with a Smooth Estimate of the Density." *Statistics in Medicine* 27 (28).

Bohrnstedt, G.W. and A.S. Goldberger. 1969. "On the Exact Covariance of Products of Random Variables." *Journal of the American Statistical Association* :1439–1442.

Boskin, Michael J. and Eytan Sheshinski. 1978. "Optimal Redistributive Taxation when Individual Welfare Depends upon Relative Income." *The Quarterly Journal of Economics* :589–601.

Bourguignon, Francois, Martin Fournier, and Marc Gurgand. 2007. "Selection Bias Corrections Based on the Multinomial Logit Model: Monte Carlo Comparisons." *Journal of Economic Surveys* 21 (1):174–205.

Brandt, Richard. 1979. *A Theory of the Good and the Right.* Oxford: Clarendon Press.

Brickman, P., D. Coates, and R. Janoff-Bulman. 1978. "Lottery Winners and Accident Victims: Is Happiness Relative?" *Journal of Personality and Social Psychology* 36.

Brown, Gordon D. A., Jonathan Gardner, Andrew .J. Oswald, and Jing Qian. 2006. "Does Wage Rank Affect Employees' Wellbeing?" Unpublished Manuscript, University of Warwick.

Brown, Sarah, Karl Taylor, and Stephen Wheatley Price. 2005. "Debt and Distress: Evaluating the Psychological Cost of Credit." *Journal of Economic Psychology* 26 (5):642–663.

Bruin, Adrian and Rainer Winkelmann. 2008. "Happiness Functions with Preference Interdependence and Heterogeneity: The Case of Altruism within the Family." Unpublished Manuscript.

Brülde, Bengt. 2007a. "Happiness and the Good Life: Introduction and Conceptual Framework." *Journal of Happiness Studies* 8:1–14.

———. 2007b. "Happiness Theories of the Good Life." *Journal of Happiness Studies* 8:15–49.

Burnham, Kenneth P. and David R. Anderson. 2003. *Model Selection and Multi-Model Inference: A Practical and Theoretical Approach.* Springer, 3rd ed.

———. 2004. "Multimodel Inference: Understanding the AIC and BIC in Model Selection." *Sociological Methods Research* 33:261–304.

Caplin, Andrew and John Leahy. 2001. "Psychological Expected Utility Theory and Anticpatory Feelings." *Quarterly Journal of Economics* 116 (1):55–80.

Cappelli, Peter and Peter D. Sherer. 1988. "Satisfaction, Market Wages, and Labor Relations: An Airline Study." *Industrial Relations* 27 (1).

Carver, Robert. 2006. "Implications of the Rental Market for UK Housing Policy." SSRN Working Paper Series, No. 954679.

Chamberlain, Gary. 1980. "Analysis of Covariance with Qualitative Data." *Review of Economic Studies* 47:225–238.

Clark, Andrew E., Ed Diener, Yannis Georgellis, and Richard E. Lucas. 2008. "Lags and Leads in Life Satisfaction: A Test of the Baseline Hypothesis." *The Economic Journal* 118 (529):F222–F243.

Clark, Andrew E., Fabrice Etile, Fabien Postel-Vinay, Claudia Senik, and Karine Van der Straeten. 2005. "Heterogeneity in Reported Well-Being: Evidence from Twelve European Countries." *Economic Journal* 115 (502):C118–C132.

Clark, Andrew E., Paul Frijters, and Michael A. Shields. 2008. "Relative Income, Happiness, and Utility: An Explanation for the Easterlin Paradox and Other Puzzles." *Journal of Economic Literature* 46 (1):95–144.

Clark, Andrew E., Yannis Georgellis, and Peter Sanfey. 1998. "Job Satisfaction, Wage Changes and Quits: Evidence from Germany." *Research in Labor Economics* 17.

———. 2001. "Scarring: The Psychological Impact of Past Unemployment." *Economica* 68:221–241.

Clark, Andrew E., Nicolai Kristensen, and Niels Westergaard-Nielsen. 2007. "Job-Satisfaction and Co-worker Wages: Status or Signal?" Unpublished Manuscript.

Clark, Andrew E. and Andrew J. Oswald. 1996. "Satisfaction and Comparison Income." *Journal of Public Economics* 61:359–381.

———. 2002. "A Simple Statistical Method for Measuring How Life Events

Affect Happiness." *International Journal of Epidemiology* 31 (6):1139–1144.

Clark, Andrew E. and Claudia Senik. 2008. "Who Compares to Whom? The Anatomy of Income Comparisons in Europe." Unpublished Manuscript.

Clements, Kenneth W., Yanrui Wu, and Jing Zhang. 2006. "Comparing International Consumption Patterns." *Empirical Economics* 31:1–30.

Cooley, Charles Horton. 1902. *Human Nature and the Social Order.* New York: Scribner's.

Das, M. and A. Van Soest. 1999. "A panel data model for subjective information on household income growth." *Journal of Economic Behaviour and Organisation* 40:409–426.

Deaton, Angus and Salman Zaidi. 2002. "Guidelines for Constructing Consumption Aggregates for Welfare Analysis." World Bank.

DeNeve, K. M. and H. Cooper. 1998. "The Happy Personality: A Meta-Analysis of 137 Personality Traits and Subjective Well-Being." *Psychological Bulletin* 124:197–229.

