

OPTIMAL POLICY AND INCONSISTENT PREFERENCES:  
BEHAVIOURAL POLICYMAKING AND SELF-CONTROL

NICHOLAS CHESTERLEY

Doctor of Philosophy in Economics  
Nuffield College, University of Oxford  
Trinity Term 2015

Submitted for the degree of Doctor of Philosophy in Economics.



*I am Legion, for we are many*

## ABSTRACT

---

This thesis takes three different perspectives, using theoretical and experimental techniques, on time-inconsistent preferences and how the existence of multiple selves can affect both consumer behaviour and policy design. Across domains such as retirement saving, health, and educational achievement, intertemporal choice presents a challenge for both individuals and policymakers.

The first paper, ‘Choosing When to Nudge: Designing Behavioural Policy around Decision-Making Costs,’ considers how behavioural policy, which has proven increasingly popular with policymakers, affects welfare. I find that for present-biased consumers, behavioural policies help some consumers but can inefficiently discourage others from optimizing. Such policies therefore have an ambiguous effect on welfare, and similar to traditional policies, can impose equity-efficiency tradeoffs. Monopolies may increase welfare given their incentive to simplify consumer decisions instead of exploit switching costs.

The second paper, ‘Virtue and Vice with Endogenous Preferences,’ considers behaviour when preferences are affected by consumer decisions. I introduce agents whose temptation to consume in the present is affected by how much they choose to save for the future. I find that differences between agents can trap them in divergent paths of self-improvement – saving more, they value the future more, making saving optimal – or binging – consuming more makes them indifferent to future costs, making consumption optimal. At the extreme, it is frequently an optimum for a consumer to consume their entire wealth.

The final paper, ‘Bet You Can’t Eat Just One: Consumption Complementarity and ‘Self’-Control’ considers an intrapersonal game between a moderate cold self and a hot self that wants to indulge. In equilibrium, sophisticated selves best respond to each other’s behaviour: the cold self over-abstains and the hot self over-indulges to avoid inducing the other state. I test these ideas in the lab, and find that subjects on a diet who were induced to consume a piece of chocolate before the experiment indulge more in chocolate during the experiment, even when the initial indulgence was imposed by the experimenter. Eating a piece of chocolate, this suggests, can induce a period during which chocolate is more appealing.

**JEL Classification:** D01; D03; D04; D11; D14; D90; C91; O38

**Keywords:** Intertemporal Choice; Preferences; Habits; Policy Design; Self-Control; Consumption; Willpower

**Word Count:** 56,823

## ACKNOWLEDGMENTS

---

As is perhaps appropriate in a thesis built around the idea of multiple selves, I had the help and support of many people throughout this thesis (and they weren't all me).

Kevin Roberts as my supervisor has been a source of constant inspiration and support through both my MPhil and DPhil work, and his ability to understand in moments what would take me weeks to piece together was both overwhelming and extraordinarily helpful. Without him, my models would be both far less interesting and far less correct, and it has been an honour to work with him.

More recently, I have also been enormously grateful to Johannes Abeler, who kindly agreed to supervise my experimental work. His insights and support throughout the experimental process have been invaluable, as has his tolerance of my entirely novice status with regard to experiments when I began. John Jensenius at CESS was also of great help and support in running the experiment.

My college advisor, Meg Meyer, has gone out of her way to be helpful throughout my time at Oxford, helping point my attention to relevant papers and speakers as well as writing innumerable reference letters and even kindly reading through several drafts.

My fellow students at Nuffield and the Department of Economics have been an essential part of the process. The only-angry-when-people-play-Risk-suboptimally Felix Pretis, the always willing to take a break from studying for the MPhil with me (even when the break was in Mongolia and I didn't play chess nearly enough) David Milligan, my (college) daughter Daniela Scur, my Turner-accomplice and diagonal-deskmate Claudia Herresthal, and the rest of the micro theory café have all provided innumerable comments and suggestions on various parts of the thesis, and all my fellow DPhil students have helped provide sanity-preserving encouragement throughout. They have also sat through numerous presentations, as have the many participants in seminars in the UK, Europe, and Canada, all of whom gave helpful advice and comments and contributed to improving the various papers.

I have also had the privilege of living with phenomenal people during my time in Oxford: Nithum Thain, Tyler Grant, Kayli Johnson, Simon Lolliot, the irrepressible Greg, JB, Neil, Jamie, and Alex of Western Road, and of course Tamara Moellenberg (who has put up with me and with economics more than any decent person should have to), as well as the other great people I've gotten to live with, I am lucky to count you among my friends, and I hope you all live with and near me again as often as possible.

As a Clarendon Scholar and SSHRC Doctoral Fellow, both the Clarendon Fund and SSHRC have been invaluable in providing me with the resources and support to focus on my work, and I am tremendously appreciative of their help. I have also benefited from the opportunity to work with both the Institute of Fiscal Studies and the C.D. Howe Institute, which have contributed enormously to my learning.

Finally, my greatest thanks must go to my family. It is perhaps trite to say that without them, I wouldn't be where I am today, but in the case of family it is quite literally true. Without their help and guidance, I would be nowhere at all.

Thank you all.

Nicholas Chesterley

Nuffield College, Oxford

March 2015

# CONTENTS

---

1	INTRODUCTION	1
2	CHOOSING WHEN TO NUDGE: DESIGNING BEHAVIOURAL POLICY AROUND DECISION-MAKING COSTS	7
2.1	Introduction	7
2.2	Literature	11
2.3	Model	14
2.3.1	Costs and Benefits of Technology Adoption	14
2.3.2	Decision Rule	16
2.4	Policy Interventions	19
2.4.1	Decision-Making Costs	19
2.4.2	Choice Intervals	22
2.4.3	Default Options	29
2.4.4	Simultaneous Policies	37
2.5	Policy Lessons	39
2.6	Firms	42
2.6.1	Monopolies	43
2.6.2	Competitive Firms	48
2.7	Conclusions	49
3	VIRTUE AND VICE WITH ENDOGENOUS PREFERENCES	52
3.1	Introduction	52
3.2	Literature	55
3.3	Model	59
3.3.1	Fixed Preference Types	61
3.3.2	Ultrasophisticates	63
3.3.3	Uniqueness	67
3.4	Consumption Patterns	69
3.4.1	Ultrasophisticates and Fixed-Preference Agents	69
3.4.2	Ultrasophisticate Behaviour	73
3.4.3	Pseudorationality	77
3.5	Welfare and Policy Implications	78
3.6	Discussion and Conclusion	79
4	BET YOU CAN'T EAT JUST ONE: CONSUMPTION COMPLEMENTARITY AND 'SELF'-CONTROL	82
4.1	Introduction	82
4.2	Theory	85
4.2.1	Model	87
4.2.2	Two Periods, Naïve Hot Selves	90
4.2.3	Infinite Horizon, Sophisticated Hot Selves	94
4.2.4	Long-Run Effects	106
4.3	Experiment	109
4.3.1	Design	111
4.3.2	Results	113
4.3.3	Hot/Cold State Transition Mechanisms	117
4.4	Discussion and Conclusions	119

5	CONCLUSION	122
	Bibliography	125
A	CHOCOLATES	132
B	EXPERIMENTAL INSTRUCTIONS	133

## INTRODUCTION

---

How well consumers plan for the future has become an increasing concern for policy-makers. Across retirement, health, education, energy use, and other issues, consumers seem to emphasize present benefits over future ones, in ways that are not always consistent or easily explained by traditional economic models. Given the importance of how we allocate resources over time – including effort, investment, and consumption – in determining long-run outcomes, governments are frequently interested in using policy to support or change consumers' decisions. Often, the goal is not to provide more options to consumers, but instead to convince them that what they might decide in the short term is not in their long-term interest. In doing so, researchers and policymakers confront the possibility that consumers behave as if they had multiple 'selves' with non-stationary preferences: that what consumers want today may not be what they want tomorrow.

In this thesis, I use both theoretical and experimental methods to better understand individual behaviour, focusing particularly on the implications of non-stationary preferences, and apply those insights to policy design. If consumers want different things at different times, then planning for the long term may be particularly difficult, and in some cases, consumers may not be able to reach their global optimum. First, using a common form of time-inconsistent preferences – Laibson's quasi-hyperbolic discounting (1997) – I study the welfare implications of behavioural policies, which have become increasingly popular with policymakers. Then, I consider behaviour and policy under two alternative approaches to time-inconsistent preferences: when preferences are endogenous and are affected by behaviour, and when consumption induces temporary 'hot' selves that have a higher marginal utility of consumption, so each self best-responds to the behaviour of the other selves. I turn to the lab to test some of these ideas, studying how consumers behave when they have been induced to indulge in temptation by the experimenter. Each of the three chapters takes a different perspective on time-inconsistency – nudges for present-biased technology adoption, behaviour when choices affect preferences, and self-control with multiple selves – but each focuses

on understanding consumer behaviour and the welfare and policy implications of time-inconsistency.

Preferences present a challenge for economics. Typically, they are treated as the final arbiter of choice: consumers make decisions based on their preferences, and by revealed preference when we observe those choices we can derive the preferences that motivated them. There is no basis, therefore, for judging preferences as better or worse, but instead economics aims to respect individual preferences as much as possible. For interpersonal decisions, this is often sensible: privileging the preferences of one person over another is by definition a controversial idea.

Behavioural economics, however, has pointed out that individuals often behave as if they have preferences that change over time, or equivalently as if they have multiple selves that take decisions in turn. The fact that consumers do not actually always optimize based on a consistent set of self-interested preferences has a long history in economics, of course: Adam Smith wrote in *The Theory of Moral Sentiments* about the importance of preferences such as fairness (1759), as did Bentham when he introduced utility (1789). Until recently, however, most of modern economics relied on the standard assumption of rational optimization, with a few exceptions such as Herbert Simon's suggestion that consumers satisfice by trying to just do well enough, rather than maximize (1955). With the surge of interest in Kahneman and Tversky's work on prospect theory and decision-making under risk (1974; 1979; 1984), however, modifications to the standard model are becoming increasingly common, and behavioural economics is popular with both economists and policymakers.

The idea of time-inconsistent preferences was first modeled in Strotz's seminal work (1955-1956), further developed by Ainslie as hyperbolic discounting (1975) and, later, motivated the quasi-hyperbolic discounting model by Laibson (1997), with work by many others along the way, such as Phelps and Pollak (1968) and Peleg and Yaari (1973).<sup>1</sup> These models suggest that consumers place a particularly high weight on the present as opposed to the future, but unlike standard exponential discounting models, weight future periods against each other without the present bias factor. As a result, today vs. tomorrow is a very different choice from a year from now vs. a year plus a day. This

<sup>1</sup> Daniel Akst reports that Strotz applied his work to his own life: as President of Northwestern University, he had his academic employer pay him in 12 monthly installments to avoid having to save for the summer (Akst, 2011).

contrasts directly with Samuelson's 1937 discounted utility model, which despite his strong reservations was adopted almost immediately after publication as the standard model in economics, and is generally still used today. In discounted utility, consumers care more about the present than the future, but do so at a stationary rate of discounting, one that does not change over time (Samuelson, 1937). For Strotz and Laibson, however, preferences are non-stationary: each time 'today' arrives, it gains a new importance in the eyes of the consumer.<sup>2</sup>

With this in mind, recent research has considered issues of self-control as a game between multiple selves, who may want very different things but frequently have access to the same technologies for commitment and opportunities for consumption. In 1978, Thomas Schelling argued in the AER that two selves were in a "continual contest for control," and that most of us have ways of trying to motivate the other self, from putting the alarm clock further from the bed to his own attempts to toughen himself to cold temperatures by going to bed without a blanket, only to wake up in the middle of the night too cold to find the blanket and cursing his earlier self (Schelling, 1978). Thaler and Shefrin make a similar argument, and discuss the implications for pensions, the timing of income flows, and individual discount rates (Thaler and Shefrin, 1981). A self can use methods such as commitment devices, delay tactics, or reaching a deal between the selves in order to maximize utility.

