"Relative Concerns for Consumption at the Top": An Intertemporal Analysis for the UK *

Climent Quintana-Domeque
University of Oxford and IZA

Johannes Wohlfart
Goethe University Frankfurt

February 2016

Abstract

This paper investigates whether the consumption of rich households provides a reference point in the intertemporal consumption choices of non-rich households. Using UK household data on food consumption, we estimate the Euler equation implied by a life-cycle model incorporating relative concerns for the consumption of rich households. According to both our OLS and GMM estimates, for the population of non-rich households as a whole, there is no evidence of such relative concerns. We also examine an alternative model of relative concerns in which households over-consume when exposed to higher reference group consumption, and find correlational evidence that this may be the case for food consumed away from home. Finally, we investigate the presence of heterogeneous relative concerns (across county and household characteristics) in both models, finding evidence of relative concerns (for consumption at the top) in counties with relatively low income inequality. This mechanism seems to operate for food consumed away from home.

JEL Classification Codes: D12, D91

Keywords: keeping up with the "rich" Joneses, inequality, upward-relative concerns

*We would like to thank Tony Atkinson, Martin Browning, Ian Crawford and two anonymous referees for their helpful comments and suggestions. We also thank Enzo Cerletti, Alexis Grigorieff, Leander Heldring, Michael Koelle, Lu Liu, Salvatore Morelli and Christopher Roth for helpful discussions. Any errors contained in the paper are our own. Johannes Wohlfart acknowledges financial support from the German Academic Exchange Service (DAAD) and the German National Merit Foundation. The views contained here are those of the authors and not necessarily those of their respective institutions. Quintana-Domeque (corresponding author): University of Oxford, Department of Economics, Manor Road Building, Manor Road, Oxford OX1 3UQ, United Kingdom; climent.quintana-domeque@economics.ox.ac.uk

1 Introduction

A growing empirical literature in economics documents the importance of relative concerns in consumption choices (Ravina, 2007; Maurer and Meier, 2008; Charles et al., 2009; Kuhn et al., 2011; De Giorgi et al., 2012; Drechsel-Grau and Schmid, 2014; Alvarez-Cuadrado et al., 2015; Bertrand and Morse, forthcoming). Models of relative concerns assume that consumers care not only about their absolute level of consumption but also about how their own consumption compares to the consumption of their reference group (Quintana-Domeque and Turino, 2016). Recent evidence suggests that comparisons may be directed upwards, i.e. that consumers primarily compare themselves to those with higher income or economic status (Ferrer-i Carbonell, 2005; Senik, 2009; Card et al., 2012; Drechsel-Grau and Schmid, 2014; Bertrand and Morse, forthcoming). Among others, such mechanisms would have important implications for the effect of increasing top income shares on welfare and economic behavior. This is the first paper that examines the importance of upward-looking relative concerns in a life-cycle model.

Starting from a model of relative concerns along the lines of Galí (1994) and Maurer and Meier (2008) we derive an Euler equation that describes the consumption growth of low- and middle-income (non-rich) households that compare themselves to high-income (rich) households. The model predicts that non-rich households adjust their consumption *growth* in order to smooth their consumption profile relative to the consumption of rich households. We estimate the Euler equation on food consumption data from the British Household Panel Survey (BHPS) over the period 1997-2008. We start by constructing the reference point for a given non-rich household as the average consumption among rich households *in its county of residence*.[1] Later we use a definition that includes demographic characteristics as well.

We find no evidence for an effect of the growth in rich consumption on the consumption growth of non-rich households, at least for the population of non-rich households as a whole. Both OLS and GMM estimates of the Euler equation yield small positive coefficient estimates that are close to zero and statistically insignificant.

[1]Households are classified into rich and non-rich according to the position of the main earner in the county-level earnings distribution as estimated out of the Annual Survey of Hours and Earnings (ASHE).

These findings are robust to alternative definitions of the reference group, the presence of exogenous or contextual effects of reference group characteristics in the sense of Manski (1993), and the interval censoring nature of our consumption data.

We also examine whether the data are consistent with alternative models of upward-looking relative concerns in which exposure to higher reference group consumption can induce households to over-consume.[2] Instead of adjusting their consumption *growth* as predicted by the life-cycle model, non-rich households may adjust their consumption *levels*. We find correlational evidence for this mechanism for food consumed away from home.

We then investigate whether there are heterogeneous relative concerns (across county and household characteristics) in both the life-cycle model and the alternative model. In particular, for the life-cycle model, we find some evidence for upward-looking relative concerns in counties with relatively low income inequality, consistent with the idea that non-rich households in low inequality areas are more likely to compare themselves to their rich co-residents. We provide evidence that the effect in low inequality counties is driven by food consumed away from home, which is a more visible subcategory of consumption than food consumed at home (Heffetz, 2011). Similarly, we find that the effect for food consumed away from home in the alternative model of relative concerns is driven by households living in low inequality areas.

In using food consumption data for the study of intertemporal consumption choices we follow a long tradition in the literature (Hall and Mishkin, 1982; Zeldes, 1989; Runkle, 1991; Dynan, 2000; Maurer and Meier, 2008; Blundell et al., 2008; Etheridge, 2015). However, this approach, dictated in part by the lack of other appropriate types of consumption data, implies that our findings do not extend to other subcategories of consumption that are perhaps more susceptible to peer effects (Kapteyn et al., 1997; Bertrand and Morse, forthcoming). One should interpret our findings in this light.

Our study contributes to the literature in at least four ways:

First, we contribute to the literature on upward-looking relative concerns in

[2] Such behaviour would be consistent with myopia (Arrow and Dasgupta, 2009) or with high relative consumption being instrumental to achieving other goals such as finding a job (Frank et al., 2014) or spouse (Hopkins, 2008).

consumption, an idea that goes back to Veblen (1899) and which was first formalised as the "Relative Income Hypothesis" by Duesenberry (1949). Recently, this idea has attracted new attention, as a number of studies argued that the increase in top income shares was partially responsible for the decline in savings rates among middle-class households in the US (Rajan, 2010; Frank et al., 2014). Bertrand and Morse (forthcoming), using data from the Consumer Expenditure Survey (CEX), provide evidence that higher income and consumption among rich households have induced non-rich households to consume a larger fraction of their income. In a related study, Drechsel-Grau and Schmid (2014) document similar patterns in consumption-savings decisions of German households using data from the German Socio-Economic Panel (GSOEP). While the existing studies have focused on an effect on consumption *levels*, we extend the literature by testing whether upward-looking relative concerns operate in a life-cycle fashion for consumption *growth*. Moreover, since our data contain information on the county of residence, we are able to construct reference groups at a finer geographic granularity than previous studies. Reference groups are defined at the state level in Bertrand and Morse (forthcoming) or at the East-West level in the study on Germany by Drechsel-Grau and Schmid (2014). Finally, when we examine whether higher reference group consumption induces non-rich households to overconsume, the panel structure of our data allows us to remove the impact of time-invariant unobserved heterogeneity. Due to their cross-sectional data, Bertrand and Morse (forthcoming) are not able account for such heterogeneity.

Second, our paper is related to the literature on the estimation of Euler equations with external habits. Most prominently, Maurer and Meier (2008) exploit a social equilibrium condition to identify peer effects and estimate their model on food consumption data from the US Panel Study of Income Dynamics (PSID). Applying a definition of reference groups along demographic dimensions, they find evidence for moderate peer effects in consumption growth.[3] The present paper adopts a similar methodology to test the more refined hypothesis of relative concerns for the consumption of high-income households. Our different research question allows us to address two potential problems of these studies. First, identification of the effect of the average behaviour in a group on the behaviour of the units comprising

[3]Similar contributions are provided by Ravina (2007) and Alvarez-Cuadrado et al. (2015).

the group is complicated by the so-called reflection problem (Manski, 1993): the units comprising each group might share similar unobserved characteristics or be exposed to common shocks, giving rise to spurious correlations in economic outcomes. Second, as recently emphasised by Angrist (2014), IV regressions of own outcomes on average group outcomes excluding the individual herself (the "leave-out mean") mechanically result in a coefficient estimate of one. This is because the first stage regresses leave-out averages on the instrument used, while the reduced form regresses own outcomes on the instrument. Thus, first stage and reduced form estimates coincide, leading to a second-stage estimate of one. The problem is that everyone is someone else's peer (see also Boozer and Cacciola (2001)).[4]

In our application we clearly separate the "treated" units (non-rich households) from the units providing the "treatment" (rich households). It could be argued that the reflection problem is mitigated if "treated" units and those providing the mechanism are less similar. Moreover, Angrist (2014) shows that such separation helps to overcome the problem that first stage and reduced form coincide in IV estimations.[5] Thus, the hypothesis of upward-looking relative concerns is easier to test than the hypothesis of relative concerns for the consumption of more similar households. However, our zero finding on a life-cycle model with upward-looking relative concerns does not imply that peer effects are generally unimportant in consumption choices. In fact, some of our results point to the importance of other forms of relative concerns that are more challenging to test empirically.

Third, our study is the first to distinguish empirically between different models of relative concerns, namely a life-cycle model and a model of myopic consumers. This

[4]Angrist (2014) also highlights the problems that motivate the use of the leave-out-mean in social interaction models. Consider an OLS regression of individual outcomes on both average group outcomes and covariates. The estimated coefficient on the average group outcome is proportional to the excess of the 2SLS estimates over the OLS estimates in a regression of outcomes only on covariates (where in the 2SLS regression, covariates are instrumented by group dummies). 2SLS estimates may exceed OLS estimates due to peer effects but also due to many other reasons unrelated to peer effects. To overcome this problem, researchers usually regress individual outcomes on the leave-out-mean. However, group-level factors such as common shocks may still lead to spurious correlations unrelated to peer effects, motivating the use of IV estimation.

[5]To address any remaining concerns of spurious correlations, we specify a county-specific discount factor and use GMM estimation based on the orthogonality conditions of the life-cycle model. Specifically, we instrument the growth in rich consumption with lagged variables, exploiting the fact that under rational expectations current shocks will be uncorrelated with all information that was available in the prior period.

distinction is of key importance, as over-consumption and financial distress can arise as a consequence of higher reference group consumption in the second model but not in the first. Although better data are needed to disentangle the two mechanisms perfectly, our study can be seen as a first step in this direction.

Finally, our paper provides an exploration of the potential implications of the increase in income inequality that has occurred in English-speaking countries over the past decades. With relative concerns for the consumption of rich households, non-rich households derive a utility loss from observing increasing consumption levels among their rich peers. Moreover, upward-looking relative concerns imply that income inequality contributes to shaping the consumption behaviour of non-rich households. This links inequality to a range of other policy-relevant topics such as debt behaviour or the design of optimal fiscal and monetary policy (Airaudo and Bossi, 2014). The UK provides a fertile ground for studying the consequences of higher top income shares. Figure 1 displays the evolution of average incomes in the UK from 1980 until 2011. The annual percentage growth rate of average incomes over this period was 1.04% for the top 5%, 0.83% for the top 10% and 0.51% for the bottom 90%. This development has made the UK the most inegalitarian among European countries (Piketty, 2014, p.325).[6]

[**Figure 1 about here**]

The remainder of this paper is structured as follows. Section 2 presents our empirical strategy based on a life-cycle model of upward-looking relative concerns. Section 3 discusses the data and describes the sample of our study. The main results from the life-cycle model are reported in section 4. Section 5 contains several robustness checks. Section 6 investigates whether the empirical patterns are consistent with alternative models of relative concerns. Section 7 examines the presence of heterogeneous relative concerns in both life-cycle and alternative models. Finally, section 8 offers a discussion of the findings of the paper.

[6]The increasing dispersion of real incomes was largely driven by an increase in wage inequality, amongst others originating from higher pay in the financial sector (Blundell and Etheridge, 2010; Bell and Van Reenen, 2013).

2 Empirical strategy based on a life-cycle model of upward-looking relative concerns

2.1 Keeping up with the "rich" Joneses

There are several reasons why middle- and low-income households could have relative concerns for the consumption of those with higher incomes. Perhaps most importantly, consumers may want to emulate the spending habits of the upper classes to signal status, an idea that goes back at least to Veblen (1899). Individuals may either directly derive utility from a high social standing or imitate the spending of the rich to achieve other goals, such as success in the competition for marriage partners or in the labour market. Empirical evidence from the happiness literature supports the idea that comparisons are mostly upward: Having low relative standing negatively affects satisfaction, but individuals do not derive much satisfaction from a high relative position (Ferrer-i Carbonell, 2005; Senik, 2009; Card et al., 2012).

The idea that relative concerns for the consumption of rich households influence the intertemporal consumption choices of non-rich households can be formalised in a life-cycle model along the lines of Galí (1994) and Maurer and Meier (2008). At the beginning of each period, non-rich household i solves the following optimisation program:

$$\max_{C_{i,t+j},A_{i,t+j+1}} U_t = E_t \sum_{j=0}^{\infty} \beta^j u\left(C_{i,t+j}, GEOM\left[C^R_{i,t+j}\right], \psi_{i,t+j}\right)$$

$$\text{s.t.} \qquad A_{i,t+j+1} = (1 + r_{i,t+j})A_{i,t+j} + Y_{i,t+j} - C_{i,t+j}, \tag{1}$$

where C is overall non-durable consumption, A denotes the real (net) wealth or (net) assets, Y is real household income, β is the discount factor, U_t is the discounted stream of expected future utility at time t, u is the intra-period sub-utility function, $r_{i,t+j}$ is the real ex-post interest rate that the household i faces, which is unknown at the beginning of period $t + j$, and $\psi_{i,t+j}$ is a vector of preference shifters, such as age and the number of children in the household. Finally, $GEOM\left[C^R_{i,t+j}\right]$ denotes the *geometric mean* of the consumption of rich households in household i's reference group at time $t + j$.