Di Tella, R., J. P. Haisken-De New, and R. MacCulloch. 2007. "Happiness Adaptation to Income and to Status in an Individual Panel." NBER Working Paper 12159.

Diener, Ed and Marissa Diener. 1995. "Cross-Cultural Correlates of Life Satisfaction and Self-Esteem." *Journal of Personality and Social Psychology* 68:653–663.

Diener, Ed, Robert. A. Emmons, Randy J. Larsen, and Sharon Griffin. 1985. "The Satisfaction with Life Scale." *Journal of Personality Assessment* 49 (1):71–75.

Diener, Ed and Richard E. Lucas. 1999. "Personality and Subjective Well-Being." In *Well-Being: The Foundations of Hedonic Psychology*, edited by D. Kahneman, E. Diener, and N. Schwarz. New York: Russell Sage.

Diener, Ed, Carol Nickerson, Richard E. Lucas, and Ed Sandvik. 2002. "Dispositional Affect and Job Outcomes." *Social Indicators Research* 59 (3):229–259.

Diener, Ed and Shigehiro Oishi. 2000. "Money and Happiness: Income and Subjective Well-Being across Nations." In *Culture and Subjective Well-Being*, edited by E. Diener and E. M. Suh. The MIT Press, 185–218.

Draca, Mirko, Raffaella Sadun, and John van Reenen. 2006. "Producitivity and ICT: A Review of the Evidence." CEP Discussion Paper NO. 749.

Duesenberry, J. S. 1949. *Income, Savings and the Theory of Consumer Behavior.* Cambridge, MA: Harvard University Press.

Dunn, Elizabeth W., Timothy D. Wilson, and Daniel Gilbert. 2003. "Location, Location, Location: The Misprediction of Satisfaction in Housing Lotteries." *Personality and Social Psychology Bulletin* 29 (11):1421–1432.

Easterlin, R. A. 2004. "Life Cycle Happiness and Its Sources." Department of Economics, University of Southern California.

Elster, Jon and George Loewenstein. 1992. "Utility from Memory and Anticipation." In *Choice Over Time*, edited by George Loewenstein and Jon Elster. New York: Russell Sage Foundation, 213–234.

European Bank for Reconstruction and Development. 2004. *Transition Report*. London: European Bank for Reconstruction and Development.

Ferrer-i-Carbonell, Ada. 2005. "Income and Well-Being: An Empirical Analysis of the Comparison Effect." *Journal of Public Economics* 89 (997–1017).

Ferrer-i-Carbonell, Ada and Paul Frijters. 2004. "How Important is Methodology for the Estimates of the Determinants of Happiness?" *The Economic Journal* 114:641–659.

Fisher, Irving. 1930. *The Theory of Interest*. London: MacMillan.

Fitzgerald, John, Peter Gottschalk, and Robert Moffitt. 1998. "An Analysis of Sample Attrition in Panel Data: The Michigan Panel Study of Income Dynamics." *Journal of Human Resources* 33 (2):251–299.

Fourcade, Marion and Kieran Healy. 2007. "Moral Views of Market Society." *Annual Review of Sociology* 33:285–311.

Frank, Robert H. 1984. *Choosing the Right Pond*. Oxford, UK: Oxford University Press.

———. 1985. "The Demand for Unobservable and Other Nonpositional Goods." *American Economic Review* 75 (1):101–116.

———. 1999. *Luxury Fever: Why Money Fails to Satisfy in an Era of Excess.* New York: The Free Press.

———. 2003. "Does Absolute Income Matter?" In *The Paradoxes of Happiness in Economics*. Milan, Italy.

Frederick, Shane and George Loewenstein. 1999. "Hedonic Adaptation." In *Well-Being: The Foundations of Hedonic Psychology*, edited by D. Kahneman, E. Diener, and N. Schwarz. New York: Russell Sage.

Frederick, Shane, George Loewenstein, and Ted O'Donoghue. 2002. "Time Discounting and Time Preference: A Critical Review." Unpublished Manuscript.

Frey, Bruno S. 2008. *Happiness: a Revolution in Economics*. MIT Press.

Friedman, Benjamin. 2005. *The Moral Consequences of Economic Growth*. New York: Alfred A. Knopf.

Frijters, Paul, Ingo Geishecker, John P. Haisken-De New, and Michael A. Shields. 2006. "Can the Large Swings in Russian Life Satisfaction Be Explained By Ups and Downs in Real Incomes?" *Scandinavian Journal of Economics* In Press.

Frijters, Paul, John P. Haisken-DeNew, and Michael A. Shields. 2004. "Money Does Matter! Evidence from Increasing Real Income and Life Satisfaction in East Germany following Reunification." *American Economic Review* 94 (3):730–740.

Furnham, Adrian and Helen Cheng. 2000. "Lay Theories of Happiness." *Journal of Happiness Studies* 1:227–246.

Gecas, Viktor and Monica A. Seff. 1990. "Social Class and Self-Esteem: Psychological Centrality, Compensation, and the Relative Effects of Work and Home." *Social Psychology Quarterly* 53 (2):165–173.

Gelman, Andrew. 2008. "Scaling Regression Inputs by Dividing by Two Standard Deviations." *Statistics in Medicine* 27 (15):2865–2873.

Ghidey, Wendimagegn, Emmanuel Lesaffre, and Paul Eilers. 2004. “Smooth Random Effects Distribution in a Linear Mixed Model.” *Biometrics* 60 (4):945–953.