Both economists and psychologists generally accept the idea that preferences are non-stationary, and that each set of preferences can be represented by a different 'self.' How preferences and selves vary is not yet well understood, however, and introduces a challenge to analysis: if multiple selves exist, which set of preferences matter in welfare analysis, and who is doing the optimizing? As a result, fixed preferences remains the typical assumption in economic work, and indeed in many situations selves may have similar preferences. Over time, however, selves appear to differ in systematic and important ways. What does this mean for consumer behaviour, in both the short and long run? And is there an opportunity for policy to improve outcomes? This thesis takes several different perspectives in an effort to illuminate some possible responses to those questions.

---

<sup>2</sup> A particularly useful review of the literature can be found in Frederick et al. (2002).

Chapter 2 considers a relatively straightforward case, when consumers are present biased and face decision costs, making them reluctant to adopt energy efficient technology that has long-term benefits. Because there is an externality, there is a relatively clear case for policy intervention, and behavioural interventions – nudges – have proved popular with policymakers as a way of changing behaviour. There is reason, however, to be worried that nudges have proved too popular. I introduce a model of decision costs with consumers who are present biased, and find that nudges may discourage consumers from optimizing and convince them to take the default even when it makes them and society worse off. In some cases, behavioural policies may create welfare distortions and equity-efficiency trade-offs. Avoiding this problem is possible, but imposes restrictive information requirements on the policymaker, which may often not be met. Nudges are potentially a useful policy tool, but require careful cost-benefit analysis to be used well, and in some cases even well-intentioned nudges may actually reduce societal welfare, rather than increase it.

Chapter 2 discusses the limitations of policies that are already popular. The other two chapters, however, focus directly on understanding consumer behaviour, potentially leading to new policy prescriptions. Chapter 3 explores a particular form of non-stationary preferences: drawing on the psychology literature, it considers the idea of endogenous preferences, so that behaviour affects preferences instead of just being determined by them. I consider a model of quasi-hyperbolic discounting in which present bias, the temptation to consume today, is affected by how much self-control consumers exert. I find that differences between agents can trap them in divergent paths of virtue or vice due to the existence of two types of optima: *self-improvement optima*, from the desire to have ‘better’ preferences, and *binge optima*, from the increased indifference to future costs induced by myopic decisions. Agents may also have multiple locally optimal consumption paths, and preference endogeneity can explain why some agents attempt to resist temptation, while others indulge in it.

Chapter 3 assumes that each self is nested in the period before: each self accounts for their future behaviour, but their choice is unaffected by their past behaviour except through their wealth. There is therefore no intrapersonal game between selves. In my final chapter, I relax this assumption and allow selves to each best respond to the behaviour of the other self, so that each must account for the expected behaviour of the

other when they make their equilibrium decision. Consuming in a cold state induces a temporary hot state of complementarity in consumption, making further consumption more appealing. Intuitively, a consumer might be welfare optimizing to eat a square of chocolate, have a single drink, or take a five minute break from work, but if doing so induces a hot state that prefers further indulgence, then they may be unable to choose their optimum consumption path and prefer instead to avoid indulgence in the first place.

I find that naïve cold selves who do not anticipate this effect overconsume relative to their optimum: the glutton who is repeatedly surprised by how much they eat, for example. Sophisticates restrict their decisions because they anticipate the costs that result from entering a different state. While in the cold state, the consumer therefore consumes less than is otherwise optimal in order to avoid entering a hot state: once they enter a hot state, they binge and consume more than is otherwise optimal to avoid returning to the cold state. Multiple equilibria may exist, in which case the optimism or pessimism of each self about the other self's behaviour can determine which equilibrium results. Behaviour also depends on whether selves are sophisticated and able to best respond to the behaviour of the other self, or naïve.

To test these ideas, I turn to the lab. Experimental work provides an important source of data for predictions about consumer behaviour, and can help refine and improve theoretical analyses. Here, I conduct a lab experiment to test for the existence of a short-term causal effect of consumption on how much the consumer indulges. Drawing on the theory of hot and cold states introduced earlier, I test whether subjects induced to consume a piece of chocolate before the experiment indulge more in chocolate during the experiment. I find some evidence of this effect for subjects on a diet, who could be expected to find chocolate more tempting. This holds even when the initial indulgence is externally imposed, ruling out explanations that rely on freedom of choice in the initial indulgence. Eating a piece of chocolate, in other words, can induce a period during which further chocolate consumption is more appealing. This could potentially result from a fixed cost to indulging in chocolate – getting sticky hands, or a loss of self-image, perhaps – or from a temporary shift in preferences, all of which could provide some explanation of the underlying mechanism of a hot state.

All three chapters aim to deepen and extend our understanding of consumer behaviour. As policymakers and educators, there is the potential for us to do better helping consumers plan for the future. In some situations, nudges may prove a useful tool to simplify the consumer decision and make it easier for them to optimize. In others, consumers are vulnerable to negative outcomes based on the non-stationarity of their preferences, and so increasing their sophistication and providing them with tools to commit to certain behaviours may be valuable. I begin by studying the nudge-based policies that are increasingly popular with policymakers in the UK, the US, and elsewhere, before turning in later chapters to the possibility for novel policies based on models of time-inconsistency in intertemporal choice.

## CHOOSING WHEN TO NUDGE: DESIGNING BEHAVIOURAL POLICY AROUND DECISION-MAKING COSTS

---

*“It is now a well-established proposition that both self-control and cognitive effort are forms of mental work.”*

*-Daniel Kahneman (2011)*

### 2.1 INTRODUCTION

A common friction for consumer choice is the presence of decision costs, either material or cognitive. Taking a decision requires effort and may also impose search costs to find and analyze information. Such costs can encourage employees to save at the default pension rate if faced with too many possible plans (Huberman, Iyengar and Jang, 2004); doctors to prescribe the standard treatment when too many alternatives are available (Redelmeier and Shafir, 1995); and consumers to refuse to insulate their lofts until the government offers a joint loft insulation and cleaning package, even when it saves them no money (DECC, 2013). This problem has only become more acute today, corresponding to the rapid multiplication of options and information facing consumers.

As austerity and cynicism about the power of government have struck, governments have turned to behavioural economics as an inexpensive way of helping consumers navigate complex decisions. Drawing on Kahneman’s seminal work in the field, recent work by the British Behavioural Insights Team, the Danish Nudging Network, the White House Social and Behavioural Sciences Team and others has focused on practical applications, including modifying defaults, adding or taking away options, or simply changing the framing of the decision (Kahneman and Tversky, 1984; Thaler and Sunstein, 2008). Unfortunately, theory has not kept pace with practice, and the justification for

and welfare implications of these policies are not always clear, particularly given their reliance on deviations from the standard economic model of preferences.

This paper analyzes how the existence of decision costs and present bias can both explain the effectiveness of behavioural policies as well as allow for the study of their limitations. The model provides lessons for a wide range of settings, from personal health to retirement saving: to illustrate the argument, I focus on consumers considering adopting energy-saving technology, but who, because they are present biased, overestimate the decision costs they will incur if they attempt to find their optimal technology rather than take the default. This paper provides a systematic model of the costs and benefits of behavioural policy, providing lessons for policymakers on both when and how such policies should be used.

I consider several commonly-used behavioural policies – simplification, changing the number of options, and defaults – each of which has a slightly different effect on consumers, but which together serve to demonstrate conclusions general to behavioural policy design. I find that behavioural policies change not just how well individuals do when they optimize, but also who optimizes, potentially creating a deadweight loss. The first effect I entitle the *payoff effect*: behavioural policies typically increase the payoff either of consumers taking the default or of consumers who optimize. The second – often neglected – effect I call the *composition effect*: the fact that the change in payoff also changes whether it is worthwhile for consumers to optimize, and so changes how many choose to optimize as opposed to take the default. Setting a higher default savings rate, for example, will help consumers who used to save less, but may harm those who used to save more. This can explain why increasing the default savings rate does not always increase the total amount saved (Choi et al., 2004), or why increasing the default donation for a charity does not increase aggregate donations (Altmann et al., 2014). Policies such as forced choice or worsening defaults may pose particularly large equity concerns by making some consumers much worse off. Careful cost-benefit analysis is needed when designing behavioural policy, and efficient policy design may impose prohibitively high information requirements on policymakers. Absent full information, several general rules emerge: that compliance with policies should be simple; that the effects of the policy should be targeted as much as possible towards either con-

sumers taking the default or optimizers rather than the population as a whole; and that policymakers should pay attention to the composition effect when designing policy.

I also briefly study the implications of the model for industrial organization, finding that monopolists can actually be efficiency-enhancing in some contexts, because they have an incentive to reduce decision costs, while more competitive firms may seek to use them as switching costs to better extract surplus from consumers. Facing a monopolist, consumers might actually benefit from increased irrationality, because it allows them to extract more surplus from the monopolist, who must offer them a strictly positive consumer surplus in order to convince them to optimize.

Relatively few studies analyze the theoretical underpinnings of behavioural policy, and the potential consequences of using it as a policy tool are not always well understood. Analysis has instead tended to be empirical, looking at issues such as the use of energy star ratings to inform consumers (Houde, 2014) and the importance of energy efficient standards in choice sets (Tsvetanov and Segerson, 2014), or in pensions, how increasing choice sets can discourage choice (Huberman, Iyengar and Jang, 2004). A number of papers have considered the use of standard policy tools when consumers are boundedly rational, such as O'Donoghue and Rabin (2003), who consider optimal taxation when consumers have self-control problems, and find that such taxation can increase social surplus by helping consumers resist temptation. Other studies have considered economic tools in other contexts, including healthcare provision, pigouvian taxation, salience in grocery store sales, and drug addiction, among others (Allcott, Mullainathan and Taubinsky, 2014; Baicker, Mullainathan and Schwartzstein, 2015; Chetty, Looney and Krof, 2009; Gul and Pesendorfer, 2007). A useful review of models of behavioural biases and how they motivate paternalistic environmental policies can be found in Allcott (2014).

What has been less studied is the importance of decision costs in understanding behavioural policy and its optimal use, and how behavioural policies themselves might interact with bounded rationality. Sallee (2014) considers whether, in the presence of decision costs, consumers could rationally choose not to optimize, and suggests rational inattention can play an important role in efficient technology adoption: Houde (2014) considers a similar question empirically, finding that product certification can act as a substitute for more accurate but more complex information. The closest paper to this

one, however, is probably Choi et al. (2003), which models defaults when consumers vary in their optimal savings rate. Choi et al. highlight the importance of heterogeneity in optimal savings rates in setting defaults - in some situations, they find it is optimal to set a default far from the mean optimum saving rate, because doing so encourages optimization. This paper takes a similar approach, but introduces the importance of heterogeneity in decision costs for efficiency and welfare considerations.