Under certain assumptions on the utility function (see the online appendix, section A.1), observed consumption growth for household i at time t can be written as:

$$\Delta \ln C_{i,t} = \frac{1}{\rho} \ln \beta + \frac{1}{\rho} \ln(1 + r_{i,t}) + \frac{1}{\rho} \Delta \psi_{i,t} + \gamma \Delta ARITM \left[\ln C^R_{i,t} \right] + \epsilon_{i,t}, \tag{2}$$

where $\quad E_{t-1} \left[\epsilon_{i,t} \right] = 0.$

The parameter γ can be interpreted as the *strength of relative concerns for the consumption of rich households* and its estimation is the main focus of this paper. Estimating equation (2) provides a direct test against a standard life-cycle model with constant relative risk aversion utility and *without* relative concerns (as in Attanasio et al. (1999) or Alan et al. (2009)), which is nested for the case $\gamma = 0$. $\epsilon_{i,t}$ denotes the forecast error that reflects innovations to permanent income. Under rational expectations the forecast error will be uncorrelated with all available information at time $t - 1$. This insight will prove useful when addressing the reflection problem that arises from the fact that rich and non-rich households living in the same county may be subject to common shocks or share similar unobservable characteristics.

The life-cycle model predicts that non-rich households adjust their consumption *growth* in order to smooth their consumption relative to the consumption of rich households. Importantly, this does not imply that non-rich households engage in overconsumption in consumption *levels*. Since households in the model are forward-looking, they recognize that any attempt to come closer to the consumption levels of their rich peers would result in lower relative consumption in the future due to the intertemporal budget constraint. Instead, non-rich households aim to keep the distance to the consumption levels of rich households constant. One implication of this is that shocks to the expected growth rate of rich consumption from the current to the next period induces non-rich households to *reduce* their current consumption levels such as to obtain a similar increase in consumption from the current to the next period.

The economic mechanism of interest assumes that non-rich households form an estimate of the likely growth path of rich consumption from the current period onward. Non-rich households may extrapolate such an estimate from past changes in rich consumption. Such changes could either be observed directly or proxied through factors such as the local supply of goods or investment opportunities influencing the interest rate rich households are facing.

2.2 Econometric specification

Preference shifters. For empirical fit it is necessary to let the growth rate of consumption depend on changes in preference shifters that are likely to affect the marginal utility of consumption (Attanasio, 1999). Therefore, we specify the vector of preference shifters according to:

$$\begin{aligned}\psi_{i,t} = {} & \alpha_{0i} + \alpha_1 Age_{i,t} + \alpha_2 Age_{i,t}^2 + \alpha_3 HeadEmployed_{i,t} \\ & + \alpha_4 SpouseEmployed_{i,t} + \alpha_5 nadults_{i,t} + \alpha_6 nchildren0\_2_{i,t} \\ & + \alpha_7 nchildren3\_4_{i,t} + \alpha_8 nchildren5\_11_{i,t} + \alpha_9 nchildren12\_15_{i,t} \\ & + \alpha_{10} nchildren16\_18_{i,t} + z_{i,t},\end{aligned}$$

where Age is the age of the household head; $HeadEmployed$ and $SpouseEmployed$ are dummies taking a value of 1 if the head is in employment or if there is an employed spouse in the household; $nadults$ is the number of adult household members; $nchildren0\_2$ refers to the number of children aged 0-2, and similarly for the remaining variables. The employment status dummies address the finding that the empirical fit of Euler equations improves much if one conditions on changes in employment status. Potential reasons for this include non-separabilities in the utility function, non-durable consumption goods being used as inputs to market production (Aguiar and Hurst, 2013) or the substitution of home production for market goods (Aguiar and Hurst, 2005, 2007). In addition, $z_{i,t}$ captures the influence of unobserved preference shifters. Note that this specification takes account of unobserved heterogeneity in consumption levels through the household-specific constant α_{0i}, which is differenced out in the estimating equation. Also, since we take consumption growth as referring to a yearly change for all observations, the change in the age of the household head will be one year across observations. Consequently, the coefficient α_1 will be buried in the intercept. Moreover, we include year dummies to capture both the common impact of macro shocks and, since we assume common interest rates, $r_{i,t} = r_t$, the influence of the term $\frac{1}{\rho}\ln(1 + r_{i,t})$.

Choice of reference groups. As indicated by Manski (1993), inference in social interaction models is not possible unless the researcher has prior information on the composition of reference groups. Lacking direct survey questions on the composition of someone's social circle, there are two approaches commonly followed in the literature. The first one postulates that individuals compare themselves to other individuals living in the same *geographical* area (Blanchflower and Oswald, 2004; Luttmer, 2005; Ravina, 2007; Alvarez-Cuadrado et al., 2015). The second approach assumes that individuals compare themselves primarily to other individuals sharing certain *demographic* characteristics such as age, gender, education or profession (Alessie and Kapteyn, 1991; Maurer and Meier, 2008; Lewbel et al., 2013). In our main analysis, we follow the first approach and use a geographical definition of reference groups. However, instead of assuming that a consumer's reference group comprises all other households living in the same area, we examine whether the consumption growth of "rich" households can explain the consumption growth of "non-rich" households within a given county. Thus, we construct $\Delta\bar{c}^R_{i,t}$ as the change in mean log consumption of the rich households living in non-rich household i's county of residence. Moreover, in some specifications we examine the role of upward-looking comparisons within more narrowly defined cells that are based on both place of residence and demographical characteristics.

The reflection problem. Our empirical specification is subject to the reflection problem (Manski, 1993). Accordingly, there are three possible sources of correlation between *own* and *reference group* outcomes. First, there might be direct effects of peers' outcomes on own outcomes *(endogenous effects)*. In our context, this refers to the genuine influence of the growth in average rich consumption on own consumption growth: relative concerns for consumption at the top. Second, there could be direct effects of reference group characteristics on own outcomes, while these characteristics also influence reference group outcomes *(exogenous or contextual effects)*. This situation would apply if, say, age and family composition of rich households influence their own consumption, while at the same time exerting a direct influence on the consumption of non-rich households.[7] In this case, reference group characteristics would be part of the preference shifters $z_{i,t}$. Third, households living in the same area face the same institutional environment, so that they may be subject to the same shocks and share similar individual characteristics *(correlated effects).* In our case, correlated effects may arise because households in the same county are exposed to the same county-specific business cycles, the same local prices or the same advertisement. These effects are likely to lead to a positive correlation of the independent variable with the error term. Also, positive sorting of households with similar unobserved characteristics such as similar discount rates into the same county might give rise to correlated effects.

Regarding the possibility of exogenous or contextual effects, we check the robustness of our findings to controlling for changes in averages of demographic characteristics of rich households. Moreover, we include county dummies to account for the impact of time-invariant correlated effects on consumption growth. Since this formulation is equivalent to specifying a county-specific discount factor, this allows us to control for unobserved heterogeneity in discount rates across counties. However, county fixed effects are not able to capture time-varying correlated effects such as county-specific price trends. For this reason, we include the county-level employment rate as a control for county-specific business cycles.[8]

Measurement error. Both own consumption growth and the growth of average rich consumption are noisily measured. Following the consumption literature, we assume that the measurement error in the food consumption data we use is multiplicative in levels and additive in logs (Dynan, 2000; Alan et al., 2009). Taking averages over all rich households in each county and year should remove the individual measurement error in rich consumption. However, the sampling variation associated with estimating average rich consumption out of a small number of observations

[7] Depending on their family size and age structure, rich households might demand services such as nannying from non-rich households.

[8] Note that any remaining spurious correlation due to correlated effects would bias our results *towards* finding an effect. Given that we do not find a significant positive effect in most of our specifications, if anything this makes our results stronger.

adds an additional layer of measurement error. Assuming that this measurement error also takes an approximately multiplicative form, we have:

$$\hat{c}_{i,t} = c_{i,t} + u_{i,t},$$
$$\hat{\bar{c}}^R_{i,t} = \bar{c}^R_{i,t} + v_{i,t},$$

where $\hat{c}_{i,t}$ and $\hat{\bar{c}}^R_{i,t}$ are *observed* own log consumption (of a non-rich household) and *observed* average log consumption among rich households in the household's reference group, respectively. Consequently, the error term in the estimating equation will take the form:

$$e_{i,t} = \epsilon_{i,t} + \Delta u_{i,t} - \gamma \Delta v_{i,t} + w_{i,t} + \frac{1}{\rho}\Delta z_{i,t},$$

where $\epsilon_{i,t}$ again denotes the forecast error, $w_{i,t}$ denotes the non-log-linear part of consumption growth, and $\Delta z_{i,t}$ captures the change in the influence of unobserved preference shifters. Under the assumptions of classical measurement error, mismeasurements induce a serial correlation with an MA(1)-structure in the error term. Measurement error in the dependent variable will reduce the precision of our estimates but will not give rise to inconsistency. However, measurement error in the independent variable, the growth of average consumption among rich households, will lead to a negative correlation between observed values and the error term. This will drive the OLS estimate of γ to zero.

Estimating equation. Taking these considerations into account and letting lower case letters denote logs, our estimating equation for the Euler equation is given by:

$$\Delta\hat{c}_{i,t} = \kappa + \gamma\Delta\bar{\hat{c}}^{R}_{i,t} + \frac{1}{\rho}\alpha_2\Delta Age^2_{i,t} + \frac{1}{\rho}\alpha_3\Delta HeadEmployed_{i,t} + \frac{1}{\rho}\alpha_4\Delta SpouseEmployed_{i,t}$$
$$+ \frac{1}{\rho}\alpha_5\Delta nadults_{i,t} + \frac{1}{\rho}\alpha_6\Delta nchildren0\_2_{i,t} + \frac{1}{\rho}\alpha_7\Delta nchildren3\_4_{i,t}$$
$$+ \frac{1}{\rho}\alpha_8\Delta nchildren5\_11_{i,t} + \frac{1}{\rho}\alpha_9\Delta nchildren12\_15_{i,t} + \frac{1}{\rho}\alpha_{10}\Delta nchildren16\_18_{i,t}$$
$$+ \alpha_{11}EmploymentRate_{i,t} + year_t + county_i + e_{i,t}. \tag{3}$$

The estimation of this equation by OLS is likely to produce a bias whose direction is unknown. While attenuation bias will drive the estimate of γ to zero, any remaining spurious correlation between own consumption and rich consumption will lead to an upward bias of the parameter estimate. Moreover, the changes in employment status of head and spouse are jointly determined with consumption growth and will therefore introduce additional bias.

GMM estimation. In order to obtain consistent estimates in the presence of these problems, we estimate the Euler equation by two-step GMM (Hansen, 1982).[9] We treat the growth in average log consumption of rich households as well as the changes in employment status as endogenous. Valid instruments must be uncorrelated with the forecast error, i.e. they must be part of the information set at time $t-1$. Moreover, they must be uncorrelated with the measurement errors in the independent and the dependent variables, the non-log-linear component of consumption growth and changes in unobserved preference shifters. We treat changes in family composition such as the arrival of children as completely foreseen at time $t-1$, so we include them directly into our set of instruments. The excluded instruments used in our baseline estimations are the lag of the average log income of rich households, the lagged growth rate of the 80th percentile of the local earnings distribution as well as the lagged employment status of head and spouse.

One might be worried about the strength of the excluded instruments since theory implies that lagged income growth of the rich should not be directly related to current consumption growth of the rich. However, lagged income growth should correlate

[9]The assumption of rational expectations together with the assumption on the structure of the error term imply a set of population orthogonality conditions. Specifically, the expectation of the products of the error term and valid instruments should be equal to zero.

with unobserved determinants of consumption growth of the rich, such as marginal tax rates, investment opportunities or borrowing rates, which should be relevant for consumption growth of the rich through their impact on the interest rate. Moreover, lagged income growth could be related to changes in unobserved preference shifters that also influence consumption growth. In the empirical analysis, we will investigate the relevance of our instrumental variables.

Inference. Given that our main independent variable varies at the county level we adjust our standard errors for clustering at this level. This also takes account of the serial correlation in the error term that is induced by measurement error. Likewise, the GMM weighting matrix is adjusted to allow for clustering and arbitrary heteroskedasticity.

The independent variable is an estimate for the growth of rich consumption in the population. Therefore, conventional standard errors are subject to the "generated regressor problem" and will be biased (Pagan, 1984). Another potential problem is that the number of observations differs across counties, which could render inference based on conventional clustering of standard errors invalid (MacKinnon and Webb, 2014). To address these concerns, we show that our results are robust to using alternative p-values obtained from a wild bootstrap exercise in the online appendix.