Ghidey, Wendimagegn, Emmanuel Lesaffre, and Geert Verbeke. forthcoming. “A Comparison of Methods for Estimating the Random Effects Distribution of a Linear Mixed Model.” *Statistical Methods in Medical Research* .

Gilbert, Daniel, Elizabeth C. Pinel, Timothy D. Wilson, Stephen J Blumberg, and Thalia Wheatley. 1998. “Immune Neglect: A source of durability bias in affective forecasting.” *Journal of Personality and Social Psychology* 75 (3):617–638.

Gilbert, Daniel T. 2006. *Stumbling on Happiness.* Harper Collins.

Gilbert, Daniel T., Matthew A. Killingsworth, Rebecca N. Eyre, and Timothy D. Wilson. 2009. “The Surprising Power of Neighborly Advice.” *Science* 323 (5921):1617.

Gilbert, Daniel T. and Timothy D. Wilson. 2000. “Miswanting: Some Problems in the Forecasting of Future Affective States.” In *Feeling and Thinking: The Role of Affect in Social Cognition*, edited by J. P. Forgas. Cambridge University Press, 178–197.

———. 2007. "Prospection: Experiencing the Future." *Science* 317 (5843):1351.

Goldberger, Arthur S. 1968. "The Interpretation and Estimation of Cobb-Douglas Functions." *Econometrica* 36 (3/4):464–472.

Goldstein, Harvey. 1997. "Methods in School Effectiveness Research." *School effectiveness and school improvement* 8 (4):369–395.

Goldstein, Harvey, Jon Rasbash, Min Yang, Geoffrey Woodhouse, Huiqi Pan, Desmond Nuttall, and Sally Thomas. 1993. "A Multilevel Analysis of School Examination Results." *Oxford review of education* 19 (4):425–433.

Goodman, Leo A. 1960. "On the Exact Variance of Products." *Journal of the American Statistical Association* :708–713.

Gorsuch, Richard L. 1983. *Factor Analysis.* Hillsdale, NJ: Lawrence Erlbaum, 2 ed.

Goskomstat. 2002. *All Russia Population Census - Rural Settlements Grouped by Population Size.* Goskomstat.

Gould, John P. 1974. "Risk, Stochastic Preference, and the Value of Information." *Journal of Economic Theory* 8 (1):64–84.

Graham, Carol, Andrew Eggers, and Sandip Sukhtankar. 2004. "Does Happiness Pay? An Exploration Based on Panel Data from Russia." *Journal of Economic Behaviour and Organisation* 55:319–342.

Graham, Carol and Andrew Felton. 2006. "Inequality and Happiness: Insights from Latin America." *Journal of Economic Inequality* 4 (1):107–122.

Gray, Robert J. 1992. "Flexible Methods for Analyzing Survival Data Using Splines, with Applications to Breast Cancer Prognosis." *Journal of the American Statistical Association* 87 (420):942–951.

Griffin, James. 1986. *Well-Being: Its Meaning, Measurement and Moral Importance*. Oxford: Clarendon Press.

Griliches, Z. 1992. "The Search for R&D Spillovers." *The Scandinavian Journal of Economics* 94:29–47.

Hamermesh, D. S. 2001. "The Changing Distribution of Job Satisfaction." *Journal of Human Resources* 36:1–30.

Hansen, Lars Peter. 1982. "Large Sample Properties of Generalized Method of Moments Estimators." *Econometrica* 50 (4):1029–1054.

Hautsch, Nikolaus and Stefan Klotz. 2003. "Estimating the Neighborhood Influence on Decision Makers: Theory and an Application on the Analysis of Innovation Decisions." *Journal of Economic Behavior and Organization* 52 (1):97–113.

Haybron, Daniel. 2005. "On Being Happy and Unhappy." *Philosophy and Phenomenological Research* 71 (2):287–317.

———. 2007. "Life Satisfaction, Ethical Reflection, and the Science of Happiness." *Journal of Happiness Studies* 8 (1):99–138.

Headey, Bruce, Ruud Muffels, and Mark Wooden. 2008. "Money Does not Buy Happiness: Or Does It? A Reassessment Based on the Combined Effects of Wealth, Income and Consumption." *Social Indicators Research* 87 (1):65–82.

Headey, Bruce and Mark Wooden. 2004. "The Effects of Wealth and Income on Subjective Well-Being and Ill-Being." *Economic Record* 80:S24–S33.

Heckman, James J. 1979. "Sample Selection Bias as a Specification Error." *Econometrica* 47 (1):153–162.

Heeringa, Steven G. 1997. "RLMS Sample Attrition, Replenishment, and Weighting in Rounds V-VII." Institute for Social Research, University of Michigan.

Helliwell, John F. and Haifang Huang. 2005. "How's the Job? Well-Being and Social Capital in the Workplace." NBER.

Hendry, David and Hans-Martin Krolzig. 2002. "New Developments in Automatic Geneneral-to-Specific Modelling." Unpublished Manuscript, Economics Department, Oxford University.

Herrnstein, Richard J. and Drazen Prelec. 1991. "Melioration: A theory of Distributed Choice." *The Journal of Economic Perspectives* :137–156.

Hilton, Ronald W. 1979. "The Determinants of Cost Information Value: An Illustrative Analysis." *Journal of Accounting Research* :411–435.

———. 1981. "The determinants of information value: Synthesizing some general results." *Management Science* :57–64.