I illustrate my theoretical model with the example of the UK Behavioural Insights Team (BIT) encouraging loft insulation by households (DECC, 2013). Relatively few consumers have been quick to adopt energy-efficient technologies, though many can pay for themselves relatively quickly and also have public benefits. To address the issue, the BIT created a combined offer of loft insulation and cleaning, finding that the joint offer increased take up even when no financial savings were realized (BIT, 2011-2012). This suggests that it was not just the financial cost of insulation and cleaning that was relevant. Instead, individuals found insulation more attractive when it was a less difficult decision to implement. Governments would like to encourage consumers to adopt such technologies, because they have externalities: consumers may be reluctant to do so, because of internalities, how their decisions affect the relative welfare of their future and present selves. Behavioural policy is one approach to resolving this issue.

Unfortunately, a lack of clear theoretical understanding has led governments to eagerly adopt behavioural policy as a replacement for traditional policy, even when inappropriate. This paper argues that such policies can bring costs of their own: increasing the default savings rate or the pharmaceutical plans available to seniors helps some to the detriment of others. Behavioural policy provides policymakers with another tool for their toolbox, but not a silver bullet to be used whenever possible.

I introduce a framework of decision-making costs (Section 2.3) to examine the impact of behavioural policy options on technology adoption, including simplifying decisions (2.4.1), reducing the choice interval (2.4.2), and modifying defaults (2.4.3). Perhaps more importantly given the surge in popularity of these policy measures, the model is also used to study the limitations of each policy. Implications for the behaviour of firms is considered in Section 2.6. Section 2.7 concludes.

## 2.2 LITERATURE

Behavioural economics has expanded greatly in recent years, and policy applications have followed suit. Starting from a relatively simple understanding of decision difficulty, in which people were assumed to make errors, the recent trend has been to study heuristics, time-saving cognitive rules that lead to an acceptable choice in many situations, without requiring actual optimization. In general, studies have tended to be empirical in nature, particularly experimental, examining the effects of a policy intervention. The findings have led to considerable enthusiasm about their potential for changing behaviour. Less work has been done, however, in modeling how such policies work, and so studies often lack a comprehensive theoretical framework with which to explore the implications of their results. It is this gap in the literature that the following work is meant to address.

### *Behavioural Economics and Bounded Rationality*

The issue of individuals finding choice difficult is not new, of course. In a modern context, Herbert Simon (1955) introduced the idea of bounded rationality, arguing that individuals are not perfect optimizers. More recently, behavioural economics has addressed many of the same questions around imperfect optimization, beginning with work by Kahneman and Tversky (1974; 1984). The results of behavioural economics have also begun to affect other areas of the discipline, including economic policy (Congdon, Kling and Mullainathan, 2011), but also suggest new issues in decision-making in general, such as mental state and identity (Akerlof and Kranton, 2010).

A typical experiment on decision-making costs concerns the availability of jam (Iyengar and Lepper, 2000). Two booths distributing free samples of and then selling jam were set up in supermarkets. One was the extensive-choice booth: it offered 24 different types of jam. The other had more limited choice, with only 6 jams. Iyengar and Lepper found that while 60% of passersby stopped at the extensive-choice booth, only 40% of passersby stopped when there was limited choice. Customers tried a similar number of jams at each booth, however, and while 10% of consumers purchased jam at the limited-

choice booth, only 3% purchased jam when presented with extensive choice. Similar results were found in the context of essay topics offered to students and chocolates.

Early theoretical work on decision-making costs includes work by Radner and Rothschild (1975). There, a decision maker faces a time constraint, and so cannot optimize in every period. In order to do better at one activity, therefore, he must do worse at another. More recently, efforts have been made to model agents choosing the magnitude of a costly deliberation input (for an overview, see Conlisk, 1996). Collectively, however, such efforts encountered a common problem: that of infinite regress. In brief, it suggests that the problem of allocating cognitive effort among activities is itself a difficult activity. In turn, that decision requires cognitive effort, and so on (for further discussion, again see Conlisk, 1996). The consumer faces a problem that rapidly increases in difficulty.

To avoid this problem, some models assumed that individuals would simply make errors. In response, it might be optimal for policymakers to reduce the choice set, or even coerce the population into making the same choice (see, for example, Mirrlees, 1987). Other responses have taken a more formal approach, showing how the issue can be resolved by making the consumer decision problem a function of all their options, including how they refine their perception of those options (Lipman, 1991). More recent work, however, argues that people use heuristics to solve the decision-making problem, relying on cues from the context in which the decision is made and simple rules to select an option (Gigerenzer and Selten, 2001; Goldstein et al., 2001). Using these simple techniques individuals can avoid the cognitive effort required for calculations on what can be very difficult problems, potentially choosing to not make a decision, to do what everyone else is doing, or to pick what seems easiest, instead of attempting to figure out what is optimal.

### *Policy Applications*

The traditional response by policymakers has been to provide more information, hoping to encourage optimization. Errors and imperfect optimizing, this assumes, are a result of inadequate information, and so increasing the volume of information available increases the quality of decisions. Unfortunately, although providing more information may help in some cases, it may also increase the cacophony of inputs, leading to information overload and even worse decision-making (Toffler, 1970).

Literature examining decision-making costs from a policy perspective often focuses on the issue of pensions, perhaps because of the relatively direct nature of defaults that can be implemented. To understand the impact of defaults, Choi et al. (2004) study the effect of automatic default saving rates in three different companies, and find that 401(k) participation rates increase from between 26% to 43% for relatively new employees before the defaults are introduced, to over 85% in all three companies, with many individuals saving at the default rate. Huberman, Iyengar, and Jang (2004), on the other hand, examine how increased choices of funds within a pension plan affects participation rates, finding that a ten-fund increase led to a 1.5% to 2% drop in participation, as people avoided making the increasingly difficult decision altogether.

More recently, policymakers have turned to Sunstein and Thaler's (2008) 'nudges': minor alterations to the environment which can affect the choices people make. Governments have adopted these techniques with enthusiasm, believing them a solution to policy problems in an age of austerity.

One of the clearest models of behavioural policy is the work of Choi et al. (2003). They introduce four elements: that there are costs of opting out of a default, that there is an option value in waiting to incur those costs, that individuals procrastinate opting out of a default, and that there is heterogeneity in optimal savings rates. Given these factors, they argue that it may sometimes be optimal to set defaults that are on the boundary of possible savings rates, due to the increased incentive to opt out of the default.

Work has also focused on another issue: encouraging the adoption of energy efficient technology, which has both private and public benefits. One example is the work by the UK Behavioural Insights Team, which aimed to increase take-up of loft insulation services (DECC, 2013). The policy proposed was to reduce the costs associated with the decision by creating a combined offer of loft insulation and cleaning. Two treatments were implemented: one in which the joint package's cost was simply the sum of the cost of the two independent services previously available in the market, providing the consumers with no monetary savings, and one where a subsidy was provided. The subsidized treatment increased take-up by a factor of four, as might be expected, but more interestingly the full-cost treatment increased take-up by only slightly less (BIT, 2011-2012). This suggests that it was not just the physical or monetary cost of cleaning

the loft that was relevant; individuals also sought to avoid a search cost of finding an appropriate service, or a cognitive cost of determining which service was best.

### 2.3 MODEL

Consider a population of households considering whether to insulate their lofts. They can either do as the government recommends, the default choice, or figure out what choice is best for their particular circumstances. To optimize, however, they incur significant cognitive costs choosing the optimal level of insulation and an appropriate firm, as well as material costs of insulation, installation, and preparing their loft for insulating. Households that choose well can gain significant savings on their energy bill in the long run but have to pay a relatively large short-run cost, so the optimal choice for each household varies with their short-run need for liquidity as well as their type of home and individual energy use.

In other words, the population of households faces a decision that may be good for them in the long run, but requires them to incur decision costs immediately. They are also present biased, which means they overestimate the scale of those short-term costs relative to the long-run benefits. This describes a wide range of decisions faced by households, and so is sufficiently general to be useful in a wide range of policy domains. Before continuing, however, it is useful to consider the costs and benefits of household technology choice in more depth.

#### 2.3.1 *Costs and Benefits of Technology Adoption*

The decision about whether to adopt a new energy efficient technology by a consumer has several interesting characteristics. In particular, it involves immediate costs to choose and install a given technology in exchange for distant benefits in decreased energy use. Agents may also face some uncertainty over the nature of the technologies they can adopt before they research them as well as how much effort it will take to select the optimal technology. Even with perfect information, they are also likely heterogeneous in the costs and benefits they experience from optimization. Finally, agents may have systematic biases in their estimation of costs and benefits, leading to inefficient choices. I discuss each of these factors in turn.

It is clear that choosing to adopt a new technology will involve immediate costs, both cognitive and material. In order to adopt a new energy-saving technology, a household must acquire information about its options (information), compare that information to its preferences and decide what is best for them (introspection), and then implement the decision, integrating the technology into their household (integration).<sup>1</sup> Material costs are relatively frequent: search costs, for example, might be expected to be present in almost any decision. There is also considerable empirical evidence for cognitive costs, as for example in tests of ego-depletion: Baumeister (1998) and others have found that consumers can become exhausted by making decisions. Decision costs thus appear to be common among consumers.

In the future, however, the decision can also bring returns: increasing home energy efficiency now means spending less on energy later. Neither of these factors is necessarily known ex-ante; the costs of adoption will vary with the technology adopted, and the returns will vary with future energy needs as well as market prices, all of which may be uncertain. Households must estimate the future returns as best they can, uncertainty exacerbated by the fact that before paying the costs of optimization, consumers do not know what their optimal technology is.

Even in a world of perfect information, challenges remain. To focus on these trade-offs, I assume consumers face no uncertainty, omitting the expected value notation from the model.<sup>2</sup> I define  $c$  as the ex-ante estimated cost of taking a given decision. This is the difficulty compared to taking the easiest option, which I will refer to as the *default choice*: this may be inaction, but could equally be following a social norm or accepting a government suggested default. The default choice therefore need not be a single choice, but can instead be applying a simple heuristic or other decision rule. For simplicity, the cost of taking the default choice is taken to be zero.

Agents vary in how hard they find decisions. Some consumers have a loft with a generation of accumulated possessions, while others have no cleaning to do; an employee of

<sup>1</sup> Estimates of these costs vary. Under basic assumptions about the savings from insulating a loft and the utility function of a consumer, I use the fact that most consumers choose not to insulate their loft to calculate what the present value of the decision cost must be in order to prevent them from adopting a new technology, as indeed most consumers do not. This estimate comes to approximately £200, providing a rough order of magnitude for the scale of decision costs in this context. Consumers are prepared to sacrifice considerable material benefit to avoid cognitive and decision costs.

<sup>2</sup> If uncertainty is formally considered, the results are qualitatively similar, with the caveat that it is the effect on consumers' expectations that now determines the impact of policy, and that consumer behaviour may also vary with their preferences over risk.

a home improvement firm may find it relatively easy to decide which insulation is best for his own loft (or may not, if he has too much information available). This individual aspect is captured as  $\alpha$ , indexed by individual. From a policymaker's perspective, it is the distribution of  $\alpha$ ,  $f(\alpha)$ , that matters: for individuals,  $\alpha c$  captures their unbiased ex-ante estimate of how difficult they will find a decision.