3 Data

3.1 Data from the BHPS

Our main source of data is the British Household Panel Survey (BHPS). The BHPS is an annual survey consisting of a nationally representative sample of about 5,500 households recruited in 1991, including about 10,300 individuals from England, Scotland and Wales. Additional samples added 1,000 low-income households in 1997, 1,500 households from Scotland and Wales each in 1999, and 2,000 households from Northern Ireland in 2001. The main topics covered in the BHPS are housing, residential mobility, health behaviour, labour market behaviour, socio-economic values and income.[10]

[10]For a detailed description of the BHPS, see Jenkins (2010).

The BHPS has been shown to be a potentially valuable resource for analyses of consumption and savings (Guariglia and Rossi, 2002; Disney et al., 2010; Etheridge, 2015). The consumption items available in the BHPS are similar to the ones available in the Panel Study of Income Dynamics (PSID), which has been widely used for studying consumption behaviour in the US (Mankiw and Zeldes, 1991; Dynan, 2000; Alan et al., 2009). In line with a large part of the consumption literature, we focus on food consumption, that is the sum of food consumed at home and food consumed at restaurants.

The question on expenditure on food consumed at home is asked at the household level and reads: *"Please tell me approximately how much your household spends each week on food and groceries"*. Households are told that this includes "take-aways" eaten in, but excludes all meals eaten outside the home. In contrast, the question on food consumed away from home is asked at the individual level. Household members are asked: *"Tell me about how much you personally spend in an average month on eating out at, or buying take-away food from a restaurant, pub or cafe, including school meals or meals at work"*.[11] To construct a measure for household consumption, we scale the two categories to give yearly figures. We deflate them to 2000 prices by the CPI for food items and non-alcoholic beverages and the CPI for catering services (including restaurants, cafes and canteens), respectively. We then sum real expenditures on food consumed at restaurants across household members and add household consumption of food at home to obtain the households's total food consumption.

The expenditures on each of the two categories are measured in 13 bands from £0 - £160.[12] Following Guariglia and Rossi (2002) and Etheridge (2015), we assign midpoints to all categories but the highest one, to which we assign £180. Since the width of the categories is increasing in the amount spent, this can be expected to induce additional measurement error that is multiplicative in levels and additive in logs.

[11]Since "take-aways" eaten in are included in both questions, there might be a tiny overlap between the two categories. However, since this should concern only a small fraction of overall food expenditure, we do not expect this overlap to induce a large bias in our estimates.

[12]The categories are £0, £0 - £9, £10 - £19, £20 - £29, £30 - £39, £40 - £49, £50 - £59, £60 - £79, £80 - £99, £100 - £119, £120 - £139, £140 - £159, £160 and over. There are no zeros in food consumed at home.

In our context, one might be worried that the bands of the consumption variables are too wide to detect a meaningful number of observations with non-zero consumption growth. In the Euler equation (3) that uses growth in total food expenditure as dependent variable this could lead to an underestimation of the effect of interest. To investigate such a concern, Table A1 in the online appendix shows a transition matrix for switches between bands of the amount spent on food consumed at home in our study sample. The table shows that there is reasonable dispersion around the main diagonal. Table A2 in the online appendix reports the number of observations with no change in expenditure separately for each year and for food consumed at home, food consumed away from home and total food. Reassuringly, the fraction of observations with no change in total food is only around 5 percent for each year. Therefore we do not expect the interval censoring to induce a severe downward bias of our estimates. Nonetheless, we will further investigate this issue in the robustness checks section by means of interval regression.

In using food consumption data for the study of intertemporal consumption choices we follow a long tradition in the literature (Hall and Mishkin, 1982; Zeldes, 1989; Runkle, 1991; Dynan, 2000; Maurer and Meier, 2008; Blundell et al., 2008; Etheridge, 2015). This approach, dictated in part by the lack of other appropriate types of consumption data, has some shortcomings. First, survey data on food consumption could be particularly noisy. Alan et al. (2009, p.321) estimate that 62 percent of the variation in food consumption growth in the PSID are due to noise. Consequently, hypothesis tests of parameter values using only food data may have low statistical power. We address the potential problem of measurement error by using GMM. Moreover, food consumption might behave fundamentally differently than overall non-durable consumption. Following Ravina (2007), we check that our consumption measure exhibits a hump-shape over the life-cycle and a positive relationship with income, which are both key features of overall consumption. Figure 2 shows the fitted values of local polynomial regressions of total food consumption on the age of the household head and on net total household income. Reassuringly, our data on total food consumption match the key characteristics of overall consumption well.[13]

[13]Note that estimating an Euler equation for a subcategory of consumption is only valid if one

On the other hand, food compares favorably to other consumption categories in some important aspects. First, on a technical level, estimating an Euler equation with expenditure data requires a measure of consumption with *low durability*, i.e. a measure for which consumption plausibly equals expenditure. This is likely to be the case for food. Second, as highlighted in Maurer and Meier (2008, p.460), in the visibility survey conducted by Heffetz (2011) food consumption at home and at restaurants are ranked in the (upper) middle of the distribution across goods in terms of their visibility.[14]

At any rate, the use of food consumption data is an obvious limitation of our study and one should not extrapolate our findings to other, perhaps even more visible consumption categories.

[**Figure 2 about here**]

3.2 Study sample

Classification into rich and non-rich. We use data from the Annual Survey of Hours and Earnings (ASHE) to classify the households in our BHPS sample into rich and non-rich. The ASHE collects data on earnings and working hours of a one percent representative sample of all individuals in paid employment in the UK, and the Office for National Statistics (ONS) publicly provides estimates of the deciles of workplace-based county-level earnings distributions. We use these estimates in our analysis.[15]

To classify the households in our sample, we first identify the main earner in each household.[16] In each county and year, a household is classified as rich if the

invokes additional separability assumptions. In our case, we need to assume that utility is separable between consumption of food and other consumption. This assumption would be violated if, for instance, the utility derived from eating at a nice restaurant varied with the amount spent on clothing (Dynan, 2000, p.401).

[14] Among 31 expenditure categories, food at restaurants is ranked 7th and food at home is ranked 14th in terms of visibility.

[15] In principle, one could attempt to estimate county-level income distributions out of the BHPS. However, the small number of observations in each cell as well as over- and under-sampling of particular income groups render such an exercise unreliable.

[16] Earnings are defined as current gross labour earnings of the main earner in the household from his or her main job before taxes, National Insurance contributions and other deductions. We take the variables for current labour earnings and total household income in our BHPS sample from a supplementary dataset of derived income variables provided by researchers of the University of

current gross labour earnings of the main earner in his or her first job exceed the 80th percentile of the distribution of individual labour earnings from the ASHE.[17] Arguably, the 80th percentile is a fairly generous definition of "rich" households. However, this choice represents a compromise between a narrow definition of rich households and having sufficient observations in the rich group to construct meaningful estimates of the mean "rich consumption" in a given county and year. In addition, by choosing this threshold, we can compare our findings to those in Bertrand and Morse (forthcoming).

Sample selection. We use the 12 waves of the BHPS from 1997 until 2008 in line with the availability of the variable on food consumed away from home and of the earnings distribution data from the ASHE. Since ASHE data are not publicly available for Northern Ireland, we ignore the Northern Irish sample of the BHPS. The initial unbalanced panel of all English, Scottish and Welsh observations completing a full interview consists of 72,437 household-year observations from 13,813 households and 140 counties (also including unitary authorities and metropolitan counties). We drop households whose head is retired or in full-time education, exclude outliers and classify observations into "rich" and "non-rich" as explained above. Moreover, in order to obtain meaningful estimates of the mean log consumption among "rich" households, we restrict the analysis to county-year cells with at least 10 rich observations. We estimate the mean (log) consumption of total food and the average (log) net disposable income in the rich group in each county-year cell that is part of this sample.[18]

Figure 3a shows the evolution of average consumption among rich and non-rich households over the sample period 1997-2008. Consumption levels peak in 2003 and 2004 and drop with the onset of the financial crisis. Figure 3b displays the evolution of the average 80-50 ratio and the average 90-50 ratio of county-level earnings distributions in our sample. Both numbers increase over the sample period. This indicates that not only national but also local inequality in the UK has grown.

Essex. This dataset has been widely used in analyses using the BHPS (Blundell and Etheridge, 2010; Etheridge, 2015). All income variables refer to current income, are deflated by the CPI and are scaled to give an annual figure.

[17] The 80th percentile of the county-level distribution of individual earnings from employee jobs includes all part-time and full-time jobs paid on adult rates and refers to gross pay before taxes, National Insurance contributions and other deductions.

[18] Current household net disposable income refers to the sum of labour income, investment income, pension income and benefit income of all household members after taxes, National insurance contributions and other deductions. Its data source is the BHPS supplementary dataset of derived income variables.

Thus, the focus on the county level seems appropriate to study the implications of increasing inequality in the UK. Figures 4a and 4b reproduce the figures above but with separate scales for the two curves in each panel, making similarities or differences more clearly visible.

[Figure 3 about here]

[Figure 4 about here]

For the estimation of the Euler equation, we select a sample of non-rich households in the age range 20-64 that are unlikely to be liquidity-constrained.[19] To account for liquidity-constraints, we exclude households from the low-income sample of the BHPS. The final sample is an unbalanced panel consisting of 13,490 non-rich household-year observations and 10,037 first differences belonging to 2,914 households and 50 counties. Details of the sample selection procedure are provided in the online appendix (Table A3). Table 1 reports summary statistics of the final sample. For example, we can see that the average household head is approximately 44 years old, that 31% of the household heads are women, that approximately 90% of the households are based in England and that the average number of rich observations per cell is around 22.

[Table 1 about here]

4 Results from the life-cycle model with relative concerns

OLS results. Can the growth of consumption among rich households explain the consumption growth of non-rich households? Table 2 presents the results of the estimation of the Euler equation (3). All columns include demographic controls (changes in family composition and age squared) and changes in the labour market

[19]The exclusion restriction in our GMM estimations is that conditional on the included exogenous variables, consumption growth is orthogonal to everything that is part of the information set at time $t-1$. This restriction is not valid for households facing liquidity constraints. Consequently, estimating an Euler equation for such households results in biased parameter estimates (Alan et al., 2009).

status of head and spouse as well as year and county dummies. OLS estimates are presented in columns 1 and 2. Column 1 shows that conditional on the included controls, there is no significant correlation between own consumption growth of non-rich households and the growth of the average consumption of rich households in a household's county of residence. The estimate of γ, i.e. the estimated coefficient on $\Delta\bar{\tilde{c}}_{i,t}^{R}$, is 0.014 with a standard error of 0.028. Column 2 adds the county-level employment rate (i.e., a business cycle indicator) as control. The coefficient estimate for the parameter γ remains at a small positive value and statistically insignificant.

[**Table 2 about here**]

GMM results. Columns 3 and 4 present re-estimations of columns 1 and 2 using GMM. The coefficient estimates in columns 3 and 4 are similar in magnitude to the OLS estimates in columns 1 and 2 and remain statistically insignificant. Column 4 displays our preferred specification, which controls for the county-level employment rate. The coefficient estimate is 0.012 with a standard error of 0.090. To illustrate the small economic magnitude, the estimate implies that a 1 percentage point increase in the growth of average rich consumption will increase non-rich consumption growth by approximately 0.012 percentage points.

Other findings. Consistent with previous findings in the literature, changes of the labour market status of head and spouse have significantly positive effects on consumption growth. For instance, if the household head switches from "out of employment" into employment, household consumption grows by around 11 percent. The estimated coefficient on age squared is significantly negative throughout, consistent with the familiar hump shape of consumption over the life-cycle (Attanasio and Browning, 1995). The rest of the estimated coefficients are reported in the online appendix (Table A4). The R^2 of all estimations is around 0.10, indicating that a non-negligible fraction of the variance of observed consumption growth in our sample can be explained by the included controls. Unfortunately, we cannot compare this figure to other studies estimating log-linearised Euler equations, such as Dynan (2000) or Alvarez-Cuadrado et al. (2015), since these studies do not report the R^2.

Threats to identification. We test the validity of our instruments using the Hansen J-statistic. The p-values of the J-test in columns 3 and 4 are 1.00 and 0.95, so we cannot reject the null hypothesis that our instruments are valid. We check whether our instruments are only weakly correlated with the endogenous explanatory variables. At the bottom of Table 2, we report first stage F-statistics constructed as proposed in Sanderson and Windmeijer (2013).[20] The value for the growth of average rich consumption is around 37 for specifications 3 and 4, indicating that the lag of the average log income of rich households and the lagged growth of the 80th percentile are informative instruments for this variable. Likewise, the values for the changes in employment status are outside the problematic range. The first stage regressions for the results in columns 3 to 6 are reported in the online appendix (Tables A5 and A6).