Hirsch, Fred. 1976. *The Social Limits of Growth.* Cambridge, MA: Harvard University Press.

Hirschman, Albert O. and Michael Rothschild. 1973. "The Changing Tolerance for Income Inequality in the Course of Economic Development." *Quarterly Journal of Economics* 87 (4):544–566.

Hynes, Stephen, Nick Hanley, and Riccardo Scarpa. 2008. "Effects on Welfare Measures of Alternative Means of Accounting for Preference Heterogeneity in Recreational Demand Models." *American Journal of Agricultural Economics* 90 (4):1011–1027.

Itami, Hiroyuki. 1977. *Adaptive Behavior: Management Control and Information Analysis.* American Accounting Association.

Jevons, Herbert S. 1905. *Essays on Economics.* London: MacMillan.

Jevons, William S. 1888. *The Theory of Political Economy.* London: MacMillan.

Johns, Helen and Paul Ormerod. 2007. *Happiness, Economics, and Public Policy.* London: Institute of Economic Affairs.

Jorgenson, Dale W. and J. Stiroh, Kevin. 2000. "Raising the Speed Limit: U.S. Economic Growth in the Information Age." *Brookings Papers on Economic Activity* 2000 (1):125–235.

Jouini, M. and R.T. Clemen. 1996. "Copula Models for Aggregating Expert Opinions." *Operations Research* 44:444–457.

Kahneman, Daniel. 1994. "New Challenges to the Rationality Assumption by." *Journal of Institutional and Theoretical Economics* 150 (1):18–36.

———. 1999. "Objective Happiness." In *Well-Being: The Foundations of Hedonic Psychology*, edited by D. Kahneman, E. Diener, and N. Schwarz. New York: Russell Sage Foundation.

Kahneman, Daniel, Ed Diener, and Norbert Schwarz, editors. 1999. *Well-Being: The Foundations of Hedonic Psychology.* New York: Russell Sage Foundation.

Kahneman, Daniel, Jack L. Knetsch, and Richard H. Thaler. 1991. "Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias." *Journal of Economic Perspectives* 5 (1):193–206.

Kahneman, Daniel and Alan B. Krueger. 2006. "Developments in the Measurement of Subjective Well-Being." *Journal of Economic Perspectives* 20 (1):3–24.

Kahneman, Daniel, Alan B. Krueger, David Schkade, Norbert Schwarz, and Arthur A. Stone. 2006. "Would You Be Happier If You Were Richer? A Focusing Illusion." *Science* 312:1908–1910.

Kahneman, Daniel and Jackie Snell. 1992. "Predicting a Changing Taste: Do People Know What They Will Like?" *Journal of Behavioral Decision Making* 5 (3).

Kahneman, Daniel, Peter P. Wakker, and Rakesh K. Sarin. 1997. "Back to Bentham? Explorations of Experienced Utility." *Quarterly Journal of Economics* 112:375–405.

Kan, Man Yee and Jonathan Gershuny. 2006. "Infusing Time Diary Evidence into Panel Data: An exercise in calibrating time-use estimates for the BHPS." ISER Working Paper 2006–19.

Kaplow, Louis. 2008. "Optimal Policy with Heteregeneous Preferences." *BE Journal of Economic Analysis and Policy: Advances* 1.

Karelis, Charles. 2007. *The Persistence of Poverty: Why the Economics of the Rich Can't Help the Poor*. New Haven: Yale University Press.

Kasser, Tim and Richard M. Ryan. 1993. "A Dark Side of the American Dream: Correlates of Financial Success as a Central Life Aspiration." *Journal of Personality and Social Psychology* 65:410–410.

———. 1996. "Further Examining the American Dream: Differential Correlates of Intrinsic and Extrinsic Coals." *Personality and Social Psychology Bulletin* 22:280–287.

Keeney, Mary J. 2007. "Measuring Irish Capital." *The Economic and Social Review* 38 (1):25–62.

Kendall, Maurice and Jean D. Gibbons. 1990. *Rank Correlation Methods.* Oxford University Press, 5 ed.

Kermer, Deborah A., Erin Driver-Linn, TImothy D. Wilson, and Daniel T. Gilbert. 2006. "Loss Aversion Is an Affective Forecasting Error." *Psychological Science* 17 (8):649–653.

Kingdon, Geeta Gandhi and John Knight. 2007. "Community, Comparisons and Subjective Well-Being in a Divided Society." *Journal of Economic Behavior and Organization* 64:69–90.

Klugman, Jeni and Jeanine Braithwaite. 1998. "Poverty in Russia During the Transition: An Overview." *The World Bank Research Observer* 13:37–58.

Knight, John, Lina Song, and Ramani Gunatilaka. 2009. "Subjective Well-Being and its Determinants in Rural China." *China Economic Review* forthcoming.

Knutson, Brian, Grace W. Fong, Shannon M. Bennet, Charles M. Adams, and Daniel Hommer. 2003. "A region of mesial prefrontal cortex tracks monetarily rewarding outcomes: characterization with rapid event-related fMRI." *NeuroImage* 18:263–272.

Komárek, Arnošt and Emmanuel Lesaffre. 2006. "Bayesian Semi-Parametric Accelerated Failure Time Model for Paired Doubly Interval-Censored Data." *Statistical Modelling* 6 (1):3.