Individuals also face systematic biases in their treatments of the costs and benefits of optimizing. Here, the benefits to optimizing are distant in time, while installation costs are incurred immediately. As in many of the interesting decisions we make in our lives, the data suggest we do not respond well to such a challenge: instead, many of us overweight immediate benefits (Laibson, 1997). We therefore sacrifice too much in the future for too little in the present. This preference for the present has been found in a variety of contexts, including the allocation of time and the choice between money-time bundles, though the exact mechanism is sometimes disputed (Benhabib, Bisin and Schotter, 2010; Zauberman and Lynch, 2005).

The costs of technology adoption occur immediately, while the benefits are generally distant in the future: consumers are thus particularly liable to present bias, and place more weight on the costs than on the benefits of technology adoption. They may then choose to take the default when it would be optimal absent present bias for them to optimize. This effect may also be exacerbated by other factors, such as ambiguity aversion: since decision costs are uncertain, while the status quo provides a certain return, ambiguity averse consumers will prefer to remain at the status quo even when it is rational to switch (Ellsberg, 1961).

This tendency to overweight costs is captured as  $\lambda$ , so that consumers anticipate their total decision costs to be  $\lambda\alpha c$ .

### 2.3.2 *Decision Rule*

I have so far assumed that (i) agents face a decision cost,  $\alpha c$ , if they wish to optimize with a new technology instead of taking the default choice. This cost represents both material search costs of making the decision as well as potential cognitive costs, and is heterogeneous over individuals, with variation that is captured by  $\alpha$  and distributed as  $f(\alpha)$ . If they do optimize, they pay their decision cost,  $\alpha c$ , but incur no loss otherwise.

Additionally, I now assume that (ii) if they take the default, they incur a heterogeneous loss measured as the present value of the utility loss as compared to their optimal choice, captured as  $L$ . This is the foregone benefit from not optimizing. The consumer loss function is distributed over the population as  $g(L)$ . For simplicity, I assume that  $\alpha$  and  $L$  are independently distributed, though as we shall see below, any correlation between them would significantly affect the optimal policy choice.

If people could accurately perceive the value of  $L$ ,  $\alpha$ , and  $c$ , they could then make an optimal trade-off between the cost to adopt a new technology and the loss if they take the default option. Unfortunately, the evidence suggests that individuals do not accurately perform this trade-off, and so I further assume that (iii) agents overweight the cost of making a decision by a factor of  $\lambda$ . Finally, I assume (iv) that agents are aware of the values of all of these parameters,  $\alpha$ ,  $c$ ,  $\lambda$ , and  $L$ , but act naïvely, taking  $\lambda$  as part of their utility functions instead of realizing they will not actually incur the cost. In this way it affects their decisions, but not their payoffs. The policymaker has full information.

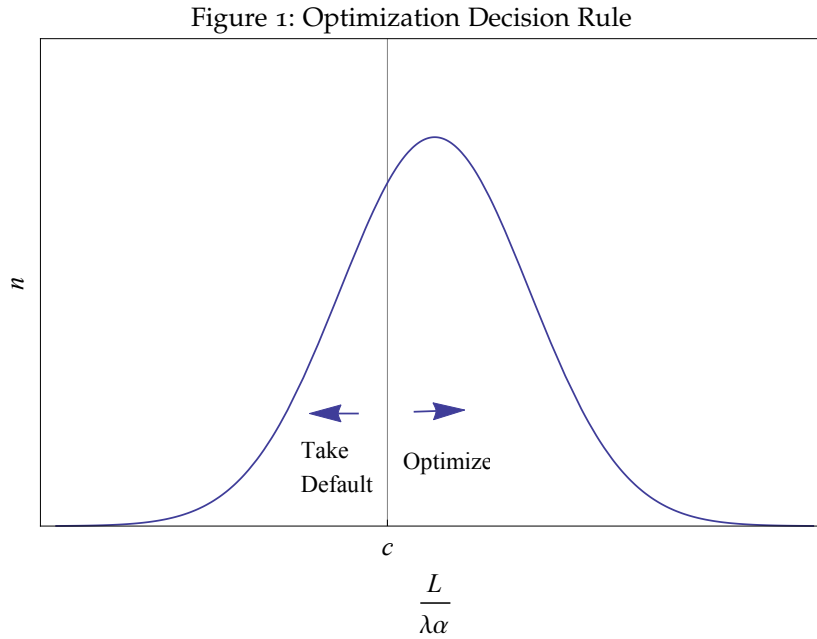
Consumers must therefore make a cost-benefit analysis: they compare their perception of the costs of the decision against the benefits. This trade-off is represented by the decision rule

$$\begin{cases} L \geq \lambda\alpha c & \text{Agent Optimizes} \\ L < \lambda\alpha c & \text{Agent Takes Default} \end{cases} \quad (1)$$

If the loss from the default,  $L$ , is larger than their perception of the costs of researching and adopting a new technology,  $\lambda\alpha c$ , then the agent chooses to optimize. If  $L$  is smaller, then they will choose to take the default. This default may simply be to preserve the status quo, continuing to have an uninsulated loft, or to follow a simple heuristic or other rule: doing what their neighbours did, or accepting the recommendation of a home builder. In all cases, however, the default represents the option with the lowest decision cost. Crucially, this decision rule is efficient only if  $\lambda = 1$ .

$L$ ,  $\lambda$ , and  $\alpha$  are all characteristics of the agent, whereas  $c$  is based on the difficulty of the decision. An agent therefore makes their decision by comparing their value of  $\frac{L}{\lambda\alpha}$  against the difficulty of the decision,  $c$ . Figure 1 captures the result of this comparison:

agents with  $\frac{L}{\lambda\alpha}$  below the critical value of  $c$  take the default: agents with a higher value of  $\frac{L}{\lambda\alpha}$  optimize.



Agents decide whether to optimize or take the default based on their values of  $L$ ,  $\lambda$ , and  $\alpha$ .

Following Schwartz (2004), I refer to individuals who prefer to optimize given all but the best of defaults as *maximizers*, and individuals who would prefer to take the default given all but the worst of defaults as *satisficers*. In the model, maximizers have a low value for  $\alpha$ , and a high value for  $L$ : satisficers have a high value for  $\alpha$ , and a low value for  $L$ . In addition, I add a third group: the *marginal optimizers*, who may currently be optimizing or taking the default, but are prepared to stop or start optimizing given a small push from policy. This group has moderate values for at least one of  $\alpha$  or  $L$ .

Every individual in the population employs the decision rule as shown in 1, with individual variation in  $\alpha$  and  $L$ . The population therefore divides into two groups based on their individual characteristics. The first group, with a sufficiently low  $\alpha$  and/or high  $L$ , will choose to optimize, and will get utility of  $-\alpha c$ . The second group, with a high  $\alpha$  and/or a low  $L$ , will choose to take the default and get utility of  $-L$ . Total utility, expressed as a loss function, is therefore the sum of these two groups:

$$U = \int_0^\infty \left[ - \int_0^{\frac{L}{\lambda c}} \alpha f(\alpha) d\alpha - \int_{\frac{L}{\lambda c}}^\infty L f(\alpha) d\alpha \right] g(L) dL \quad (2)$$

Note that the policymaker is not present biased ( $\lambda = 1$ ), but that present bias still plays a role in social utility by determining how many consumers optimize and how

many take the default. This also implicitly assumes a utilitarian social welfare function, an aspect which will be discussed later. For simplicity, I initially assume a scalar value for  $L$ , instead of treating it as a distribution over the population. This helps clarify the trade-offs inherent in behavioural policies and their effect on the number of consumers who optimize. This is presumably also an accurate reflection of situations in which there is little or no heterogeneity in loss functions. This assumption is relaxed in Section 2.4.3.

## 2.4 POLICY INTERVENTIONS

I examine four separate cases related to the model: when the policymaker controls the complexity of a decision; when the policymaker controls the size of the choice set; when the policymaker controls a default; and finally, combinations of the preceding three policies. All of these can arguably be referred to as nudges, given their behavioural economic motivations, but defaults are perhaps the most common case. For each, I examine their impact on welfare and their optimal design.

First, note that traditional economic policies may not be effective in addressing decision costs. If consumers are taking a default bundle of environmental technologies – a government recommendation, or a simple heuristic such as the  $\frac{1}{n}$  rule (Benartzi and Thaler, 2001) – that contains too many inefficient technologies, a policymaker might consider taxing carbon or inefficient technologies to incentivize efficient technology adoption. Doing so encourages some consumers to stop taking the default and adopt a more energy efficient technology bundle, but also distorts the capital allocation decision of the optimizers, creating a deadweight loss. If consumers are rational, such that  $\lambda = 1$ , the net effect will be negative, as follows from the First Theorem of Welfare Economics. Even when  $\lambda > 1$ , governments often consider behavioural economics and behavioural policy as an alternative to traditional monetary instruments.

### 2.4.1 *Decision-Making Costs*

A natural alternative to taxing energy is to inform the public which appliances use the most energy and how consumers can best reduce energy use, simplifying the consumer

decision. In the model, that is reflected in a fall in  $c$ . I find that simplifying decisions brings two complementary benefits, both encouraging individuals to optimize and increasing the benefits to optimizers, in contrast to the policy options considered below.

It is perhaps the most obvious thing a policymaker can do to help a decision-maker faced with a difficult choice: make it less difficult, at some cost to society. This could entail making the actual problem simpler, by improving labeling or sorting of options, or by reducing the difficulty a person finds with a problem, perhaps by increasing knowledge about the problem or familiarity with the trade-offs required. For energy-saving technologies, this could include energy star labeling, providing resources to inform the consumer decision, or the ‘tiering’ of options (Huberman, Iyengar and Jang, 2004). Note that the cost to society of reducing  $c$  is likely convex, since as a problem becomes easier it will become increasingly hard to simplify further. In addition, some costs of the decision may be inescapable: regret over lost opportunities, or the need for introspection, for example. Otherwise, if  $c$  could be reduced to zero, all individuals could choose optimally at no cost to themselves. These costs may be direct, as when the government incurs costs to provide resources to consumers, or indirect, as when firms are forced to increase prices to compensate for conforming to regulations simplifying consumer decisions.

The policymaker has the welfare function from equation (2) with a cost to society of decision simplification,  $G(c)$ , included. I initially treat  $L$  as a scalar, for simplicity: this is relaxed in section 2.4.3.

$$U = - \int_{\frac{L}{\lambda c}}^{\infty} Lf(\alpha) d\alpha - c \int_0^{\frac{L}{\lambda c}} \alpha f(\alpha) d\alpha - G(c) \quad (3)$$

Individuals with a low  $\alpha$  optimize and incur a cost of  $\alpha c$ , while the remainder of the population take the default and incur  $L$ . Society incurs  $G(c)$ , the cost in social welfare of reducing  $c$  further, with  $G'(c) < 0$ . The total social welfare loss is the sum of these three types of costs.

The policymaker wishes to reduce the complexity of the problem when the derivative with respect to  $c$  is negative: that is, when reducing  $c$  will increase welfare. Taking the first derivative gives us:

$$\frac{dU}{dc} = - \underbrace{\frac{L^2}{\lambda c^2} f\left(\frac{L}{\lambda c}\right) \left(1 - \frac{1}{\lambda}\right)}_{\text{composition effect}} - \underbrace{\int_0^{\frac{L}{\lambda c}} \alpha f(\alpha) d\alpha}_{\text{payoff effect}} - G'(c) \quad (4)$$

The first term in (4) captures the change in who optimizes as a result from the change in  $c$ : a lower  $c$  means more people will choose to adopt their optimal technology. Since optimizing is now easier, fewer consumers are discouraged from optimizing, and so the number taking the default falls. This is the composition effect. Because agents overestimate the size of the costs due to  $\lambda$ , too few of the marginal optimizers optimize, and so reducing  $c$  helps compensate for this, leading to a welfare gain. The larger the number of consumers who are convinced to start optimizing and the more they overestimate decision costs, the larger the benefits from the composition effect.