Just-identified specifications. Another way to address the concern of weak instruments is to reduce the degree of overidentification. We therefore estimate just-identified specifications which are more robust in the presence of weak instruments (Angrist and Pischke, 2009, p.209). Column 5 reports the results of a GMM estimation using only lagged average income of the rich as instrument for the growth in rich consumption. Column 6 uses only the lagged change in the 80th percentile of the local earnings distribution. In both cases, the estimate of γ remains at a small positive value and statistically insignificant.[21]

Food at home vs. food away from home. Finally, food consumed at home and food consumed away from home may exhibit different degrees of visibility (Heffetz, 2011). In particular, one would expect food consumed away from home to be more visible and therefore be subject to stronger relative concerns than food consumed at home. We therefore repeat our estimations separately for the two subcategories

[20] The Sanderson-Windmeijer F-statistic is based on the Angrist-Pischke F-statistic (Angrist and Pischke, 2009). A given endogenous explanatory variable is regressed on the first-stage predictions of all other endogenous regressors and on all exogenous variables. The residuals of this regression are regressed on the excluded instruments. The Angrist-Pischke F-statistic is then computed as the F-statistic on the hypothesis that the coefficients on the excluded instruments are jointly zero in this regression. Critical values of this statistic are unavailable, but as a rule of thumb, values greater than 10 indicate that weak instruments should not be a major concern (Angrist and Pischke, 2009, p.213). Sanderson and Windmeijer (2013) adjust this test statistic taking into account that a linear projection of the other endogenous variables is partialled out instead of the variables themselves.

[21] We checked the robustness of our findings to using a Limited Information Maximum Likelihood (LIML) estimator which is known to be less subject to finite sample bias than GMM in the presence of weak instruments (Flores-Lagunes, 2007). Throughout, the estimates of γ remained virtually unchanged.

of food consumption.[22] The results are reported in the online appendix (Tables A7 and A8). Throughout, the coefficient estimates are insignificant, indicating that the effects do not differ by the type of food consumption, at least not in the full sample.

5 Robustness checks

5.1 Alternative definitions of reference groups

A potential reason for the zero finding in the life-cycle model is that a purely geographic definition of reference groups is inadequate to capture the relative concerns that are actually important in the choices of non-rich households. Other studies of interdependent preferences often assume that households compare themselves primarily with households that share certain *demographic* characteristics (Alessie and Kapteyn, 1991; Maurer and Meier, 2008; Lewbel et al., 2013). To explore this idea, we synthesize the geographic definition of reference groups used so far with definitions along two demographic dimensions, namely *educational attainment* and *age cohort*. Both dimensions have the advantage that they are almost constant over our sample period.[23]

First, we classify households into five groups according to the highest education level achieved by the head over the sample period (less than GCSE/O-level, GCSE/O-level, A-level, HND, university degree). Within each educational group, government office region and year, we compute the average consumption of those whose current labour earnings exceed the 80th percentile of the region-level earnings distribution. We subsequently drop these "rich" observations from our sample and re-estimate the Euler equation (3).[24] Columns 1-3 in Table 3 report the results of this exercise. OLS results are shown in column 1. Column 2 shows GMM results instrumenting the change in rich consumption with lagged average rich income, while column 3

[22]The separability assumption necessary to estimate an Euler equation on a subcategory of consumption is likely violated in the cases of food at home and food away from home. However, this exercise should give an initial idea about the relevance of relative concerns that depend on the visibility of a good.

[23]Only around 2 percent of the observations in our sample are out of households for which the educational attainment of the head changes over the sample period.

[24]We conducted the sample selection in the same way as before. Among others, we again deleted cells with less than 10 rich observations.

adds the lagged change in the fraction of spouses employed in the rich group as an instrument.[25] Across specifications, the coefficient estimates are of small economic size and statistically insignificant. Second, we classify households into 8-year cohorts according to the age of the household head in 1997. The results of this exercise are shown in columns 4-6. Although the GMM results are now larger in size than previously, the coefficients are imprecisely estimated. Overall, we conclude that our main results are robust to alternative definitions of reference groups.

[**Table 3 about here**]

5.2 Exogenous or contextual effects

It is possible that exogenous or contextual effects in the sense of Manski (1993) lead to biased estimates of the parameter γ due to omitted variables. That is, characteristics of rich households that influence their consumption behaviour could exert a direct influence on the consumption behaviour of non-rich households. To account for this possibility, we add changes in the average numbers of adults and children, the change in the average age of the household head as well as the growth of average log income among rich households as controls in our specifications. We encountered problems of weak instruments in this exercise. To mitigate these problems, we only estimate just-identified specifications. We instrument the change in rich consumption with the lagged growth of the 80th percentile of the earnings distribution. As an instrument for the change in rich income, we use lagged average log total family hours worked among rich households. We treat changes in average age and family composition among rich households as exogenous. Table 4 presents the results of this exercise. Throughout, the estimate of γ remains at a small positive value and statistically insignificant, suggesting that exogenous effects are not an important confounding factor in our analysis.

[**Table 4 about here**]

[25] We could no longer use the lagged growth in 80th percentile of the region-level earnings distribution as an instrument here because this variable does not vary across educational groups or age cohorts within a given region.

5.3 Interval censoring

Although we only observe a small fraction of observations with no change in consumption bands, it could still be the case that these observations cause a downward bias in our estimates. Moreover, the interval censoring could induce non-classical measurement error in the dependent variable, giving rise to a bias whose direction is unknown (Angrist and Krueger, 1999, p.1341). To address these concerns, we reestimate the OLS specifications of the Euler equation using interval regression (Stewart, 1983), a generalisation of the Tobit model that allows the dependent variable to be measured in potentially overlapping bands of arbitrary size.[26] Table 5 shows the results of this exercise. Columns 1-2 show OLS estimates and columns 3-4 report the corresponding results of interval regressions. The coefficient estimates on $\Delta\bar{\hat{c}}_{i,t}^{R}$ increase in the latter specifications, even though they remain of small economic size and statistically insignificant. The coefficient estimates on the other independent variables in the interval regressions are very similar to their OLS counterparts. While this exercise indicates that taking midpoints biases our estimate of γ toward zero, the bias does not seem big enough to account for the lack of a meaningful effect in our main estimations. This is reassuring, and consistent with the results in Etheridge (2015), who demonstrates that taking midpoints of the banded BHPS consumption data should not give rise to economically meaningful biases.

[Table 5 about here]

5.4 Additional checks

We conduct a range of additional robustness checks that are reported in the online appendix. First, we use the top 30% as an alternative definition of rich households (Table A9). Second, we repeat our main estimations in a smaller sample that applies a narrower definition of liquidity constraints (Table A10). Third, we control for county- and region-specific trends in consumption growth (Table A11). Finally, we adjust our p-values for the fact that the independent variable is an estimate for the

[26] We calculate the minimum and maximum possible consumption growth for each observation according to the expenditure bins, setting the values to missing if expenditure in any of the two periods is in the highest category. This slightly reduces the number of observations available for this robustness check.

growth of rich consumption in the population. Given that the number of counties is relatively small (at most 50), and the number of observations differs across counties, we provide Wild bootstrap p-values (MacKinnon and Webb, 2014) in the online appendix (Table A12). All these additional adjustments are immaterial for our main estimates.

6 Alternative models of relative concerns

The analysis so far has focused on a life-cycle model in which consumers are forward-looking and have rational expectations. These consumers understand that higher relative consumption today comes at the cost of lower relative consumption in the future due to the intertemporal budget constraint (Arrow and Dasgupta, 2009). Therefore, instead of trying to *come closer* to the consumption levels of their rich peers, these households *smooth* their relative consumption, keeping the expected distance to rich households constant over time.

An alternative version of the Relative Income Hypothesis (RIH) postulates that exposure to higher consumption levels of the reference group induces households to overconsume and to potentially run into financial distress (Georgarakos et al., 2014; Frank et al., 2014; Bertrand and Morse, forthcoming). Such mechanisms require either a certain degree of myopia or that high relative consumption is instrumental to other goals such as finding a job (Frank et al., 2014) or spouse (Hopkins, 2008). This section presents an attempt to distinguish between the two models and to examine whether the data are consistent with the alternative version of the RIH.

The life-cycle model and the myopic model predict different behavioral responses to an unexpected permanent increase of reference group consumption in the current period. Life-cycle consumers do not react to such a shock if the expected growth rate of reference group consumption from the current to the next period is unchanged. Myopic consumers, by contrast, may react to such a shock. Later, the intertemporal budget constraint could lead to a reversal of their spending response. Both types of consumers may react to transitory shocks to the consumption of their reference group.

Thus, the key challenge in testing the alternative version of the RIH against the life-cycle model is to disentangle empirically permanent from transitory shocks to the consumption of rich households. Unfortunately, this is very difficult in non-experimental data. To see this, suppose that the average consumption of rich households at the county level follows the following process:

$$\begin{aligned} \bar{\tilde{c}}^R_{c,t} &= \bar{\tilde{c}}^R_{c,t-1} + \epsilon_{c,t} + \eta_{c,t} - \eta_{c,t-1} \\ \text{so} \quad \Delta\bar{\tilde{c}}^R_{c,t} &= \epsilon_{c,t} + \Delta\eta_{c,t} \end{aligned} \tag{4}$$

It is impossible to isolate the permanent shock $\epsilon_{c,t}$ from the change in the transitory shock $\Delta\eta_{it}$.

In practice, if the process describing the consumption of rich households has a high persistence, one may be willing to assume that changes in the average consumption of rich households are due to permanent shocks.[27] We therefore regress the change in average rich consumption on the lagged level and test the Null hypothesis that the coefficient on lagged rich consumption equals zero against the one-sided alternative that it is less than zero.[28] We conduct this exercise separately for total food, food at home and food away from home and cannot reject the Null hypothesis at conventional significance levels in any of these tests.[29] This lends support to an interpretation of changes in the average consumption of rich households as the result of permanent shocks.

In view of this, one can regard our earlier OLS results of no effect of a change in reference consumption on own consumption growth as first evidence against the alternative version of the RIH. However, to allow for a lagged spending response with a potential later reversal we estimate additional specifications that also include shocks to rich consumption in the previous period. The results are shown in Table

[27] The empirical literature has converged on the above specification to describe income processes (Carroll et al., 2014). According to the Permanent Income Hypothesis, only permanent changes in income should affect consumption. This further supports the interpretation of changes in rich consumption as permanent shocks.

[28] The sample consists of 352 county-year observations out of the 50 counties used in the final analysis over the period 1997-2008, excluding county-year observations with less than ten observations classified as rich. Bond et al. (2005) show that this test for a unit root behaves well in panels with large N and small T.

[29] The p-values are 0.983 for total food, 0.993 for food at home and 0.680 for food away from home.

A13 in the online appendix. In line with the baseline results, we find no significant effect for total food consumption or subcategories of food consumption in the full population.

As it is unclear whether the interpretation of changes in rich consumption as permanent shocks is fully valid, we provide additional evidence on consumption *levels*.[30] Following Bertrand and Morse (forthcoming), we regress own consumption on average rich consumption while controlling for time fixed effects, household demographics and household income. To reduce potential bias from omitted correlated effects we include household fixed effects in the estimations.[31] Moreover, in some specifications we include lagged reference consumption to allow for lagged responses or a potential spending reversal.[32]

Table 6 shows the results of this exercise. While neither total food consumption of the rich (columns 1 and 2) nor their consumption of food at home (columns 3 and 4) is statistically significantly related to the consumption of the non-rich, we obtain significantly positive coefficient estimates of current and previous year reference consumption for food consumed away from home (columns 5 and 6). These differences are consistent with food consumed away from home having a higher visibility than food consumed at home (Heffetz, 2011). We find no evidence for a reversal of the spending response due to the intertemporal budget constraint. Potential reasons include that the reversal happens at a later stage or that households reduce consumption in other categories that we do not observe.[33]

[30] Among others, an interpretation of all changes in rich consumption as shocks is problematic because lagged income variables, which we have used as instruments in the analysis of the life-cycle model, have predictive power for consumption growth. After partialling out the influence of these variables, the process exhibited less persistence, which made an interpretation of changes as permanent more problematic.

[31] Whereas the baseline estimations exploit within-county variation in consumption growth of the rich, these specifications look at within-household variation in reference consumption levels. The inclusion of household fixed effects is an advantage relative to Bertrand and Morse (forthcoming) who are only able to control for state fixed effects due to their cross-sectional data. Another advantage is the construction of reference group consumption at the finer geographical level of the county as compared to the state.

[32] Note that the number of observations differs from our previous estimations. This is because in the earlier tables the number of observations refers to consumption growth, while here we look at consumption levels. Moreover, since we want to highlight differences between food consumed at home and away from home, we focus on the smaller sample where the log of food away from home is not missing. This is also the case for the following analysis of heterogeneity.

[33] We also estimated specifications including additional lags that are omitted for brevity. We did not find evidence for a spending reversal in these estimations. However, including further lags reduced the sample size and therefore statistical power.

[Table 6 about here]

Overall, while our data are not ideal to test the myopic model against the life-cycle model, we find correlational evidence for myopic consumers with upward-looking relative concerns for food consumed away from home.

7 Heterogeneous relative concerns

In this section, we investigate the possibility that some households, depending on their own characteristics or those of their location, might be more prone than others to respond to the consumption of rich households. We explore heterogeneity across six dimensions, namely: inequality in the county, population density in the county, age of the household head, marital status of the household head, educational attainment of the household head, and income of the household head.[34]

7.1 Life-cycle model with heterogeneous relative concerns

[Table 7 about here]

Inequality. Column 1 of Table 7 exploits the idea that the strength of relative concerns might differ according to the degree of inequality in an area. A higher degree of inequality means that rich households are "further away" from non-rich households and the society in a county is more segregated. Non-rich households in high inequality areas may therefore be less likely to compare themselves to their rich co-residents. Additionally, one could argue that consumption as a signal for status will be more effective if the difference in income between the two groups is less clear-cut. Accordingly, we would expect the effect to be greater if inequality is lower. We measure inequality in a given county by the 90-50 ratio of the earnings distribution in 1997, the baseline year of our sample.[35] While this value is predetermined with respect

[34]We checked for the presence of such heterogeneous effects using both OLS and GMM. However, including interaction terms increased the number of endogenous explanatory variables, leading to weak instrument problems in many of the GMM estimations. We therefore only report OLS results.