———. 2008. "Generalized Linear Mixed Model with a Penalized Gaussian Mixture as a Random Effects Distribution." *Computational Statistics and Data Analysis* 52 (7):3441–3458.

Komárek, Arnošt, Emmanuel Lesaffre, and Joan F. Hilton. 2005. "Accelerated Failure Time Model for Arbitrarily Censored Data with Smoothed Error Distribution." *Journal of Computational and Graphical Statistics* 14 (3):726–745.

Köszegi, Botond. 2005. "Utility from Anticipation and Personal Equilibrium." Unpublished Manuscript, University of California, Berkely.

Köszegi, Botond and Matthew Rabin. 2006. "A Model of Reference-Dependent Preferences." *Quarterly Journal of Economics* 121 (4):1133–1166.

———. forthcoming. "Reference-Dependent Risk Attitudes." *American Economic Review* .

Kreft, Ita G. G., Jan De Leeuw, and Leona S. Aiken. 1995. "The Effect of Different Forms of Centering in Hierarchical Linear Models." *Multivariate Behavioral Research* 30 (1):1–21.

Kullback, S. and R. A. Leibler. 1951. "On Information and Sufficiency." *Annals of Mathematical Statistics* 22:79–86.

Laird, Nan M. 1978. "Nonparametric Maximum Likelihood Estimation of a Mixing Distribution." *Journal of the American Statistical Association* 73:805–811.

Laird, Nan M. and James H. Ware. 1982. "Random-effects Models for Longitudinal Data." *Biometrics* 38:963–974.

Layard, Richard. 1980. "Human Satisfactions and Public Policy." *The Economic Journal* 90 (360):737–750.

———. 2005a. *Happiness: Lessons from a New Science*. London: Allen Lane.

———. 2005b. "Rethinking Public Economics: The Implications of Rivalry and Habit." In *Economics and Happiness - Framing the Analysis*, edited by Luigino Bruni and Pier Luigi Porta. Oxford, UK: Oxford University Press.

Layard, Richard, Guy Mayraz, and Stephen J. Nickell. 2008. "The Marginal Utility of Income." *Journal of Public Economics* 92 (8–9):1846–1857.

Lelkes, Orsolya. 2006. "Tasting Freedom: Happiness, Religion and Economic Transition." *Journal of Economic Behavior and Organization* 57:173–194.

Levinson, Arik. 2009. "Valuing Public Goods Using Happiness Data: The Case of Air Quality." NBER Working Paper 15156.

Lévy-Garboua, L. and C. Montmarquette. 2004. "Reported Job Satisfaction: What Does it Mean?" *Journal of Socio-Economics* 33:135–151.

Liberman, Nira and Yaacov Trope. 2008. "The Psychology of Transcending the Here and Now." *Science* 322 (5905):1201–1205.

Lindley, Dennis V. 1985. *Making Decisions*. Wiley, 2nd edition ed.

Linville, Patricia W. and Gregrory W. Fischer. 1991. "Preferences for Separating or Combining Events." *Journal of Personality and Social Psychology* 60:5–23.

Loewenstein, George. 1987. "Anticipation and the Value of Delayed Consumption." *The Economic Journal* 97:666.

Loewenstein, George and Daniel Adler. 1995. "A Bias in the Prediction of Tastes." *The Economic Journal* 105:929–937.

Loewenstein, George and Erik Angner. 2003. "Predicting and Indulging Changing Preferences." In *Time and Decision: Perspectives on Intertemporal Choice*, edited by George Loewenstein, D. Read, and R. Baumeister. New York: Russell Sage.

Loewenstein, George and Shane Frederick. 1997. "Predicting Reactions to Environmental Change." In *Environment, Ethics, and Behavior: The Psychology of Environmental Valuation and Degradation*, edited by Max H. Bazerman, David M. Messick, Ann E. Tenbrussel, and Kimberly A. Wade-Benzoni. New York: Russell Sage, 52–72.

Loewenstein, George, Ted O'Donoghue, and Matthew Rabin. 2003. "Projection Bias in Predicting Future Utility." *Quarterly Journal of Economics* 118:1209–1248.

Loewenstein, George and David Schkade. 1999. "Wouldn't It Be Nice? Predicting Future Feelings." In *Well-Being: The Foundations of Hedonic Psychology*, edited by Daniel Kahneman, Ed Diener, and Norbert Schwarz. New York: Russell Sage Foundation.

Lucas, Richard E. 2005. "Time Does Not Heal All Wounds." *Psychological Science* 16 (12):945.

Lucas, Richard E., Andrew E. Clark, Yannis Georgellis, and Ed Diener. 2003. "Reexamining Adaptation and the Set Point Model of Happiness: Reactions to Changes in Marital Status." *Journal of Personality and Social Psychology* 84 (3):527–539.

Luttmer, Erzo F. P. 2005. "Neighbors as negatives: Relative earnings and well-being." *Quarterly Journal of Economics* 120:963–1002.

Lykken, David and Auke Tellegen. 1996. "Happiness is a Stochastic Phenomenon." *Psychological Science* 7:186–189.

Lykova, Tatiana, Ekaterina Petrova, Sergei Sivaev, and Raymond Struyk.

2004. "Participation in a Decentralised Housing Allowance Programme in a Transition Economy." *Housing Studies* 19 (4):617–634.

Lyubomirksy, Sonja, Chris Tkach, and M. Robin DiMatteo. 2006. "What Are the Differences between Happiness and Self-Esteem?" *Social Indicators Research* :363–404.