The first term reflects the welfare impact of convincing people to start optimizing. The second term captures the benefit to optimizers from incurring a lower cost, the payoff effect. Since optimizing is now less costly, anyone who does optimize pays a lower cost to do so, increasing total welfare. Finally, recalling that  $G'(c) < 0$ , the third term reflects the cost to society of lowering the decision cost. Since the first term is negative, the composition effect provides an incentive to lower equilibrium decision costs below what an analysis that only accounts for the payoff effect would suggest: the size of this benefit depends on the distribution of decision costs, how accurately consumers weight future costs, and their relative loss from taking the default and optimizing.

How does optimal policy compare to when consumers are not present biased? Consider the benchmark case:  $\lambda = 1$ , and individuals weight costs accurately.

$$\frac{dU}{dc} = - \int_0^{\frac{L}{c}} \alpha f(\alpha) d\alpha - G'(c) \quad (5)$$

Here, the first term disappears and the second is strictly negative, and so selecting the optimal value of  $c$ ,  $c^*$ , is a matter of equating the marginal cost of reducing  $c$  against the marginal benefit in terms of reduced complexity for those who do optimize. When people optimize correctly, there is no welfare gain from individuals choosing whether or not to optimize, since reducing  $c$  no longer compensates for an insufficiently large incentive to optimize, and the composition effect does not affect welfare.

What about when  $\lambda > 1$ ? When agents overweight costs, the marginal benefit from encouraging them to optimize is positive, but the number of agents who are optimizing

as a fraction of the total population is also smaller, *ceteris paribus*, reducing the size of the payoff effect. As a result, the benefits to reducing the decision cost also increase. The new optimum value of  $c$  is therefore ambiguous relative to the optimum when  $\lambda = 1$ . Formally, the net benefit of reducing the decision cost increases for a marginal increase in  $\lambda$  when an elasticity condition holds:

$$-f' \left( \frac{L}{\lambda c} \right) \frac{L}{\lambda c} > f \left( \frac{L}{\lambda c} \right) \quad (6)$$

When the increase in benefits from the composition effect outweighs the fall in the payoff effect, then a policymaker should simplify a problem beyond what they would in the benchmark case. If we compare a situation where the same number of consumers optimize in both cases, however, the composition effect adds a strict incentive to reduce the decision cost further.

The optimal level of decision complexity depends on the marginal benefit from both the composition effect and the payoff effect, and analyses that account for only one of these effects will underestimate the optimal level of simplification appropriate in a given context. A policymaker who realizes that consumers are present biased and face decision costs, but does not account for the composition effect in their cost-benefit analysis, will unambiguously set the decision complexity too high because they do not account for the benefit to consumers who were taking the default, and begin optimizing.

Reducing the difficulty of a decision doesn't just help those who were already optimizing: it also encourages consumers to optimize. Equally, creating a policy that is complicated to obey doesn't just hurt those who must now struggle through it: it also imposes costs on those who now decide not to optimize at all. Policymakers, who often seem to prefer the complicated over the simple, should account for this.

#### 2.4.2 *Choice Intervals*

A vast array of energy efficient technologies are available for adoption. Even if well labeled and explained to the consumer, the selection can be intimidating, and a policymaker may naturally consider whether reducing the number of options available might increase welfare. A second policy alternative is therefore to restrict the choice set. At an extreme, this could entail forcing all individuals to make the same choice, per-

haps requiring all households to install the same technology. More moderately, it could mean restricting the number of options available for a given type of technology, such as insulation or window glazing, or creating categories of technology such as high-cost, high-savings technologies and low-cost, low-savings technology, but allowing only one option within each type. This may make it easier for decision-makers to choose between the remaining options by reducing the effort required to optimize, but potentially at the cost of preventing some agents from fully optimizing, as their optimal choice might not be in the choice interval provided. As in the case of reducing  $c$ , the net effect of changing the choice interval is composed of several parts that I examine in turn. A policymaker requires substantial information about the preferences of a population to use choice intervals effectively, though this challenge can be simplified by looking only at the preferences of certain critical subgroups.

#### 2.4.2.1 *Welfare Effects of Changing the Interval Size*

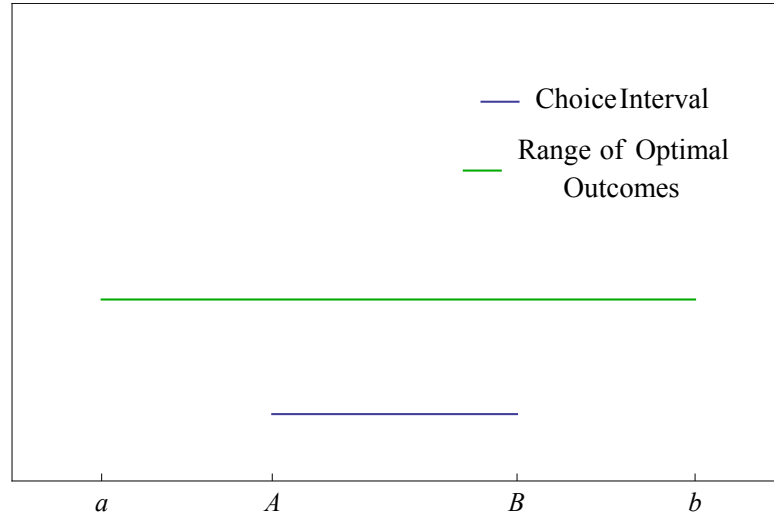
Before considering policy design, a brief discussion may help illustrate how consumers are affected by the size of the interval.

Intervals are useful because of the heterogeneity of preferences. If some individuals in a population would prefer choices outside the current interval, their payoff to optimizing would increase if the interval was expanded. Equally, if everyone wanted the same point in the choice set, then expanding the size of the interval beyond that point has no benefit.

Consider a population with optimal choices uniformly distributed between  $a$  and  $b$ : the ratio of yearly savings to installation costs, perhaps. Depending on their individual characteristics, each individual may have a slightly different optimal level. A policymaker offers an interval of an arbitrary size between  $A$  and  $B$ , as shown in Figure 2.

Individuals who have an optimal point outside the choice interval do the best they can by selecting a point on the boundary of  $(A, B)$ . An individual whose optimal choice is  $a$ , for example, would select  $A$ . Absent decision costs, it is intuitive that increasing the size of the interval will increase the number of agents who can select their optimum, and so make them better off.

Figure 2: Choice Interval



The policymaker chooses a choice interval out of the range of all possible outcomes to allow consumers to select from.

Total utility for the population is equal to the utility of those individuals who can optimize, which without loss of generality I set to zero, plus individuals whose preferences lie outside the interval, and do the best they can by selecting the boundaries. The cost of choosing away from the agent's optimal point is taken as quadratic, so that being forced farther away from their optimum imposes increasing costs.

Total social welfare is therefore:

$$U = - \int_a^A \frac{1}{b-a} (A-x)^2 dx - \int_B^b \frac{1}{b-a} (x-B)^2 dx \quad (7)$$

Setting the population interval to  $[0, 1]$  such that  $a = 0$ ,  $b = 1$ , and defining the choice interval to be  $A = \frac{1}{2} - \varepsilon$ ,  $B = \frac{1}{2} + \varepsilon$ , with  $\varepsilon \leq \frac{1}{2}$ ,

$$U = -\frac{2}{3} \left( \frac{1}{2} - \varepsilon \right)^3 \quad (8)$$

Finally, differentiating utility with respect to the size of the interval,  $\varepsilon$ ,

$$\frac{dU}{d\varepsilon} = 2 \left( \frac{1}{2} - \varepsilon \right)^2 \quad (9)$$

Increasing the size of the choice interval increases utility. In this simple case, the interval should encompass the preferences of the entire population, so that  $\varepsilon = \frac{1}{2}$ . Since  $\frac{d^2U}{d\varepsilon^2} = -2 + 4\varepsilon$  and  $\varepsilon \leq \frac{1}{2}$ , the second derivative is negative and increasing  $\varepsilon$  has

diminishing marginal returns. More generally, utility as a function of the interval is cubic: for small values of  $\varepsilon$ , the impact of a change is small, but as  $\varepsilon$  increases the effects of a change in  $\varepsilon$  increase, as seen in equation (9), where an increase in  $\varepsilon$  translates to an increase in  $[A, B]$ . When the interval is small, it is less likely to be optimal to increase the interval.

Increasing the size of the interval can also impose costs on boundedly rational consumers, however, in the form of research to learn about new options as well as effort to compare possible outcomes. The second step can be particularly difficult when comparisons need to be made about non-quantifiable properties, or along multiple dimensions when no one option is best in all dimensions. Choosing a higher salary at the same job may be simple, but changing jobs, which may involve changing location, coworkers, and responsibilities, is often a difficult choice.

Consider a decision-maker faced with  $n$  different technologies which he could adopt to reduce energy consumption in his home. I assume that the decision-maker must research each option, at an average cost  $r$  per option, to determine its level of energy savings, maintenance costs, installation costs, and other characteristics. The expected marginal research cost to increasing the interval by one option is therefore  $r$ . Second, the decision-maker must compare this technology to the other options, determining his preferences on issues including pay-back time, upfront costs, and level of disruption from installation.

For simplicity, I model this second part of the decision as pairwise comparison: the decision-maker compares in turn all different pairs of options.<sup>3</sup> The total number of pairs,  $p$ , is therefore  ${}_n C_2 = \frac{n(n-1)}{2}$ , and  $p$  is convex in  $n$ . Maintaining the simplicity of the example, we can then define the complexity of the decision,  $c$ , as

$$c = \beta rn + p = \beta rn + \frac{n(n-1)}{2} \quad (10)$$

<sup>3</sup> This provides an upper bound for the possible costs of comparison, as it ignores the possibility of transitivity in choices. It seems reasonable that preferring  $a$  to  $b$ , and  $b$  to  $c$ , provides information about  $a$  and  $c$ . Since the data is not conclusive on to what extent real preferences are transitive (for example, Kahneman and Tversky, 1979), the issue of transitivity is left for further study. Another approach would be to follow the optimal search literature: agents search at random through options, and adding another option adds a linear cost based on the chance it might be chosen instead of the optimal one (for example, Stigler, 1961). This cost is therefore diminishing as  $n$  increases. This could provide a more detailed justification for the research cost suggested above, but does not account for decision costs from comparing options.

where  $\beta$  is a weighting variable reflecting the relative importance of  $r$  and  $p$ . This example is done in discrete space, but an analogous process could be conducted over continuous decisions by breaking the choice interval into discrete sections (saving roughly 8% against roughly 10%, for example) and comparing those sections. The rate of change in  $c$  is therefore;

$$\frac{dc}{dn} = \beta r + n - \frac{1}{2} \quad (11)$$

The complexity of the decision is convex in the size of the interval, and increasing the number of available options by one increases decision costs by  $\beta r + n - \frac{1}{2}$ . The larger the current number of choices, the larger the effect on decision complexity of adding or removing a choice.