[35]A different strategy would be to instrument the current level of inequality with its value in 1997. Since following this strategy led to weak instrument problems, we decided to instead use the baseline value directly.

to the sample period, persistent omitted factors such as local tastes or institutions that correlate with both inequality and consumption growth could lead to a bias of our estimates. The inclusion of county fixed effects should account for the influence of such factors as long as it is time-invariant. However, the reaction of non-rich consumption to changes in rich consumption could still differ systematically across counties with different levels of inequality for reasons unrelated to inequality. To address this concern we checked the robustness of our findings to using the baseline 80-50 ratio or the 90-50 ratio of the previous period instead of the baseline 90-50 ratio, and obtained qualitatively similar results. However, we view the evidence in this section as suggestive and stay away from attaching a strict causal interpretation.[36] We define the dummy variable $LowIneq$, which takes a value of 1 if the observation is from a county that was classified as a low inequality area according to a median split of our sample along the 90/50-ratio of the county-level earnings distribution in 1997, and 0 otherwise.[37] The estimated coefficient on the interaction of the growth of average rich consumption with the low inequality dummy is positive and significantly different from zero at the 5 percent level. Accordingly, in low inequality areas an increase in the growth of rich consumption by 1 percentage point leads to an increase in own consumption growth by about 0.17 percentage points. This finding suggests that households primarily compare themselves to those closer within their reach.

Population density. One mechanism behind the potential effect of rich consumption on non-rich consumption is the search for status by non-rich households (Frank, 1985). In his "Theory of the Leisure Class", Veblen (1899) stressed that such costly ways of status signalling will have a more prominent role if social interactions are more anonymous, as it is the case in cities. Moreover, recent evidence in the happiness literature suggests that city dwellers compare their incomes to the incomes of their peers more intensely than residents of more rural areas (Clark and Senik, 2010, p.584). These points imply that the effect should be increasing in the population density of an area. We again measure population density by its baseline value in 1997

[36] So far, the literature has not been successful in finding a broadly applicable instrument for inequality. One notable exception is Agarwal et al. (2015) who use lottery winnings as a source of exogenous variations in inequality at the very micro level.

[37] The 90th percentile of the earnings distribution is not available for all counties in our sample, leading to a slight reduction in sample size. The data source is the ASHE.

and define the dummy variable $LowDens$, which takes a value of 1 if the observation is from a county which was classified as a low density area according to a median split of our sample along residents per square kilometer in 1997, and 0 otherwise.[38] The estimates shown in column 2 indicate that, consistent with the ideas in Veblen (1899), relative concerns for consumption at the top are more important in areas with higher population density.

Age. Relative concerns could vary by age group. Younger individuals switch their job and their social environment more often. Therefore, they engage more often in anonymous social interactions. Also, younger individuals may feel a stronger need to signal status through consumption to set themselves apart from their peers whose average incomes are still comparatively low. These points imply that younger individuals could be more prone to emulate the spending habits of the rich. We construct the dummy variable *Old* that takes the value 1 if the household head is aged above 42 and 0 otherwise. This value is in the middle of the age range in our sample (20-64) and again roughly divides the sample by half. The results shown in column 3 give no indication for different effects across age groups.

Marital status. Competition for marriage partners might induce individuals to act as if they had relative concerns in consumption (Hopkins, 2008). In the present case, married individuals may feel less need to signal status through imitating the consumption of rich households. Column 4 examines whether the effect differs between households whose head is married and households whose head is unmarried. We define the dummy variable $Married_{-1}$, which takes a value of 1 if the household head was married in the previous year.[39] We find no empirical support for this hypothesis: The point estimate on $\Delta\bar{\tilde{c}}^R \times Married_{-1}$ is positive, albeit not statistically significant.

Education. While emulating the consumption of the rich may be one way of achieving status, another way could be the accumulation of human capital.[40] Thus, individuals with higher educational attainment may feel less need to signal status

[38]The data source is the mid-year population estimates from the ONS.

[39]The married dummy is lagged because if the hypothesized mechanism was true, marital status would be endogenous.

[40]See the recent work by Moav and Neeman (2012) based on a trade-off between conspicuous consumption and human capital as signals for unobserved income.

through consumption. In column 5, the dummy variable $LowEduc_{-1}$ takes the value 1 if the lagged educational attainment of the head is below A-level and 0 otherwise. As expected, the coefficient on the growth of rich consumption decreases in size and the coefficient on the interaction term is positive and marginally significant. This suggests that relative concerns might play a moderate role for households whose head has lower educational attainment.

Income. Frank et al. (2014) argue that households compare their consumption to the consumption of those who are located only slightly higher in the income distribution. Thus, while the consumption of the rich may be a reference point for the middle class, lower income households may instead compare their consumption to the consumption of middle-income households. We explore this possibility in column 6. The dummy $MiddleInc_{-1}$ takes a value of 1 if the lagged earnings of the main earner exceeded the median of the county-level earnings distribution and 0 otherwise. We find no evidence for the ideas in Frank et al. (2014): none of the estimated coefficients are significantly different from zero.

Adjusting for multiple testing. If we estimate equations (1)-(6) simultaneously and test the null hypothesis that the coefficients on $\Delta\bar{\tilde{c}}^R \times LowIneq$, $\Delta\bar{\tilde{c}}^R \times LowDens$, $\Delta\bar{\tilde{c}}^R \times Old$, $\Delta\bar{\tilde{c}}^R \times Married_{-1}$, $\Delta\bar{\tilde{c}}^R \times LowEduc_{-1}$, and $\Delta\bar{\tilde{c}}^R \times MiddleInc_{-1}$ are all zero, the Wald test generates a χ^2 statistic (with 6 degrees of freedom) of 13.19, with an associated p-value of 0.0126. Hence, we reject the null hypothesis at the 5% level.

Margins of adjustment: Food at home vs. food away from home. Further, we examine whether the heterogeneous effects we find are driven by food consumed at home or by food consumed away from home. In the visibility survey conducted by Heffetz (2011), food at home and food away from home rank 14th and 7th among 31 broad categories of goods, respectively. This suggests that, if anything, the effects should be slightly stronger for food away from home. Table 8 shows the results for food at home. The only significant interaction term is the one corresponding to the interaction with the low educational attainment of the household head, suggesting that the effect for this subpopulation is driven by food consumed at home. However, the p-value of the test that all the interactions in Table 8 are zero is 0.1613, so we cannot reject the null hypothesis that there are no heterogeneous effects.[41]

[Table 8 about here]

Table 9 shows the results for food away from home. Most strikingly, an increase in rich consumption growth by one percentage point leads to an increase in own consumption growth by 0.141 percentage points in low inequality areas. This effect is significant at the 5 percent level. The effect in more densely populated areas, by contrast, is insignificant, indicating that the effect detected for total food consumption for this subpopulation is driven by both food consumed at home and food consumed away from home. Here, narrowing the analysis to subcategories of food consumption may lead to lower power and insignificant parameter estimates. Finally, the p-value from testing whether the coefficients on the interaction terms are jointly zero is 0.0276 for food away from home, so we reject the null hypothesis at the 5 percent level.

[Table 9 about here]

7.2 Alternative models with heterogeneous relative concerns

In section 6 (Table 6) we provided correlational evidence for an alternative model of relative concerns in which consumers adjust their consumption *levels* instead of their consumption *growth*. We repeat our analysis of heterogeneity for the different food consumption categories using this alternative model. The results are shown in Tables A14-A16 in the online appendix. While we find no heterogeneous effects for total food or for food consumed at home, we find that the significant effect for food consumed away from home is driven by consumers in low inequality counties. This provides further evidence that if class distinctions are more vague, people compare themselves more intensely to their rich neighbors and perceive a higher return from emulating the spending habits of their rich peers.

All in all, we find some correlational evidence of heterogeneous effects in consumption growth across county and household characteristics. The evidence is strongest

[41] The coefficient on food consumed at home for the rich is unstable in sign, magnitude and statistical significance.

in the case of low inequality areas, where the effect seems to be driven by food consumed away from home, which has a higher visibility than food consumed at home. Moreover, the significant effect for levels of food consumed at home detected earlier is driven by consumers resident in low inequality counties. This suggests that upward-looking relative concerns may have some role in consumer choices in these areas, and that these relative concerns operate both in a life-cycle and in a myopic fashion.

8 Discussion

Among the population of non-rich households as a whole, we find no evidence of an effect of the change in the average consumption among rich households on the consumption growth of non-rich households. Our findings therefore differ qualitatively from the related studies by Ravina (2007) and Alvarez-Cuadrado et al. (2015), who report statistically significant estimates of γ between 0.20 and 0.30. However, while they examine relative concerns that apply to *everyone* in a household's surroundings regardless of their positions in the income distribution, we test whether "non-rich" households have relative concerns for the consumption of "rich" households. Indeed, we obtain virtually the same point estimates if we change the reference point for a given non-rich household from average rich consumption to average non-rich consumption (excluding the household itself).[42]

This divergence in findings could reflect two things. On the one hand, separating households providing the reference point (rich households) from the households forming the study sample (non-rich households) may help to overcome mechanical correlations between own and reference group outcomes in social interaction models (Angrist, 2014). On the other hand, it could be that non-rich households compare themselves primarily to others that are more within their reach than to rich households. Our finding that, in low inequality areas, an increase in the growth of rich consumption by 1 percentage point leads to an increase in own consumption growth by about 0.17 percentage points supports such an interpretation.

[42]The coefficient estimates in the GMM specifications increased to sizes between 0.23 and 0.27 with p-values between 0.03 and 0.06.

One implication of this finding is that increasing top income shares could have an ambiguous effect on welfare: middle-class households may derive negative utility from observing rising consumption levels among their rich co-residents; at the same time, however, increasing segregation of the society could render this group less relevant for consumption comparisons. It is important to stress that our results only extend to welfare effects of higher income inequality that operate through relative concerns in consumption. In fact, welfare effects from other channels may become more important for higher levels of inequality (Krueger, 2002; Deaton, 2013). For instance, greater segregation of societies per se could be seen as a negative consequence of increasing inequality.

Our paper contributes to the literature on relative concerns for the consumption of rich households among non-rich households. Whereas previous contributions study such relative concerns in reduced-form settings (Drechsel-Grau and Schmid, 2014; Bertrand and Morse, forthcoming), this is the first paper that tests this mechanism in a life-cycle model. Our findings differ qualitatively from those in Bertrand and Morse (forthcoming) for the US and those in Drechsel-Grau and Schmid (2014) for Germany, who document that higher incomes and consumption among rich households can lead to lower savings rates of those with lower incomes. These differences can respond to: (1) geographical differences (different countries) or different geographical granularity levels (states in the US, East vs. West in Germany, counties in the UK); (2) data structure/availability differences (cross-section in Bertrand and Morse, panel data here; the use of an actual consumption variable here vs. the construction of consumption as the difference between income and savings in Drechsel-Grau and Schmid (2014)).

There are certain caveats of this study that must be acknowledged. First, due to data limitations we only consider food consumption. Relative concerns may primarily apply to other goods such as jewellery, clothing and cars. Bertrand and Morse (forthcoming) examine how budget shares in the consumption of non-rich households react to increasing top income levels and only document a moderate effect for food. Similarly, Kapteyn et al. (1997) find that among six broad categories of consumption items, peer effects are the weakest for food. Therefore, one should be cautious in extrapolating our findings to other types of consumption.

Second, our consumption data are coded in intervals of varying size, leading to additional measurement error and hiding smaller changes in consumption. While the fraction of observations with zero consumption growth is small and our results are robust to using interval regression, the combination of attenuation bias and lowered precision of our estimates could prevent us from finding a potential moderate but significant effect.

Third, our results depend on an accurate specification of peer groups. We have assumed that the relevant reference group of a given non-rich household consists of all rich households living in the same county. Among the few studies reporting direct survey evidence on the relevance of different peer groups, Dahlin et al. (2014) find that individuals in the US compare their incomes more intensely to the incomes of friends, family and co-workers than to the incomes of people living in the same geographic area. Recently, data-driven approaches to identifying relevant peer groups have emerged. Among others, Graham (2015) discusses methods for the identification of peer groups in network data, and Manresa (2013) develops a method for identifying individuals that generate spillovers in panel data for a setting in which outcomes depend on own characteristics and on characteristics of other individuals in the sample. This approach does not require an ad-hoc specification of peer groups. However, such methods are generally at an early stage and they are not applicable in a setting, like the current one, where no direct data on connected individuals are available.