Lyubomirsky, Sonja, Laura King, and Ed Diener. 2005. "The Benefits of Frequent Positive Affect: Does Happiness Lead to Success?" *Psychological Bulletin* 131 (6):803–805.

Lyubomirsky, Sonja, Kennon M. Sheldon, and David Schkade. 2005. "Pursuing Happiness: The Architecture of Sustainiable Change." *Review of General Psychology* 9 (2):111–131.

Magder, Laurence S. and Scott L. Zeger. 1996. "A Smooth Nonparametric Estimate of a Mixing Distribution Using Mixtures of Gaussians." *Journal of the American Statistical Association* 91 (435):1141–1151.

Manski, Charles F. 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *The Review of Economic Studies* 60 (3):531–542.

Marschak, Jacob and Roy Radner. 1972. *Economic Theory of Teams*. Yale Univ. Press.

Maslow, Abraham H. 1943. "A Theory of Human Motivation." *Psychological Review* 50:370–396.

McBride, Michael. 2001. "Relative Income Effect on Subjective Well-Being in the Cross-Section." *Journal of Economic Behaviour and Organisation* 45 (251–278).

McLachlan, Geoffrey J. and Kaye E. Basford. 1988. *Mixture Models.* New York: Marcel Dekker.

Mead, George Herbert. 1934. *Mind, Self and Society.* Chicago: University of Chicago Press.

Melenberg, Betrand. 1992. *Micro-Economic Models of Consumer Behaviour and Welfare.* Ph.D. thesis, Tilburg University.

Meyer, Bruce D. and James X. Sullivan. 2003. "Measuring the Well-Being of the Poor Using Income and Consumption." *Journal of Human Resources* 38. (Special Issue on Income Volatility and Implications for Food Assistance Programs).

———. 2007. "Further Results on Measuring the Well-Being of the Poor Using Income and Consumption." NBER Working Paper 13413.

Mill, John Stuart. 1861. *Utilitarianism.* London: Longmans.

Mu, Ren. 2003. "Risk, Consumption, Wealth and Human Capital: Evidence from Russia." Michigan State University.

Mundlak, Yair. 1987. "On the Pooling of Time Series and Cross Section Data." *Econometrica* 46 (1):69–85.

Nadaraya, E. A. 1964. "On Estimating Regression." *Theory of Probability and its Applications* 9:141.

Nickell, Stephen J. 1981. "Biases in Dynamic Models with Fixed Effects." *Econometrica* 49:1417–1426.

Nussbaum, Martha. 1993. "Non-Relative Virtues: An Aristotelian Approach." In *The Quality of Life*, edited by Martha Nussbaum and Amartya Sen. Oxford University Press, 242–269.

———. 2000. *Women and human development: The capabilities approach.* Cambridge Univ Press.

Nussbaum, Martha and Amartya Sen, editors. 1993. *The Quality of Life.* Oxford: Oxford University Press.

Oswald, Andrew J. 1997. "Happiness and Economic Performance." *Economic Journal* 107:1815–1831.

———. 2008. "On the Curvature of the Reporting Function from Objective Reality to Subjective Feelings." *Economics Letters* 100 (3):369–372.

Oswald, Andrew J. and Nattavudh Powdthavee. 2008. "Does Happiness Adapt? A Longitudinal Study of Disability with Implications for Economists and Judges." *Journal of Public Economics* 92 (5–6):1061–1077.

Parfit, Derek. 1984. *Reasons and Persons.* Oxford: Clarendon Press.

Powdthavee, Nattavudh. 2008. "Putting a Price Tag on Friends, Relatives, and Neighbours: Using Surveys of Life-Satisfaction to Value Social Relationships." *Journal of Socio-Economics* 37:1459–1480.

———. 2009. "Think Having Children Will Make You Happy?" *The Psychologist* 22 (4):308–311.

Rabe-Hesketh, Sophia and Anders Skrondal. 2008. *Multilevel and Longitudinal Modeling Using Stata.* Stata Corp.

Rabin, Matthew. 2000. "Diminishing Marginal Utility of Wealth Cannot Explain Risk Aversion." In *Choices, Values, and Frames*, edited by Daniel Kahneman and Amos Tversky. New York: Cambridge University Press.

Railton, Peter. 1986. "Moral Realism." *Philosophical Review* 95 (2):163–207.

Ravallion, Martin and Michael Lokshin. 2001. "Identifying Welfare Effects from Subjective Questions." *Economica* 68:335–357.

Rice, Nigel and Andrew Jones. 1997. "Multilevel models and health economics." *Health Economics* 6 (6):561–575.

Rice, Nigel and Alastair H. Leyland. 1996. "Multilevel models: applications to health data." *Journal of health services research & policy* 1 (3):154.

Richins, Marsha L. and Scott Dawson. 1992. "A Consumer Values Orientation for Materialism and its Measurement: Scale Development and Validation." *Journal of Consumer Research* 19 (3):303.

Robins, James M. and Andrea Rotnitzky. 1995. "Semiparametric Efficiency in Multivariate Regression Models with Missing Data." *Journal of the American Statistical Association* 90 (429):122–129.

Robins, James M., Andrea Rotnitzky, and Lue Ping Zhao. 1995. "Estimation of Regression Coefficients When Some Regressors Are Not Always Observed." *Journal of the American Statistical Association* 89 (427):846–866.