This suggests that the benefit to reducing the size of the interval diminishes as the size of the interval falls, while the costs to decreasing the size of the interval increase as the interval shrinks, as established above. There will therefore be an optimal choice interval that balances these two effects.

#### 2.4.2.2 *Optimal Choice Interval*

Let us assume that the set of possible technologies, interval  $I \in \mathbb{R}^\ell$ , maps to a set of optimal savings-to-cost ratios. Each agent therefore seeks to select their optimal point from within the interval. For convenience, let us define the welfare effects of the interval size as  $Z(I)$ , which modifies the loss the consumers get from taking the default as opposed to optimizing. The difference in welfare between optimizing and taking the default is then  $L + Z(I)$ . The decision costs are a function of interval size,  $c(I)$ . The policymaker maximizes the social welfare function:

$$U = \int_0^{\frac{L+Z(I)}{\lambda c(I)}} (Z(I) - \alpha c(I)) f(\alpha) d\alpha - \int_{\frac{L+Z(I)}{\lambda c(I)}}^{\infty} L f(\alpha) d\alpha \quad (12)$$

As before, total welfare is the sum of the welfare of two groups: those who optimize and those who take the default.<sup>4</sup> The welfare impact of a change in the choice interval is therefore the derivative of the utility function with respect to  $I$ : when it is positive,

<sup>4</sup> Note that the societal cost to a policy, earlier  $G(c)$ , has been omitted. This remains an issue, but is omitted in the next two cases in order to highlight the innate costs to such a policy, even if it were free to implement. Adding it would only increase the trade-offs presented here.

reducing the choice interval will reduce welfare. Taking the derivative with respect to the interval,

$$\frac{dU}{dI} = \underbrace{\left[ \left(1 - \frac{1}{\lambda}\right) (L + Z(I)) f\left(\frac{L + Z(I)}{\lambda c(I)}\right) \right]}_{\text{composition effect}} \underbrace{\left[ \frac{Z'(I) c(I) - (L + Z(I)) c'(I)}{\lambda c^2} \right]}_{\text{payoff effect}} + \underbrace{\int_0^{\frac{L+Z(I)}{\lambda c(I)}} (Z'(I) - \alpha c'(I)) f(\alpha) d\alpha}_{\text{composition effect}} \quad (13)$$

Though less straightforward to interpret than the earlier case, the change in utility again reflects two basic effects. In the benchmark case when  $\lambda = 1$  and individuals optimize correctly, the first term is zero. The impact therefore depends only on the direct effect of a change in the interval on optimizers, captured in the second term. Reducing  $I$  reduces their cost of optimizing,  $c$ , but also means some agents will no longer be able to select their optimal choices. This is the payoff effect.

When  $\lambda > 1$ , the first term is non-zero. It captures the utility benefit from convincing individuals who were not optimizing because of their over-weighting of costs to begin doing so. The result is a net benefit to welfare, but only when individuals do not optimize correctly. As in the previous case, this is the composition effect.

Reducing the choice interval brings two benefits, reducing the costs paid by optimizers and reducing the disincentive to optimize. As before, the composition effect presents an important factor to account for in cost-benefit analysis. Even if the direct benefits of providing more alternatives outweigh the costs, the loss in welfare from reducing the number of people who optimize may lead to a net negative impact. This may have been the case with Medicare Plan B, in which the Bush administration's attempt to allow seniors to select their own drug plans ended up discouraging participation (Thaler and Sunstein, 2008). Similarly, it may be beneficial to reduce the choice interval even when doing so removes options that some consumers value, since the net effect may be to encourage people to optimize.

If  $Z'(I) \geq \alpha c'(I)$  for  $\alpha \in [0, \frac{L+Z(I)}{\lambda c(I)}]$  and  $\lambda$  is sufficiently small, then both terms are positive, and so increasing the interval will make agents better off. Intuitively, this condition holds when increasing the interval brings more benefits than costs to every

individual who chooses to optimize, and the composition effect is relative small: the net effect is therefore positive.

The net effect can be negative, however, even if some agents are made better off. This is true whenever increasing the interval discourages sufficient consumers from optimizing, or when it makes consumers who do optimize sufficiently worse off. The net effect will then depend on the relative densities of the population and  $\lambda$ , as can be seen from equation 13. In some cases, it will be optimal for the policymaker to reduce the size of the choice interval, leading to a potential Pareto improvement, or a full Pareto improvement if compensation to agents who are made worse off is possible.

Finally, note that the key consumers in equation 13 are those whose decision difficulty is between  $[0, \frac{L+Z(I)}{\lambda c(I)}]$ . The effects of reducing the choice interval matter only in their impact on those who are already optimizing, or who are likely to change their decision of whether to optimize: those who consistently take the default are irrelevant, because they will simply continue to take the default. Policymakers considering changing the choice interval can simplify their problem by considering only agents who are currently optimizing or are likely to begin doing so. As a simple example, allowing teaser-rate mortgages as a non-default option may bring benefits to society even if they are harmful to unsophisticated buyers, because the only agents who optimize and use them may be sufficiently sophisticated to take on the risks, while unsophisticated buyers take the default, simple, mortgage. The overall information requirements are therefore not as large as they initially seem. Instead of requiring information about the entire population's preferences, the policymaker is interested only in the preferences of a certain subset. Nevertheless, this remains a relatively strong requirement on the policymaker, and in many situations it is unlikely that the policymaker will have access to sufficient information.

Note that if a policymaker could remove options that no agent found optimal, then  $Z'(I) = 0$ , and so  $\frac{dU}{dI} \leq 0$ : reducing the interval will always improve welfare. If such areas of the choice set can be found, therefore, then it is optimal to remove them from the available options. This will not, of course, always be possible: if agents are sufficiently heterogeneous, or if in a Hayekian sense a policymaker is uncertain about which options are best for agents, then  $E[Z'(I)]$  may be nonzero everywhere. Uncertainty over which areas are optimal for agents does not preclude optimization by the policymaker,

however, who can weight the expected value of  $Z(I)$  based on the information the policymaker does possess.

Two lessons for policymakers emerge from this analysis. One, that only certain sections of the choice interval matter to welfare, depending on who is optimizing and who is taking the default; and two, that an optimal choice interval will have to account for the fact that reducing the interval encourages people to optimize, and so may be of particular benefit when people overestimate costs. These lessons are discussed in detail below.

### 2.4.3 *Default Options*

In some cases, it may be difficult to affect the complexity of a decision or the number of choices. As a result, a government may consider trying to change the default option: implementing a single recommended insulation type and level, for example. The goal is for the government to reduce the losses people incur from taking the default by altering what default people take. This approach is seen as desirable over coercion or choice restriction because they still allow those who strongly prefer another choice to select it, instead of requiring a homogeneous choice for the whole population.

So what exactly is a default in this context? It is the low cognitive-cost option, the easiest choice. Not participating, taking the same option as last time, or imitating others are all possible defaults. Importantly, even in the absence of policy, defaults can and generally do still exist, options that I will call *natural defaults*. In pensions, this can be the equal allocation of capital across funds using the  $\frac{1}{n}$  rule (Huberman, Iyengar and Jang, 2004), or to invest in the fund with the highest historical returns (Cronqvist and Thaler, 2004).<sup>5</sup> For energy efficient technology, it is often to remain with whatever technology is currently installed or to follow an official recommendation. This explains why the best time to target households with energy efficient advice is when they are already moving home or replacing technology (Verplanken and Wood, 2006), or why almost all houses in the UK that install insulation do so to the 270mm of mineral wool insulation

<sup>5</sup> Not always sensible: in Sweden's 2001 reform of their social security system, the firm with the highest historical returns received a large market share, but also lost 70% of its value over the following three years.

the UK government recommends (Trust, 2013). These may or may not be good criteria on which to select technologies: they are, however, examples of natural defaults, and are likely chosen not on the basis of a personalized future-looking analysis, but rather by a simple heuristic. In some cases the natural default may remain in existence even as official defaults are introduced, potentially leading to multiple modes within a distribution, as people mentally establish one or the other option as the default.

Even basic defaults are therefore not simple. There is, however, a further complication: changing the default changes the expected benefit from optimizing, and so affects who optimizes. Similar to reducing the choice interval, improving the default option (modeled as reducing the losses people incur from taking the default,  $L$ ) can have significant indirect as well as direct effects that may move in opposite directions. The net impact of defaults depends on the distribution of the population. For that reason, I first consider a simple case with a homogeneous value of  $L$  before allowing it to vary across the population.

#### *Homogeneous L*

As before, I begin with aggregate welfare from (2). I consider a policymaker able to improve the default option for agents, thereby reducing  $L$ .

Clearly,  $L = 0$  is optimal, but that may not be a feasible choice. I therefore examine whether it is optimal to improve the default when the default cannot be set to be optimal for all consumers:  $L$  is strictly positive. Why would this be the case? Intuitively, if people incur different losses from the same default, there is unlikely to be one best default. Instead, any default will be good for some and bad for others.

The policymaker has a welfare function:

$$U = - \int_{\frac{L}{\lambda c}}^{\infty} L f(\alpha) d\alpha - c \int_0^{\frac{L}{\lambda c}} \alpha f(\alpha) d\alpha \quad (14)$$

To find the impact of changing the default, I differentiate with respect to  $L$ , giving:

$$\frac{dU}{dL} = \underbrace{f\left(\frac{L}{\lambda c}\right) \left(\frac{L}{\lambda c}\right) \left(1 - \frac{1}{\lambda}\right)}_{\text{direct effect}} + \underbrace{F\left(\frac{L}{\lambda c}\right) - 1}_{\text{indirect effect}} \quad (15)$$

Where  $F(X)$  is the cumulative distribution function of  $\alpha$  evaluated at  $X$ .

As before, we see two distinct effects. The first term captures the change in who decides to optimize, the composition effect, while the second term captures the benefit to those who take the default and now incur a lower loss, the payoff effect. When  $\lambda = 1$ , only the second term is non-zero, since decisions on whether to optimize or not are optimal. When  $\lambda > 1$ , however, making a default worse encourages more consumers to optimize when it is efficient for them to do so, and so has a net benefit.

Rearranging equation (15), we see that the value of  $\frac{dU}{dL}$  depends on the hazard rate,  $\frac{f(\frac{L}{\lambda c})}{1-F(\frac{L}{\lambda c})}$ : the chance that a given individual has  $\alpha = \frac{L}{\lambda c}$  given that he does not have a lower value of  $\alpha$ .  $\frac{dU}{dL}$  is positive when

$$\left(1 - \frac{1}{\lambda}\right) \left(\frac{L}{\lambda c}\right) \frac{f\left(\frac{L}{\lambda c}\right)}{\left(1 - F\left(\frac{L}{\lambda c}\right)\right)} > 1 \quad (16)$$

The hazard rate is critical because it reflects the relative weight of individuals who are already taking the default against those who will decide to start taking the default based on the new value of  $L$ . The optimal policy is therefore determined by the ratio of individuals who do better because they are still taking the default to individuals who do worse by having their decision whether to optimize distorted.

What distinguishes this policy from varying the interval? Now, the payoff effect benefits those who continue to take the default, instead of those who optimize. It is therefore the number of individuals who take the default, rather than the number who optimize, that matters relative to the number who inefficiently stop optimizing. In consequence, when a policymaker improves a default, individuals who take the default are better off, and more people take the default, leading to a self-reinforcing cycle.