Finally, at a more general level, our results raise the question of how realistic the life-cycle model augmented with external habits is as a description of consumer behaviour. In particular, the model predicts that non-rich households smooth their consumption *growth* relative to the consumption growth of rich households. This implies that positive shocks to the growth rate of rich consumption from the current to the next period induce non-rich households to reduce their initial consumption levels. This appears overly demanding to the rationality and foresight of non-rich households. Instead, consumers may try to keep their consumption *levels* today close to the consumption of their reference group, even at the cost of lower relative consumption in the future. This would require either a certain degree of myopia on the part of consumers (Arrow and Dasgupta, 2009) or that emulation of the spending habits of the rich is instrumental to other goals such as finding a spouse (Hopkins, 2008) or a job (Frank et al., 2014). Such models would be in line with the evidence on savings rates in Drechsel-Grau and Schmid (2014) and Bertrand and Morse (forthcoming) or with recent findings on a positive relationship between reference group income and financial distress (Georgarakos et al., 2014). We have provided correlational evidence for such alternative mechanisms operating for food consumed away from home. Ultimately, the use of quasi-experimental data could help to distinguish empirically between these two classes of models, which we believe to be an important avenue for future research.

References

Agarwal, S., V. Mikhed, and B. Scholnick (2015): "Does Inequality Cause Financial Distress? Evidence from Lottery Winners and Neighboring Bankruptcies," *Working Paper*.

Aguiar, M., and E. Hurst (2005): "Consumption versus Expenditure," *Journal of Political Economy*, 113(5), 919–948.

——— (2007): "Life-Cycle Prices and Production," *American Economic Review*, 97(5), 1533–1559.

——— (2013): "Deconstructing Life Cycle Expenditure," *Journal of Political Economy*, 121(3), 437–492.

Airaudo, M., and L. Bossi (2014): "Trickle-Down Consumption, Monetary Policy, and Inequality," *Working Paper*.

Alan, S., O. Attanasio, and M. Browning (2009): "Estimating Euler Equations with Noisy Data: Two Exact GMM Estimators," *Journal of Applied Econometrics*, 24, 309–324.

Alessie, R., and A. Kapteyn (1991): "Habit Formation, Interdependent Preferences and Demographic Effects in the Almost Ideal Demand System," *Economic Journal*, 101(406), 404–419.

Alvarez-Cuadrado, F., J.-M. Casado, and J. Labeaga (2015): "Envy and Habits: Panel Data Estimates of Interdependent Preferences," *Oxford Bulletin of Economics and Statistics*.

Angrist, J. D. (2014): "The Perils of Peer Effects," *Labour Economics*, 30, 98–108.

Angrist, J. D., and A. B. Krueger (1999): "Empirical Strategies in Labor Economics," in *Handbook of Labor Economics*, ed. by O. C. Ashenfelter, and D. Card, vol. 3a, chap. 23, pp. 1277–1366. Elsevier Science B.V.

Angrist, J. D., and J.-S. Pischke (2009): *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.

Arrow, K. J., and P. S. Dasgupta (2009): "Conspicuous Consumption, Inconspicuous Leisure," *Economic Journal*, 119, 497–516.

Attanasio, O. (1999): "Consumption," in *Handbook of Macroeconomics*, ed. by J. Taylor, and M. Woodford, vol. 1, chap. 11, pp. 741–812. Elsevier Science B.V.

Attanasio, O., J. Banks, C. Meghir, and G. Weber (1999): "Humps and Bumps in Lifetime Consumption," *Journal of Business and Economic Statistics*, 17, 22–35.

Attanasio, O., and M. Browning (1995): "Consumption over the Life Cycle and over the Business Cycle," *American Economic Review*, 85(5), 1118–1137.

Bell, B. D., and J. Van Reenen (2013): "Extreme Wage Inequality: Pay at the Very Top," *American Economic Review: Papers and Proceedings*, 103(3), 153–57.

Bertrand, M., and A. Morse (forthcoming): "Trickle-Down Consumption," *Review of Economics and Statistics*.

Blanchflower, D., and A. Oswald (2004): "Well-being over Time in Britain and the USA," *Journal of Public Economics*, 88, 1359–1386.

Blundell, R., and B. Etheridge (2010): "Consumption, Income and Earnings Inequality in the UK," *Review of Economic Dynamics*, 13, 76–102.

Blundell, R., L. Pistaferri, and I. Preston (2008): "Consumption Inequality and Partial Insurance," *American Economic Review*, 98(5), 1887–1921.

Bond, S., C. Nauges, and F. Windmeijer (2005): "Unit Roots: Identification and Testing in Micro Panels," *Cemmap Working Paper*, 07.

Boozer, M., and S. E. Cacciola (2001): "Inside the 'Black Box' of Project STAR: Estimation of Peer Effects using Experimental Data," *Yale Economic Growth Center Discussion Paper*, (832).

Card, D., A. Mas, E. Moretti, and E. Saez (2012): "Inequality at Work: The Effect of Peer Salaries on Job Satisfaction," *American Economic Review*, 102(6), 2981–3003.

Carroll, C. D., J. Slacalek, and K. Tokuaoka (2014): "The Distribution of Wealth and the MPC: Implications of New European Data," *American Economic Review: Papers and Proceedings*, 104(5).

Charles, K., E. Hurst, and N. Roussanov (2009): "Conspicuous Consumption and Race," *Quarterly Journal of Economics*, 124(2), 425–467.

Clark, A., and C. Senik (2010): "Who Compares to Whom? The Anatomy of Income Comparisons in Europe," *Economic Journal*, 120, 573–594.

Dahlin, M. B., A. Kapteyn, and C. Tassot (2014): "Who are the Joneses?," *CESR-Schaeffer Working Paper*, (2014-004).

De Giorgi, G., A. Frederiksen, and L. Pistaferri (2012): "Consumption Network Effects," *Working Paper*.

Deaton, A. (2013): *The Great Escape: Health, Wealth, and the Origins of Inequality*. Princeton University Press.

Disney, R., J. Gathergood, and A. Henley (2010): "House Price Shocks, Negative Equity and Household Consumption in the United Kingdom," *Journal of the European Economics Association*, 8(6), 1179–1207.

Drechsel-Grau, M., and K. Schmid (2014): "Consumption-savings decisions under upward-looking comparisons," *Journal of Economic Behavior and Organization*, 106, 254–268.

Duesenberry, J. (1949): *Income, Saving, and the Theory of Consumer Behavior*. Harvard University Press.

Dynan, K. (2000): "Habit Formation in Consumer Preferences: Evidence from Panel Data," *American Economic Review*, 90(3), 391–406.

Etheridge, B. (2015): "A Test of the Household Income Process Using Consumption and Wealth Data," *European Economic Review*, 78, 129–157.

Ferrer-i Carbonell, A. (2005): "Income and Well-being: An Empirical Analysis of the Comparison Income Effect," *Journal of Public Economics*, 89, 997–1019.

Flores-Lagunes, A. (2007): "Finite Sample Evidence of IV Estimators under Weak Instruments," *Journal of Applied Econometrics*, 22, 677–94.

Frank, R. (1985): "The Demand for Unobservable and Other Nonpositional Goods," *American Economic Review*, 75(1), 101–116.

Frank, R., O. Dijk, and A. S. Levine (2014): "Expenditure Cascades," *Review of Behavioral Economics*, 1(1-2), 55–73.

Galí, J. (1994): "Keeping Up with the Joneses: Consumption Externalities, Portfolio Choice, and Asset Prices," *Journal of Money, Credit and Banking*, 26(1), 1–8.

Georgarakos, D., M. Haliassos, and G. Pasini (2014): "Household debt and social interactions," *Review of Financial Studies*, 27(5), 1404–1433.

Graham, B. S. (2015): "Methods of Identification in Social Networks," *Annual Review of Economics*, (7), 465–485.

Guariglia, A., and M. Rossi (2002): "Consumption, Habit Formation, and Precautionary Saving: Evidence from the British Household Panel Survey," *Oxford Economic Papers*, 1, 1–19.

Hall, R., and F. Mishkin (1982): "The Sensitivity of Consumption to Transitory Income," *Econometrica*, 50(2), 461–481.

Hansen, L. P. (1982): "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50(4), 1029–1054.

Heffetz, O. (2011): "A Test of Conspicuous Consumption," *The Review of Economics and Statistics*, 93(4), 1101–1117.

Hopkins, E. (2008): "Inequality, Happiness and Relative Concerns: What actually is their Relationship?," *Journal of Economic Inequality*, 6, 351–372.

Jenkins, S. P. (2010): "The British Household Panel Survey and its Income Data," *IZA Discussion Paper Series*, 5242.

Kapteyn, A., S. Van de Geer, H. Van de Stadt, and T. Wansbeek (1997): "Interdependent Preferences: An Econometric Analysis," *Journal of Applied Econometrics*, 12(6), 665–686.

Krueger, A. (2002): "Inequality, Too Much of a Good Thing," *Working Paper*.

Kuhn, P., P. Kooreman, A. Soetevent, and A. Kapteyn (2011): "The Effects of Lottery Prizes on Winners and Their Neighbors: Evidence from the Dutch Postcode Lottery," *American Economic Review*, 101, 2226–2247.

Lewbel, A., S. Norris, and K. Pendakur (2013): "Necessary Luxuries," *Working Paper*.

Luttmer, E. (2005): "Neighbors as Negatives: Relative Earnings and Well-being," *Quarterly Journal of Economics*, 120, 963–1002.

MacKinnon, J. G., and M. D. Webb (2014): "Wild bootstrap inference for wildly different cluster sizes," *Queen's Economics Department Working Paper*.

Mankiw, N., and S. Zeldes (1991): "The Consumption of Stockholders and Nonstockholders," *Journal of Financial Economics*, 29, 97–112.

Manresa, E. (2013): "Estimating the Structure of Social Interactions Using Panel Data," *Working Paper*.

Manski, C. (1993): "Identification of Endogenous Social Effects: The Reflection Problem," *Review of Economic Studies*, 60, 531–542.

Maurer, J., and A. Meier (2008): "Smooth it like the 'Joneses'? Estimating Peer-Group Effects in Intertemporal Consumption Choice," *Economic Journal*, 118, 454–476.

Moav, O., and Z. Neeman (2012): "Saving Rates and Poverty: The Role of Conspicuous Consumption and Human Capital," *Economic Journal*, 122(563), 933–956.

Pagan, A. (1984): "Econometric Issues in the Analysis of Regressions with Generated Regressors," *International Economic Review*, 25(1), 221–247.

Piketty, T. (2014): *Capital in the Twenty-First Century*. The Belknap Press of Harvard University Press.

Quintana-Domeque, C., and F. Turino (2016): "Relative Concerns on Visible Consumption: A Source of Economic Distortions," *The B.E. Journal of Theoretical Economics*, 16(1), 33–45.

Rajan, R. (2010): *Fault Lines: How Hidden Fractures Still Threaten the World Economy*. Princeton University Press.

Ravina, E. (2007): "Habit Persistence and Keeping Up with the Joneses: Evidence from Micro Data," *Working Paper*.

Runkle, D. (1991): "Liquidity Constraints and the Permanent-Income Hypothesis," *Journal of Monetary Economics*, 27, 73–98.

Sanderson, E., and F. Windmeijer (2013): "A Weak Instrument F-test in Linear IV Models with Multiple Endogenous Variables," *cemmap Working Paper*, 58(13).

Senik, C. (2009): "Direct Evidence on Income Comparisons and their Welfare Effects," *Journal of Economic Behavior and Organization*, 72, 408–424.

Stewart, M. B. (1983): "On Least Squares Estimation when the Dependent Variable is Grouped," *Review of Economic Studies*, 50(3), 737–753.

Veblen, T. (1899): *The Theory of the Leisure Class*. The Modern Library.

Zeldes, S. (1989): "Consumption and Liquidity Constraints: An Empirical Investigation," *Journal of Political Economy*, 97(2), 305–346.

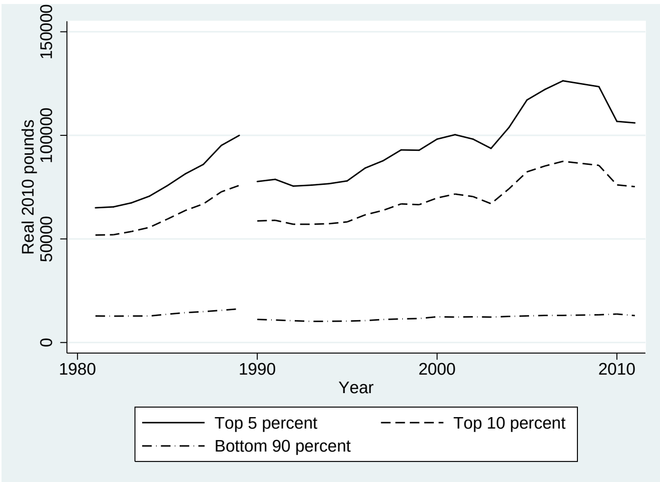

Figure 1: Average income of different income classes in the UK 1981-2011

Source: World Top Income Database. Figures until 1989 refer to average gross income of married couples and single adults. Figures from 1990 refer to the average gross income among adults (in line with the change in the tax unit).