Romanik, Clare T. 1998. "Analysis of Affordability of Cost Recovery: Communal and Network Energy Services." World Bank.

Rosenberg, Morris and Leonard I. Pearlin. 1978. "Self-Esteem Among Children and Adults." *American Journal of Sociology* 84 (1):53–77.

Sagers, Matthew J. 1989. "News Notes." *Soviet Geography* 30 (1):75–76.

Samuelson, Paul. 1937. "A Note on the Measurement of Utility." *Review of Economic Studies* 4:155–161.

Sargan, J. D. 1958. "The Estimation of Economic Relationships Using Instrumental Variables." *Econometrica* 26:393–415.

Schwarz, Gideon E. 1978. "Estimating the dimension of a model." *Annals of Statistics* 6 (2):461–464.

Schwarz, Norbert. 1990. "Feelings as Information: Informational and Motivational Functions of Affective States." In *Handbook of Motivation and Cognition*, vol. 2, edited by E. T. Higgins and R. M. Sorrentino. New York: Guildford, 527–561.

Sen, Amartya. 1985. "Well-Being, Agency and Freedom: The Dewey Lectures 1984." *Journal of Philosophy* 82 (4):169–221.

———. 1987. *On Ethics and Economics.* Wiley-Blackwell.

Senik, Claudia. 2002. "When Information Dominates Comparison. A Panel

Data Analysis Using Russian Subjective Data." DELTA Working Paper 2002–02.

———. 2004. "When Information Dominates Comparison. Learning from Russian Subjective Panel Data." *Journal of Public Economics* 88:1099–2133.

———. 2008. "Ambition and Jealousy: Income Interactions in the Old Europe Versus the New Europe and the United States." *Economica* 75 (299):495–513.

———. 2009. "Direct Evidence on Income Comparisons and their Welfare Effects." *Journal of Economic Behavior and Organization* In Press, Corrected Proof.

Shackleford, Todd K. 2001. "Self-Esteem in Marriage." *Personality and Individual Differences* 30:371–390.

Shefrin, Hersch M. and Richard H. Thaler. 1988. "The Behavioral Life-Cycle Hypothesis." *Economic Inquiry* 26:609–643.

———. 1992. "Mental Accounting, Saving, and Self-Control." In *Choice Over Time*, edited by George Loewenstein and Jon Elster. New York: Russell Sage Foundation, 287–330.

Shen, Wei and Thomas A. Louis. 1999. "Empirical Bayes Estimation via the Smoothing by Roughening Approach." *Journal of Computational and Graphical Statistics* 8 (4):800–823.

Sidgwick, Henry. 1913. *The Methods of Ethics.* London: MacMillan.

Sieff, Elaine M., Robyn M. Dawes, and George Loewenstein. 1999. "Anticipated Versus Actual Reaction to HIV Test Results." *The American journal of psychology* :297–311.

Sirgy, M. Joseph. 1998. "Materialism and Quality of Life." *Social Indicators Research* 43 (3):227–260.

Sirgy, M. Joseph, Dong-Jin Lee, Val Larsen, and Newell Wright. 1998. "Satisfaction with Material Possessions and General Well-Being: The Role of Materialism." *Journal of Consumer Satisfaction, Dissatisfaction and Complaining Behavior* 11:103–118.

Sloane, P. J. and H. Williams. 2000. "Job Satisfaction, Comparison Earnings, and Gender." *Labour* 14 (3):473–502.

Stillman, Steven. 2001. "The Response of Consumption in Russia to Economic Shocks." William Davidson Institute Working Paper No. 142.

Stram, Daniel O. and Jae Won Lee. 1994. "Variance Components Testing in the Longitudinal Mixed Effects Model." *Biometrics* 50:1171–1177.

Stutzer, Alois and Bruno S. Frey. 2006. "Does Marriage Make People Happy, or Do Happy People Get Married?" *Journal of Socio-Economics* 35 (2):326–347.

Sumner, Leonard W. 1996. *Welfare, Happiness, and Ethics*. Oxford: Clarendon Press.

Tao, Huageng, Mari Palta, Brian S. Yandell, and Michael A. Newton. 1999. "An Estimation Method for the Semiparametric Mixed Effects Model." *Biometrics* 55 (1):102–110.

Tesliuc, Emil Daniel and Lilia Ovtcharova. 2006. "Sensitivity of Poverty and Inequality Statistics to Alternative Defninitions of Household Welfare." World Bank.

Thaler, Richard H. 1980. "Toward a Positive Theory of Consumer Choice." *Journal of Economic Behavior and Organization* 1:39–60.

———. 1985. "Mental Accounting and Consumer Choice." *Marketing Science* 4 (3):199–214.

———. 1990. "Anomalies: Saving, Fungibility, and Mental Accounts." *Journal of Economic Perspectives* 4 (1):193–205.

———. 1999. "Mental Accounting Matters." *Journal of Behavioral Decision Making* 12:183–206.

Thaler, Richard H. and Eric J. Johnson. 1990. "Gambling with the House Money and Trying to Break Even: The Effects of Prior Outcomes on Risky Choice." *Management Science* 36 (6):643–660.

Timmer, Marcel P. and Bart van Ark. 2005. "Does Information and Communication Technology Drive EU-US Productivity Growth Differentials?" *Oxford Economic Papers* 57 (4):693–716.