To see this, consider the critical value of  $c$ ,  $c^*$ : when  $c \geq c^*$ , it is optimal for a policymaker to improve the default. Rearranging equation (16), we get;

$$c^* = \frac{L \left(1 - \frac{1}{\lambda}\right) f\left(\frac{L}{\lambda c}\right)}{\lambda \left(1 - F\left(\frac{L}{\lambda c}\right)\right)} \quad (17)$$

When  $c \geq c^*$ , then it is optimal for the policymaker to improve the default, since  $\frac{dU}{dL} \leq 0$ . If the decision is sufficiently difficult, then many individuals take the default, and so the policymaker should improve it. When  $c < c^*$ , the default should be made worse: then, relatively few people take the default, and so a policymaker can do better by encouraging them to optimize by increasing the loss from taking the default. Depending

on the nature of  $f(\alpha)$ ,  $c^*$  may be a local or global crossing point, though in some cases the single crossing property will hold, as when  $\alpha$  has a uniform density.

Consider the case when  $\lambda = 1$ . Above, I noted that this meant encouraging individuals to optimize had no effect on social welfare. As a result,  $c^* = 0$ , and so  $c > c^*$  and it is always optimal to improve the default, regardless of its initial value (ignoring possible societal costs of default improvement). This reflects the fact that when  $\lambda = 1$ , the first term in equation (15) is zero: if everyone is optimizing correctly, then there is no composition effect, and the default should be made as good as possible.

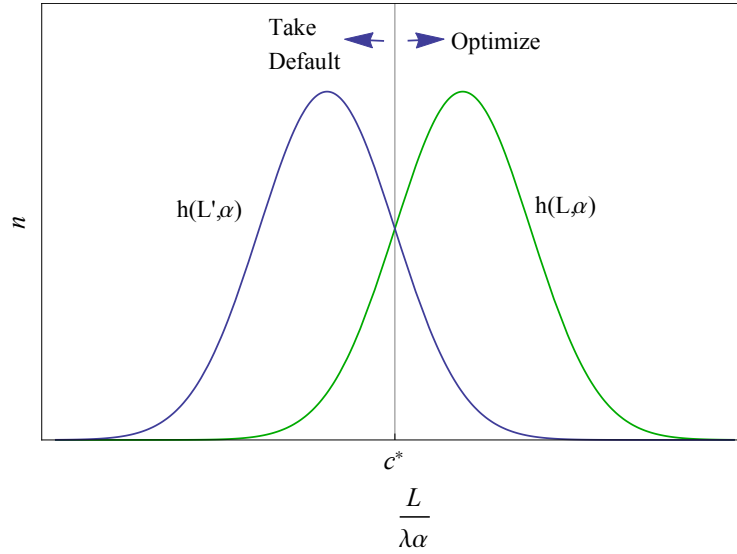
When  $\lambda > 1$ , however,  $c^*$  is nonzero. It is then possible that  $c < c^*$ . It is then optimal to implement a worse default, rather than a better one.

The intuition behind this result is that unlike in the simple case of reducing decision complexity, but similar to the case of controlling the choice interval, the conflicting composition and payoff effects from improving a default mean the sign of  $\frac{dU}{dL}$  is ambiguous. When a default is already bad, or a decision is particularly complicated, then improving the default only helps the few people currently taking the default. It also distorts the optimization decision, however, convincing some consumers to stop optimizing when they would be better off doing so. If there are many consumers who choose to stop optimizing, then imposing a default will lower social welfare. When a default is already good, in contrast, many people are taking the default, and so the benefits of improving a default are large. Figure 3 presents the effects of an improvement in the default from  $L$  to  $L'$  for an arbitrary distribution of  $\frac{L}{\lambda\alpha}$ .

First,  $L$  falls to  $L'$ , moving  $h(L, \alpha)$  left to  $h(L', \alpha)$ . This increases the welfare of those who take the default, but also increases the number of people who take the default. The policymaker must therefore balance the benefits to those beneath the green curve on the left hand side of the graph against the potentially negative effects to the group between the green and blue curves on the left hand side, which represents individuals who started taking the default in response to the policy. Depending on the shape of the curves, the net effect could be positive or negative.

Equation (17) has implications not just for when the impact of a default is positive, but also for the nature of the optimal solutions. In particular, note that when the  $\frac{dc^*}{dL}$  is positive, then the optimal values of  $L$  will be boundary solutions. To see why, note that  $\frac{dc^*}{dL} > 0$  when:

Figure 3: Optimization in Response to an Improved Default



When a default is improved, the share of the population that chooses to take the default increases because their loss from doing so falls. The population shifts from the green to the blue distribution.

$$\left[ f\left(\frac{L}{\lambda c}\right) + f'\left(\frac{L}{\lambda c}\right) \left(\frac{L}{\lambda c}\right) \right] \left(1 - F\left(\frac{L}{\lambda c}\right)\right) + f^2\left(\frac{L}{\lambda c}\right) \left(\frac{L}{\lambda c}\right) > 0 \quad (18)$$

This holds when  $f\left(\frac{L}{\lambda c}\right) + f'\left(\frac{L}{\lambda c}\right) \left(\frac{L}{\lambda c}\right)$  is either positive or small and negative. When  $f'\left(\frac{L}{\lambda c}\right) > 0$  (when the density of the population with decision cost  $\alpha$  increases as  $\frac{L}{\lambda c}$  increases), improving the default causes  $c^*$  to fall, and worsening the default causes  $c^*$  to increase. This occurs because improving the default encourages an increasingly small number of people to start optimizing when  $f'\left(\frac{L}{\lambda c}\right) > 0$ . The cost in terms of lost optimizers is therefore diminishing as the default is improved, and so the net benefit increases.

As a result, improving the default when  $c \geq c^*$  loosens the relevant condition. To illustrate, if  $c$  is above  $c^*$ , then the default should be improved. This causes  $c^*$  to fall, and so it remains optimal to continue improving the default, reducing  $c^*$  further. If it was optimal to improve the default before, it remains optimal to do so. The result is a boundary solution: the default should be made as good as possible. The opposite also holds: if it was ex-ante optimal to worsen the default, then as it increases fewer people will take the default, and the benefits to worsening the default increase further: the default should be made as bad as possible.

If  $f'\left(\frac{L}{\lambda c}\right) < 0$ , in contrast, an increasingly large number of individuals start taking the default when the default is improved. This implies that the marginal cost in lost

optimizers from improving the default increases. If it increases sufficiently, the cost in reduced optimizers outweighs the marginal benefits, and it may no longer be optimal to continue improving it. The solution is therefore interior. This is particularly likely to be true when the loss from the default is high, since the cost to encouraging people to take the default is then larger.

It can therefore be optimal to improve or worsen a default, and in some cases an extreme default may be welfare-maximizing. If many people are already optimizing and relatively few taking the default, then improving the default helps only a few people, while potentially discouraging a larger number from optimizing at all.<sup>6</sup> Marginal optimizers are harmed by an improvement to the default, while individuals who continue to optimize are unaffected. The only group to benefit are those who are already taking the default and whose decision on whether to optimize or not is unlikely to change. When a default is worsened, in contrast, optimizers remain unaffected and individuals who take the default do worse, but agents who are convinced to start optimizing do better. Designing an optimal policy requires balancing these effects.

#### *Heterogeneous L*

I have so far omitted the possibility of heterogeneity in the losses from taking the default,  $L$ . This allows for some individuals to choose to optimize, and others to take the default, but enforces uniformity within each of those groups other than over  $\alpha$ . If we allow for heterogeneity in  $L$ , however, then not only does the population divide into those who optimize and those who take the default, but individuals within each group will vary in how bad they default is for them. The policymaker therefore has an interest not just in how many people respond to the policy, but who within each group decides to start or stop optimizing, since some of them incur a higher loss from the default.

I earlier assumed  $L$  to be non-zero. Now the rationale for this follows directly from the model, since changing a default will make some people better off and others worse off: ‘improving’ the default from  $L$  to  $L'$  implies that  $L$  has first-order stochastic dominance

<sup>6</sup> A related concern is that a government induced default may often be simpler than the previous, natural default. Not opting out of a government sponsored renovation, for example, poses a smaller cognitive cost than finding out what your neighbours did. This suggests that  $c$ , thought of as the difference between taking the default and optimizing, may increase when a stronger default is created, further exacerbating the composition effect.

over  $L'$ . If the default is to install a minimal amount of insulation, for example, leading to energy wastage but incurring a low installation cost, creating a new default with a higher insulation level will be good for those who have disposable income, but bad for young people or others who require short-term liquidity, even if the increased insulation pays for itself several years later. These groups may in turn have different reactions to the same policy.

Since the policymaker can no longer choose a single level of loss for the population, let us define the distribution of the loss function over the population as  $g(L, k)$ , where  $k$  is the policymaker's choice of default. By altering  $k$ , the policymaker is able to change the distribution of  $L$ .<sup>7</sup>

The policymaker chooses a default,  $k$ , to maximize the social welfare function:

$$U = \int_0^\infty \left[ -c \int_0^{\frac{L}{\lambda c}} \alpha f(\alpha) d\alpha - \int_{\frac{L}{\lambda c}}^\infty L f(\alpha) d\alpha \right] g(L, k) dL \quad (19)$$

Differentiating with respect to  $k$  gives:

$$\frac{dU}{dk} = -c \int_0^\infty \int_0^{\frac{L}{\lambda c}} \alpha f(\alpha) d\alpha g_k(L, k) dL - \int_0^\infty \int_{\frac{L}{\lambda c}}^\infty L f(\alpha) d\alpha g_k(L, k) dL \quad (20)$$

When the policymaker changes the default, they affect the number of consumers who take the default as well as the losses from those who do. The sign of  $\frac{dU}{dk}$  determines when it is optimal for a policymaker to reduce or increase  $k$ . An improvement in the default (reduction in  $k$ ) is optimal when the losses to the group that stops optimizing are outweighed by the benefits to the group that take the improved default:

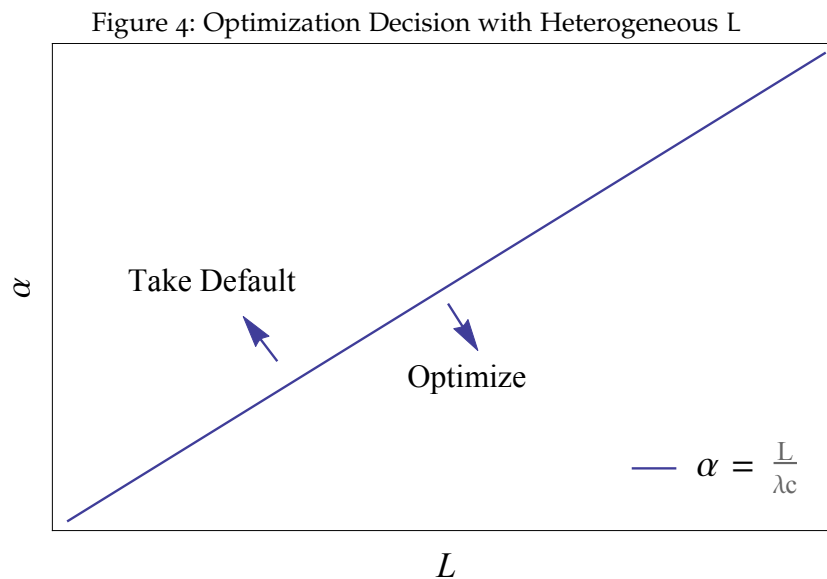
$$c \int_0^\infty \int_0^{\frac{L}{\lambda c}} \alpha f(\alpha) d\alpha g_k(L, k) dL > - \int_0^\infty \int_{\frac{L}{\lambda c}}^\infty L f(\alpha) d\alpha g_k(L, k) dL \quad (21)$$

When both the distribution of  $L$  and the distribution of  $\alpha$  are considered, the same issues as in the homogeneous- $L$  case are relevant. In addition, however, the interaction between the two variables now matters. Intuitively, it is the ratio between the loss they

<sup>7</sup> This implicitly imposes independence between  $L$  and  $\alpha$ , since  $f(\alpha)$  is not a function of  $k$ : a dependent relationship could be captured through the use of a joint distribution of  $L$  and  $\alpha$ .

incur from the default and their decision cost that determines an agent's choice to take the default or optimize, and so the interaction between their relative densities will determine who responds to a policy measure. Clearly, if  $L$  and  $\alpha$  were correlated, the nature of the correlation would play a large role in determining the welfare impact of a policy.