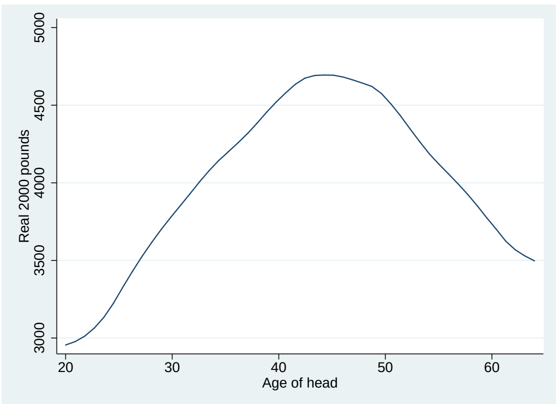

(a) Yearly food consumption on age of the head of household

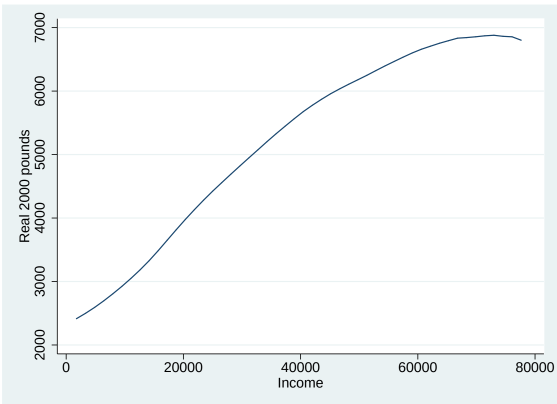

(b) Yearly food consumption on net total household income

Figure 2: Characteristics of total food consumption

Source: BHPS, 1997-2008. Kernel-weighted polynomial regressions of total food consumption on age and on net total household income using an Epanechnikov kernel. The bandwidth is 1.5 for age and 5000 for income. The sample includes all English, Scottish and Welsh households whose head is aged 20-64 that completed a full interview, having removed the top and bottom 1 percent of the distribution of net disposable household income.

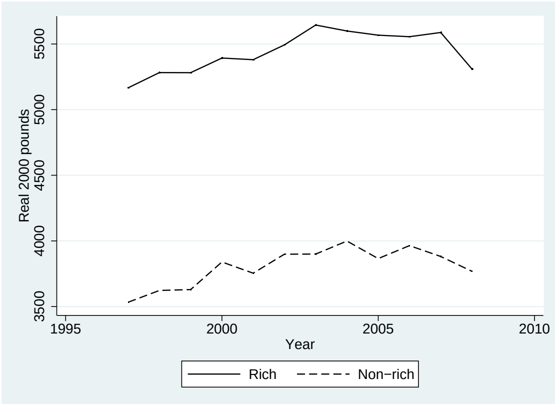

(a) Average yearly food consumption of rich and non-rich households

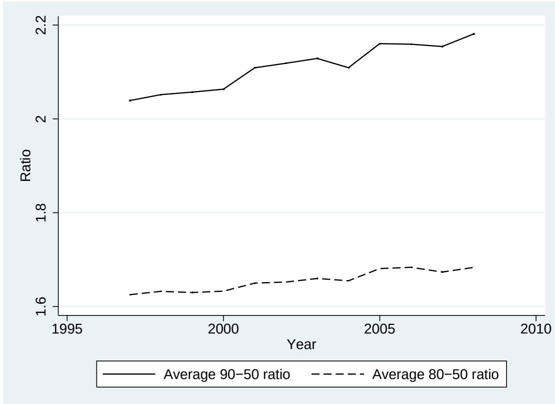

(b) Labour earnings inequality at the county level

Figure 3: Consumption and earnings inequality

Source: BHPS and ASHE, 1997-2008. Figure a) plots average consumption among rich households (above the 80th percentile) and among non-rich households (below the 80th percentile) against year. Figure b) plots the average 90/50-ratio and the average 80/50-ratio of the county-level distributions of gross earnings from paid employment against year. The sample is the study sample of county-year cells containing at least 10 households that are classified as rich. Figures are computed using the BHPS sampling weights.

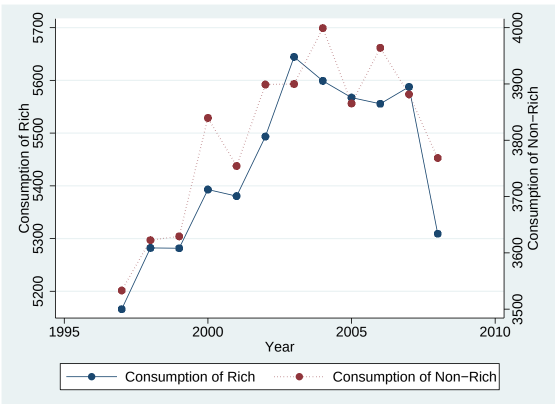

(a) Average yearly food consumption of rich and non-rich households

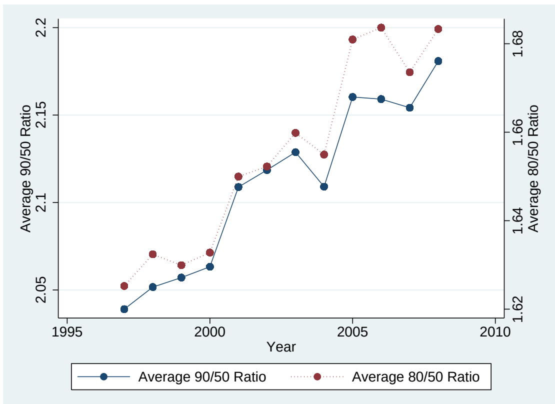

(b) Labour earnings inequality at the county level

Figure 4: Consumption and earnings inequality (different scales)

Source: BHPS and ASHE, 1997-2008. Figure a) plots average consumption among rich households (above the 80th percentile) and among non-rich households (below the 80th percentile) against year. Figure b) plots the average 90/50-ratio and the average 80/50-ratio of the county-level distributions of gross earnings from paid employment against year. The sample is the study sample of county-year cells containing at least 10 households that are classified as rich. Figures are computed using the BHPS sampling weights.

Table 1: Summary statistics

Variable	Mean	Std. dev.	Min	Max
$\hat{C}$	3,870.79	1,945.33	212.641	14,218.59
$\hat{c}$	8.121	0.565	5.360	9.562
$\Delta\hat{c}$	0.008	0.339	-2.693	2.130
Y	19,526.67	9,899.47	821.08	82,852.13
$MainEarnings$	14,141.98	8,555.02	0	39,987.07
Age	43.891	11.239	20	64
$Female$	0.306	0.461	0	1
$nadults$	1.869	0.806	1	6
$nchildren$	0.627	0.952	0	7
$Married$	0.465	0.499	0	1
$GCSE$	0.361	0.480	0	1
$ALevel$	0.227	0.419	0	1
$AboveALevel$	0.180	0.385	0	1
$PaidEmployed$	0.719	0.450	0	1
$SelfEmployed$	0.123	0.329	0	1
$SpouseEmployed$	0.427	0.495	0	1
$England$	0.895	0.307	0	1
$Scotland$	0.073	0.260	0	1
$Wales$	0.032	0.176	0	1
$\ln(80thPercentile)$	10.220	0.153	9.918	10.597
Number of rich observations per cell	21.946	10.348	10	52
$\bar{\hat{c}}^R$	8.525	0.101	8.135	8.831
$\Delta\bar{\hat{c}}^R$	0.005	0.083	-0.285	0.361
$\bar{y}^R$	10.503	0.138	10.046	10.846
N	13,490			
N first differences	10,037			
N households	2,914			
N counties	50			

Sample period: 1997-2008. All monetary variables are deflated to 2000 prices using the overall CPI and the two components for food consumption. ˆ denotes observed. Definition of variables: $\hat{C}$ is the total household food consumption (it includes both food consumed at home and at restaurants) of households classified as non-rich (below the 80th percentile); $\hat{c}$ is the log of $\hat{C}$; $\Delta\hat{c}$ is the change (first-difference) in $\hat{c}$; Y is the current household net disposable income (sum of labour income, investment income, pension income and benefit income of all household members after taxes, National insurance contributions and other deductions) of households classified as non-rich; $\bar{\hat{c}}^R$ is the average log consumption of food among households classified as rich (above the 80th percentile); $\Delta\bar{\hat{c}}^R$ is the change (first-difference) in $\bar{\hat{c}}^R$; $\bar{y}^R$ is the average log of current household net disposable income among rich households. Summary statistics are computed using the BHPS sampling weights.

Table 2: Main results

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	$\Delta\hat{c}$					
	OLS	OLS	GMM	GMM	GMM	GMM
$\Delta\bar{\hat{c}}^R$	0.014	0.012	0.007	0.012	0.009	0.023
	(0.028)	(0.027)	(0.090)	(0.090)	(0.101)	(0.187)
ΔAge^2 / 1000	-0.448	-0.447	-0.441	-0.440	-0.439	-0.440
	(0.086)***	(0.086)***	(0.087)***	(0.087)***	(0.087)***	(0.087)***
$\Delta HeadEmployed$	0.104	0.104	0.111	0.112	0.112	0.112
	(0.016)***	(0.016)***	(0.043)***	(0.043)***	(0.043)***	(0.043)***
$\Delta SpouseEmployed$	0.079	0.079	0.110	0.111	0.110	0.110
	(0.014)***	(0.014)***	(0.030)***	(0.030)***	(0.030)***	(0.030)***
Demographic controls	Yes	Yes	Yes	Yes	Yes	Yes
Business cycle indicator	No	Yes	No	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
County fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
N observations	10,037	10,037	10,037	10,037	10,037	10,037
N households	2,914	2,914	2,914	2,914	2,914	2,914
N counties	50	50	50	50	50	50
R^2	0.10	0.10	0.10	0.10	0.10	0.10
Hansen J-stat			0.00	0.00	0.00	0.00
J-stat p-value			1.00	0.95		
Sanderson-Windmeijer first stage F-stat:						
$\Delta\bar{\hat{c}}^R$			37.24	36.70	44.12	24.46
$\Delta HeadEmployed$			115.94	117.02	230.49	233.02
$\Delta SpouseEmployed$			220.64	219.82	437.30	438.88
95-percent confidence interval for γ	(-.041, .070)	(-.042, .065)	(-.169, .183)	(-.165, .189)	(-.189, .207)	(-.344, .389)

Sample period: 1997-2008. $\hat{}$ denotes observed. $\Delta\hat{c}$ is the change (first difference) in the log of the total household food consumption of households classified as non-rich (below the 80th percentile). $\Delta\bar{\hat{c}}^R$ is the change in the average log consumption of food among households classified as rich (above the 80th percentile). Demographic controls include the change in the number of adults and changes in the number of children aged 0-2, 3-4, 5-11, 12-15 and 16-18 in the household. The business cycle indicator is the county-level employment rate. In the GMM estimations 3-4 the change in rich consumption is instrumented with the lagged average log income of rich households and the lagged growth in the 80th percentile of the county-level earnings distribution, while 5 only uses lagged rich income and 6 only uses the lagged growth in the 80th percentile. Changes in labour market status are instrumented with the lagged labour market status. Standard errors are clustered at the county level and are reported in parentheses. *, ** and *** denote significance at the 10 percent, 5 percent and 1 percent level.

Table 3: Alternative definitions of reference groups

	(1)	(2)	(3)	(4)	(5)	(6)
Reference group definition:	Region and education			Region and birth cohort		
Dependent variable:	$\Delta\hat{c}$					
	OLS	GMM	GMM	OLS	GMM	GMM
$\Delta\bar{\hat{c}}^R$	0.036	0.023	0.073	0.006	0.172	0.183
	(0.032)	(0.086)	(0.071)	(0.035)	(0.197)	(0.197)
Employment status changes	Yes	Yes	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes	Yes	Yes
Business cycle indicator	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Reference group fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
N observations	13,513	13,513	13,513	13,059	13,059	13,059
N households	4,022	4,022	4,022	3,728	3,728	3,728
N reference groups	40	40	40	42	42	42
R^2	0.09	0.09	0.09	0.08	0.08	0.08
Hansen J-stat		0.00	1.24		0.00	0.11
J-stat p-value			0.27			0.74
Sanderson-Windmeijer first stage F-stat:						
$\Delta\bar{\hat{c}}^R$		23.80	18.29		9.63	9.79
$\Delta HeadEmployed$		224.87	112.44		168.57	92.90
$\Delta SpouseEmployed$		436.48	228.12		224.66	111.51

Sample period: 1997-2008. ˆ denotes observed. The study sample includes only non-rich observations (below the 80th percentile of the region-level earnings distribution) in their region-education-year cell (columns 1-3) or region-cohort-year cell (columns 4-6). $\Delta\hat{c}$ is the change (first difference) in the log of the total household food consumption of households classified as non-rich (below the 80th percentile). $\Delta\bar{\hat{c}}^R$ is the change in the average log consumption of food among households classified as rich (above the 80th percentile) in a given cell. In the GMM estimations 2 and 5 the change in rich consumption is instrumented with the lagged average log income of rich households. Columns 3 and 6 add the lagged change in the fraction of spouses employed in the rich group. Changes in labour market status are instrumented with the lagged labour market status. Standard errors are clustered at the reference group level and are reported in parentheses. *, ** and *** denote significance at the 10 percent, 5 percent and 1 percent level.