Tinbergen, Jan. 1991. "On the Measurement of Welfare." *Journal of Econometrics* 50:7–13.

Train, Kenneth E. 1998. "Recreation Demand Models with Taste Differences over People." *Land economics* :230–239.

Twenge, Jean M. and W. Keith Campbell. 2002. "Self-Esteem and Socio-Economic Status: A Meta-Analytic Review." *Personality and Social Psychology Review* 6 (1):58–71.

Ubel, Peter A., George Loewenstein, and Christopher Jepson. 2005. "Disability and Sunshine: Can Hedonic Predictions Be Improved by Drawing Attention to Focusing Illusions or Emotional Adaptation?" *Journal of Experimental Psychology Applied* 11 (2):111.

Ueki, Masao and Kaoru Fueda. 2008. "Optimal Tuning Parameter Estimation in Maximum Penalized Likelihood Method." *Annals of the Institute of Statistical Mathematics* :1–26.

University of Essex, Institute for Social and Economic Research. 2007. *British Household Panel Survey; Waves 1-15, 1991-2006.* Colchester, Essex: UK Data Archive, 2nd ed.

Van Boven, Leaf. 2005. "Experientialism, Materialism and the Pursuit of Happiness." *Review of General Psychology* 9:132–142.

Van Boven, Leaf and Thomas D. Gilovich. 2003. "To Do or To Have? That is the Question." *Journal of Personality and Social Psychology* 85:1193–1202. Peak-End Problems?

Van Houwelingen, J. C. 2001. "Shrinkage and Penalized Likelihood as Methods to Improve Predictive Accuracy." *Statistica Neerlandica* 55 (1):17–34.

Van Praag, Bernard M. S. and Ada Ferrer-i Carbonell. 2004. *Happiness Quantified: A Satisfaction Calculus Approach.* Oxford: Oxford University Press.

Veblen, T. 1899. *The Theory of the Leisure Class.* London: George Allen and Unwin.

Venturini, Francesco. 2006. "How Much Does IT Consumption Matter for

Growth? Evidence from National Accounts." In *Central Issues in Contemporary Economic Theory and Policy*, edited by M. Malgarini and G. Piga. London: Palgrave-Macmillan.

Verbeek, Marno and Theo Nijman. 1992. "Testing for Selectivity Bias in Panel Data Models." *International Economic Review* 33 (3):681–703.

Verbeke, Geert and Emmanuel Lesaffre. 1996a. "Large Sample Properties of the Maximum Likelihood Estimators in Linear Mixed Models with Misspecified Random-Effects Distributions." Technical Report 1996.1, Catholic University of Leuven, Biostatistical Centre for Clinical Trials.

———. 1996b. "A Linear Mixed-Effects Model with Heterogeneity in the Random-Effects Population." *Journal of the American Statistical Association* 91 (433):217–221.

———. 1997. "The Effect of Misspecifying the Random-Effect Distribution in Linear Mixed Models for Longitudinal Data." *Computational Statistics and Data Analysis* 23:541–556.

Verweij, Pierre J. M. and Hans C. Van Houwelingen. 1994. "Penalized Likelihood in Cox Regression." *Statistics in medicine* 13 (23–24):2427–2436.

Von Böhm-Bawerk, Eugene. 1889. *Positive Theory of Capital*, *Capital and Interest*, vol. 2. South Holland, Illinois: Libertarian Press.

Watson, G. S. 1964. "Smooth Regression Analysis." *Sankhyā: The Indian Journal of Statistics, Series A* :359–372.

Weinzierl, Matthew. 2006. "Estimating a Relative Utility Function." Unpublished manuscript, Harvard University.

Wilkinson, Will. 2007. "In Pursuit of Happiness Research - Is It Reliable? What Does It Imply for Policy?" Cato Institute, Policy Analysis No. 590.

Wilson, Timothy D. and Daniel T. Gilbert. 2003. "Affective Forecasting." In *Advances in experimental social psychology*, vol. 35, edited by M.P. Zanna. San Diego, CA: Academic Press, 345–411.

———. 2005. "Affective Forecasting: Knowing What to Want." *Current Directions in Psychological Science* 14:131–134.

Wilson, Timothy D., Thalia Wheatley, Jonathan M. Meyers, Daniel T. Gilbert, and Danny Axsom. 2000. "Focalism: A Source of Durability Bias in Affective Forecasting." *Journal of Personality and Social Psychology* 78 (5):821–836.

Winkelmann, Liliana and Rainer Winkelmann. 1998. "Why Are the Unemployed so Unhappy? Evidence from Panel Data." *Economica* 65:1–15.

Wooldridge, Jeffrey M. 1995. "Selection Corrections for Panel Data Models

Under Conditional Mean Independence Assumptions." *Journal of Econometrics* 68:115–132.

———. 2002. "Inverse Probability Weighted M-Estimation for Sample Selection, Attrition, and Stratification." *Portuguese Economic Journal* 1:117–139.

Wu, Stephen. 2001. "Adapting to heart conditions: A test of the hedonic treadmill." *Journal of Health Economics* 20 (4):495–507.

Yemtsov, Ruslan. 2007. "Housing Privatization and Household Wealth in Transition." World Institute for Development Economics Research, Research Paper No. 2007/02.

Zhang, Daowen and Marie Davidian. 2001. "Linear Mixed Models with Flexible Distributions of Random Effects for Longitudinal Data." *Biometrics* 57 (3):795–802.