Table 4: Exogenous effects

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:			$\Delta\hat{c}$			
	OLS	OLS	OLS	GMM	GMM	GMM
$\Delta\bar{\hat{c}}^R$	0.012	0.042	0.038	0.023	0.027	0.030
	(0.028)	(0.044)	(0.048)	(0.187)	(0.224)	(0.227)
$\Delta\overline{Age}^R$		-0.000	-0.000		-0.000	-0.000
		(0.002)	(0.002)		(0.003)	(0.003)
$\Delta\overline{nadults}^R$		-0.053	-0.055		-0.050	-0.048
		(0.033)	(0.032)*		(0.059)	(0.069)
$\Delta\overline{nchildren}^R$		0.015	0.015		0.017	0.017
		(0.021)	(0.021)		(0.033)	(0.033)
$\Delta\overline{y}^R$			0.016			-0.011
			(0.046)			(0.213)
Employment status changes	Yes	Yes	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes	Yes	Yes
Business cycle indicator	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
County fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
N observations	10,037	10,037	10,037	10,037	10,037	10,037
N households	2,914	2,914	2,914	2,914	2,914	2,914
N counties	50	50	50	50	50	50
R^2	0.10	0.10	0.10	0.10	0.10	0.10
Sanderson-Windmeijer first stage F-stat:						
$\Delta\bar{\hat{c}}^R$				24.46	21.18	19.35
$\Delta\overline{y}^R$						16.36
$\Delta HeadEmployed$				233.02	234.59	235.90
$\Delta SpouseEmployed$				438.88	429.89	439.55

Sample period: 1997-2008. $\hat{}$ denotes observed. $\Delta\hat{c}$ is the change (first difference) in the log of the total food consumption of households classified as non-rich (below the 80th percentile). $\Delta\bar{\hat{c}}^R$ is the change in the average log consumption of food among households classified as rich (above the 80th percentile). In the GMM estimations the change in rich consumption is instrumented with the lagged growth in the 80th percentile of the county-level earnings distribution. In column 6 the change in rich income is instrumented with lagged average total family hours worked among rich households. Changes in labour market status are instrumented with the lagged labour market status. Standard errors are clustered at the county level and are reported in parentheses. *, ** and *** denote significance at the 10 percent, 5 percent and 1 percent level.

Table 5: Interval regressions

	(1)	(2)	(3)	(4)
Dependent variable:		$\Delta\hat{c}$		
	OLS	OLS	INTREG	INTREG
$\Delta\bar{\hat{c}}^R$	0.017	0.015	0.038	0.036
	(0.029)	(0.028)	(0.028)	(0.027)
ΔAge^2 / 1000	-0.460	-0.459	-0.469	-0.468
	(0.089)***	(0.088)***	(0.093)***	(0.093)***
$\Delta HeadEmployed$	0.105	0.105	0.104	0.104
	(0.016)***	(0.016)***	(0.015)***	(0.015)***
$\Delta SpouseEmployed$	0.080	0.080	0.076	0.076
	(0.014)***	(0.014)***	(0.014)***	(0.014)***
Demographic controls	Yes	Yes	Yes	Yes
Business cycle indicator	No	Yes	No	Yes
Year fixed effects	Yes	Yes	Yes	Yes
County fixed effects	Yes	Yes	Yes	Yes
N observations	9,826	9,826	9,826	9,826
N households	2,873	2,873	2,873	2,873
N counties	50	50	50	50
R^2	0.10	0.10		
Log pseudolikelihood			-8736.3034	-8735.9402

Sample period: 1997-2008. ˆdenotes observed. $\Delta\hat{c}$ is the change (first difference) in the log of the total household food consumption of households classified as non-rich (below the 80th percentile). $\Delta\bar{\hat{c}}^R$ is the change in the average log consumption of food among households classified as rich (above the 80th percentile). In the interval regressions 3-4 the interval for the dependent variable for each observation is bounded by the minimum and maximum possible consumption growth according to the bins. Standard errors are clustered at the county level and are reported in parentheses. *, ** and *** denote significance at the 10 percent, 5 percent and 1 percent level.

Table 6: Alternative models of relative concerns: Consumption levels

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	$\hat{c}$		$Food\hat{H}ome$		$Food\hat{A}way$	
	OLS	OLS	OLS	OLS	OLS	OLS
$\bar{\hat{c}}^R$	0.055	0.063				
	(0.039)	(0.040)				
$\bar{\hat{c}}^R_{-1}$		-0.003				
		(0.038)				
$\overline{Food\hat{H}ome}^R$			0.031	0.042		
			(0.040)	(0.041)		
$\overline{Food\hat{H}ome}^R_{-1}$				-0.037		
				(0.033)		
$\overline{Food\hat{A}way}^R$					0.092	0.114
					(0.045)**	(0.047)**
$\overline{Food\hat{A}way}^R_{-1}$						0.103
						(0.044)**
Employment status dummies	Yes	Yes	Yes	Yes	Yes	Yes
Income controls	Yes	Yes	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes	Yes	Yes
Business cycle indicator	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Household fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
N observations	12,768	12,034	12,768	12,034	12,768	12,034
N households	2,884	2,874	2,884	2,874	2,884	2,874
N counties	50	50	50	50	50	50
R^2	0.19	0.19	0.13	0.13	0.14	0.14

Sample period: 1997-2008. ˆdenotes observed. $\hat{c}$, $Food\hat{H}ome$ and $Food\hat{A}way$ are the log of total household food consumption, log food at home and log food away from home of households classified as non-rich (below the 80th percentile). $\bar{\hat{c}}^R$, $\overline{Food\hat{H}ome}^R$ and $\overline{Food\hat{A}way}^R$ are the average log consumption of total food, food at home and food away from home among households classified as rich (above the 80th percentile). Standard errors are clustered at the county level and are reported in parentheses. *, ** and *** denote significance at the 10 percent, 5 percent and 1 percent level.

Table 7: Heterogeneous relative concerns

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:			$\Delta\hat{c}$			
	OLS	OLS	OLS	OLS	OLS	OLS
$\Delta\bar{\hat{c}}^R$	0.017 (0.031)	0.128 (0.040)***	0.012 (0.057)	0.047 (0.055)	-0.017 (0.050)	0.077 (0.048)
$\Delta\bar{\hat{c}}^R \times LowIneq$	0.150 (0.060)**					
$\Delta\bar{\hat{c}}^R \times LowDens$		-0.116 (0.060)*				
$\Delta\bar{\hat{c}}^R \times Old$			0.107 (0.090)			
Old			0.002 (0.008)			
$\Delta\bar{\hat{c}}^R \times Married_{-1}$				0.034 (0.075)		
$Married_{-1}$				-0.006 (0.005)		
$\Delta\bar{\hat{c}}^R \times LowEduc_{-1}$					0.136 (0.070)*	
$LowEduc_{-1}$					-0.007 (0.003)**	
$\Delta\bar{\hat{c}}^R \times MiddleInc_{-1}$						-0.030 (0.061)
$MiddleInc_{-1}$						-0.005 (0.005)
Employment status changes	Yes	Yes	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes	Yes	Yes
Business cycle indicator	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
County fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
N observations	8,858	9,225	9,225	9,225	9,225	9,225
N households	2,672	2,791	2,791	2,791	2,791	2,791
N counties	47	50	50	50	50	50
R^2	0.10	0.10	0.10	0.10	0.10	0.10

Sample period: 1997-2008. ˆdenotes observed. $\Delta\hat{c}$ is the change (first difference) in the log of the total household food consumption of households classified as non-rich (below the 80th percentile). $\Delta\bar{\hat{c}}^R$ is the change in the average log consumption of food among households classified as rich (above the 80th percentile). $LowIneq$ and $LowDens$ are dummy variables indicating whether an observation is from a county classified as a low inequality or a low density area according to a median split of the sample along the values in 1997. These estimations do not directly control for the dummy variables because these variables do not vary within counties and we include county dummies. Old is a dummy variable indicating whether the head is aged above 42. $Married$ is a marital status dummy and $LowEduc$ takes a value of 1 if the educational attainment of the head is below A-Level. $MiddleInc$ is a dummy variable indicating whether the labour earnings of the main earner exceed the 50th percentile of the earnings distribution. The controls are as explained previously. Standard errors are clustered at the county level and are reported in parentheses. *, ** and *** denote significance at the 10 percent, 5 percent and 1 percent level.

Table 8: Heterogeneous relative concerns: Food consumed at home

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	$\Delta Food\hat{H}ome$					
	OLS	OLS	OLS	OLS	OLS	OLS
$\overline{\Delta Food\hat{H}ome}^R$	0.090 (0.038)**	0.062 (0.029)**	0.015 (0.045)	0.031 (0.053)	-0.026 (0.049)	0.097 (0.043)**
$\overline{\Delta Food\hat{H}ome}^R \times LowIneq$	-0.029 (0.061)					
$\overline{\Delta Food\hat{H}ome}^R \times LowDens$		0.006 (0.056)				
$\overline{\Delta Food\hat{H}ome}^R \times Old$			0.104 (0.076)			
Old			-0.001 (0.007)			
$\overline{\Delta Food\hat{H}ome}^R \times Married_{-1}$				0.072 (0.072)		
$Married_{-1}$				-0.003 (0.005)		
$\overline{\Delta Food\hat{H}ome}^R \times LowEduc_{-1}$					0.158 (0.071)**	
$LowEduc_{-1}$					-0.006 (0.005)	
$\overline{\Delta Food\hat{H}ome}^R \times MiddleInc_{-1}$						-0.071 (0.057)
$MiddleInc_{-1}$						0.000 (0.006)
Employment status changes	Yes	Yes	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes	Yes	Yes
Business cycle indicator	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
County fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
N observations	8,858	9,225	9,225	9,225	9,225	9,225
N households	2,672	2,791	2,791	2,791	2,791	2,791
N counties	47	50	50	50	50	50
R^2	0.05	0.05	0.05	0.05	0.05	0.05

Sample period: 1997-2008. ˆdenotes observed. $\Delta Food\hat{H}ome$ is the change (first difference) in the log of the household consumption of of food at home of households classified as non-rich (below the 80th percentile). $\overline{\Delta Food\hat{H}ome}^R$ is the change in the average log consumption of food among households classified as rich (above the 80th percentile). *LowIneq* and *LowDens* are dummy variables indicating whether an observation is from a county classified as a low inequality or a low density area according to a median split of the sample along the values in 1997. These estimations do not directly control for the dummy variables because these variables do not vary within counties and we include county dummies. *Old* is a dummy variable indicating whether the head is aged above 42. *Married* is a marital status dummy and *LowEduc* takes a value of 1 if the educational attainment of the head is below A-Level. *MiddleInc* is a dummy variable indicating whether the labour earnings of the main earner exceed the 50th percentile of the earnings distribution. The controls are as explained previously. Standard errors are clustered at the county level and are reported in parentheses. *, ** and *** denote significance at the 10 percent, 5 percent and 1 percent level.

Table 9: Heterogeneous relative concerns: Food consumed away from home

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:			$\Delta Food\hat{A}way$			
	OLS	OLS	OLS	OLS	OLS	OLS
$\overline{\Delta Food\hat{A}way}^R$	-0.142 (0.073)*	0.032 (0.091)	0.092 (0.069)	0.039 (0.070)	0.035 (0.088)	-0.027 (0.077)
$\overline{\Delta Food\hat{A}way}^R \times LowIneq$	0.283 (0.108)**					
$\overline{\Delta Food\hat{A}way}^R \times LowDens$		-0.023 (0.113)				
$\overline{\Delta Food\hat{A}way}^R \times Old$			-0.160 (0.075)**			
Old			0.003 (0.021)			
$\overline{\Delta Food\hat{A}way}^R \times Married_{-1}$				-0.048 (0.070)		
$Married_{-1}$				-0.013 (0.010)		
$\overline{\Delta Food\hat{A}way}^R \times LowEduc_{-1}$					-0.032 (0.095)	
$LowEduc_{-1}$					-0.010 (0.012)	
$\overline{\Delta Food\hat{A}way}^R \times MiddleInc_{-1}$						0.092 (0.090)
$MiddleInc_{-1}$						-0.027 (0.014)*
Employment status changes	Yes	Yes	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes	Yes	Yes
Business cycle indicator	Yes	Yes	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
County fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
N observations	8,858	9,225	9,225	9,225	9,225	9,225
N households	2,672	2,791	2,791	2,791	2,791	2,791
N counties	47	50	50	50	50	50
R^2	0.09	0.08	0.08	0.08	0.08	0.08

Sample period: 1997-2008. $\hat{}$ denotes observed. $\Delta Food\hat{A}way$ is the change (first difference) in the log of the household consumption of of food away from home of households classified as non-rich (below the 80th percentile). $\overline{\Delta Food\hat{A}way}^R$ is the change in the average log consumption of food among households classified as rich (above the 80th percentile). $LowIneq$ and $LowDens$ are dummy variables indicating whether an observation is from a county classified as a low inequality or a low density area according to a median split of the sample along the values in 1997. These estimations do not directly control for the dummy variables because these variables do not vary within counties and we include county dummies. Old is a dummy variable indicating whether the head is aged above 42. $Married$ is a marital status dummy and $LowEduc$ takes a value of 1 if the educational attainment of the head is below A-Level. $MiddleInc$ is a dummy variable indicating whether the labour earnings of the main earner exceed the 50th percentile of the earnings distribution. The controls are as explained previously. Standard errors are clustered at the county level and are reported in parentheses. *, ** and *** denote significance at the 10 percent, 5 percent and 1 percent level.