Previously, households who found decisions difficult would take the default, while those who found them easy would optimize. Now, households that find decisions difficult may still optimize if they stand to lose a lot if they take the default, and households that find decisions easy may choose not to optimize if taking the default costs them little, as in Figure 4. As before, if the group that stops optimizing is large, then the net effect of the policy could be negative, but the size of the effect now also depends on which agents stop or start optimizing, not just how many of them do. The interaction between the decision cost and the loss from taking the default determines the welfare impact for a consumer of taking the default or optimizing.



When both  $L$  and  $\alpha$  vary in the population, a consumer will decide whether to optimize or take a default based on their relative values.

When might these distributional concerns be significant? Consider a group of individuals attempting to choose an energy-saving technology. Low net-worth individuals rely on their income for subsistence, making savings relatively important, and so have a high  $L$ : since many may also be less experienced with house renovations, they also have a high  $\alpha$ . High net-worth individuals, on the other hand, have a relatively low value of  $L$  due to their reduced need for savings, and can also afford to hire experts to help them with their renovations, meaning their value of  $\alpha$  is relatively low. A policy-

maker that improves the default slightly will encourage low net-worth individuals to take the new, still high-loss default instead of optimizing and lower their welfare, while high net worth individuals continue to optimize. The negative effects of the policy are exacerbated because it is high L individuals who begin taking the default when it is improved.

#### 2.4.4 *Simultaneous Policies*

In some situations, it may be valuable to combine policies. In taxation, for example, several jurisdictions have experimented with the idea of 'prefilled' tax returns, essentially implementing a default return. Simultaneously, however, the tax code is often used in an attempt to change behaviour by providing incentives for various types of activities. Simplifying the tax return can therefore encourage people to take advantage of incentive-based taxes, maximizing their effect, and also ensures more people optimize their tax return. Without simplification, many individuals who might otherwise respond to the incentives may just take the default, reducing the impact of the incentives.

Determining whether behavioural policies are complementary is a broad question, but some basic observations can be made. Perhaps the largest advantage of such combinations is that even if combining policies has only a small effect on aggregate welfare, it can have significant distributional consequences. Combining policies can allow the policymaker to target policies better, such that the individuals who are best affected by the policy can benefit from it, while reducing its impact on others. This, in turn, means that combining approaches may minimize the distortions from a single policy tool. Establishing a default while reducing the complexity of a decision can help both those who optimize and those who take the default, while mitigating the intergroup incentives for some individuals to stop optimizing, for example. The same applies to changing the choice interval and establishing a default: since the interval affects optimizers, and the default affects those who take the default, the policies can be used to obtain their benefits on the target groups while to some extent mitigating their consequences in intergroup movement.

Achieving such results is not necessarily simple. The effect of using two policies simultaneously varies heavily with the characteristics of the population, and so making broad generalizations is not straightforward. To illustrate, consider the submodularity of  $\{(L, c)\}$ , as in the case of prefilled tax returns:

$$\frac{\partial^2 U}{\partial L \partial c} = - \left(1 - \frac{1}{\lambda}\right) \left(\frac{L}{\lambda c}\right) \left[ f' \left(\frac{L}{\lambda c}\right) \left(\frac{L}{\lambda c^2}\right) + f \left(\frac{L}{\lambda c}\right) \right] - f \left(\frac{L}{\lambda c}\right) \left(\frac{L}{\lambda c^2}\right) \quad (22)$$

Note that both terms are negative when  $f' \left(\frac{L}{\lambda c}\right) \geq 0$ , and  $U$  is submodular in  $c$  and  $L$ . Then, a policymaker choosing to reduce  $c$  will want to increase  $L$ , and vice versa: they are strategic substitutes. Intuitively, reducing the cost of optimizing,  $c$ , will make optimizers better off and encourage an increasingly large group of people to start optimizing (since  $f \left(\frac{L}{\lambda c}\right)$  is locally increasing). This reduces the number of agents taking the default, and so reduces the marginal cost to increasing  $L$ . It is therefore optimal to increase the loss from taking the default in order to encourage even more people to optimize.

If  $f' \left(\frac{L}{\lambda c}\right)$  is negative and large, then  $U$  is supermodular in  $c$  and  $L$ , and decreasing  $c$  means the policymaker should also decrease  $L$ : they are strategic complements. When this condition holds, reducing the cost of optimizing makes optimizers better off, but increasingly few new people start optimizing. As the number of borderline optimizers falls, the marginal cost of discouraged optimization from reducing  $L$  decreases. It is optimal to improve the default and reduce  $L$  so that the group of people still taking the default do better, since this will convince only a small group to stop optimizing.

The value of  $f' \left(\frac{L}{\lambda c}\right)$  plays a similarly important role in determining the supermodularity of  $U$  in the case of interval size and defaults or interval size and decision costs, though with different specifics. Since  $f' \left(\frac{L}{\lambda c}\right)$  determines how the population density will change, it is a critical parameter in determining how policies should be employed.

The issue of how policies interact is a complex one, and heavily dependent on the nature of the population. If we additionally consider the distribution of  $L$ , then a policymaker is concerned not just with the slope of  $f'$ , but with its interaction with  $L$ . In general, however, as a policymaker encourages an agent to take the default some individuals will do worse than others, and so there is an incentive for the policymaker to try to target policies more precisely. Combining policies can allow for more such targeting by controlling several parameters within the model. Any such combination is sensitive

to the distribution of the population, however, and so broad conclusions are difficult to draw. For the moment, it is sufficient to note that policies may well have beneficial interaction effects, and that it is an area in need of further exploration.

## 2.5 POLICY LESSONS

A number of policy implications can be drawn from the preceding analysis. The fundamental lesson is that behavioural policies can have costs and benefits similar to traditional economic policy, and must be carefully designed to fit a given population and problem. Balancing the payoff and composition effects creates a trade-off that can play a significant role in choosing when – and how – to nudge.

Decision costs are the most straightforward of the policies I examine. Simplifying decisions brings two complementary benefits, encouraging individuals to optimize and increasing the benefits to optimizers, and so is unambiguously useful to the policymaker. Rather than taxing energy and creating a deadweight loss, for example, policymakers should first turn to policies like simplifying the decision to insulate a loft. Doing so both helps those who optimize as well as encourages households to optimize, both of which increase welfare.

Too often, in areas like taxation, pensions, welfare programs, and environmental regulation, policymakers seem to err on the side of increased complexity in order to accommodate all interests, failing to realize that they may discourage the very populations they are trying to help from taking advantage of them. In the case of Medicare Plan B, for example, its complexity reduced its potential benefits (Thaler and Sunstein, 2008). Policies like the British one of bundling loft cleaning with insulation, in contrast, provide an unambiguous benefit and are a desirable line of policy inquiry.

The other two policies I examine are more complex. When designing a choice interval, the policymaker must account for both the payoff as well as the composition effect, realizing that adding or removing choices will affect the incentives to optimize as well as the welfare of the households that do optimize. Fortunately, optimal policy design is simplified by the need to address the preferences of only a subsection of the population.

Households that always take the default, for example, are unaffected by the size of the choice interval as long as it contains the default. Choices that are only optimal for them

can thus be removed from the interval without affecting welfare. The important section of the population, in contrast, is those who are or may become optimizers. Removing options that this section of the population would choose reduces welfare by denying them their optimal choice. The optimal interval, therefore, removes all options optimal only for default-takers, while potentially leaving those optimal for optimizers.

Beyond this minimum, it may also be optimal for a policymaker to further restrict choice if the benefits in terms of reduced decision costs are sufficient to compensate for the loss of options. In addition, the employer may wish to eliminate options even if they are good for some optimizers, in order to encourage more individuals to optimize. Whether this is the case will depend on how many households lose their optimal choices and how many are convinced to optimize: the net impact of the payoff effect and the composition effect.

Finally, I consider the effect of establishing a default. In general, defaults are seen as useful policy tools because they allow individuals with a lot to gain from optimizing to opt out and continue to optimize. If the distribution of preferences is wide, then a default helps those in the center of the distribution while allowing for the extremes to continue optimizing. This simple conception of defaults omits a critical second parameter, however: the distribution of decision costs within the population. Similar to the case of choice intervals, a policymaker considering improving a default must balance a payoff effect from improving the welfare of consumers who take the default, and a composition effect from convincing individuals who were previously optimizing to take the default. If consumers overweight their decision costs, improving the default will encourage too many to stop optimizing, which can have a negative impact on welfare. The hazard rate allows the measurement of the distribution of these trade-offs.

Analyses that consider only the payoff effect and not the composition effect thus overestimate how often defaults are useful. Because improving defaults helps consumers who take the default and hurts those who are convinced not to optimize, it is inefficient when few consumers are taking the default but many could be inefficiently convinced to start if it was slightly improved, as occurs when problems are simple, when the suggested default is still quite poor for most consumers, or when people do not overweight costs significantly. In addition, the impact of using a default is determined not just by how many households change behaviour in response to a default, but also which ones

do: a household which incurs a large loss from taking the default may still be convinced to do so if they also have high decision costs, incurring a larger loss than a household with a low loss from taking the default and low decision costs.

Ideally, a policymaker should select a default that improves the outcomes only for individuals who were already not optimizing. This avoids incentivizing individuals to stop optimizing, while still helping those who take the default, and so negates the composition effect. Defaults should thus be designed to target as exclusively as possible the preferences of households already taking the default. Creating the default level of energy efficiency to reflect the needs of older households, though they may represent the majority of the population, will be ineffective if it encourages them to stop optimizing, while younger households, who do not optimize, are faced with a default that makes them worse off.

In some cases it may be welfare improving to worsen a default, in order to encourage optimization. With a sufficiently bad default, only a very few people with extremely large decision-making costs will choose the default. This, however, leads concerns over equity. Though it may be efficient to set a worse default, the group with high decision-making costs who select the default will be significantly worse off than those who choose to optimize, based on characteristics that are arguably not theirs by choice. This may lead to a conflict between utilitarian welfare functions, which I have implicitly assumed, and Rawlsian functions, which would support worsening defaults only when the worst off in society are helped by it.

Some analyses have suggested forced choice as equivalent to a default which is so bad as to force agents to optimize (Choi et al., 2003). Unfortunately, this suggestion fails to acknowledge the presence of natural defaults and heuristics outside of official control. It may be that the number of people who would employ such a strategy is limited, but they may incur significant utility losses, suggesting that forced choice does not resolve the problem faced by a Rawlsian policymaker.

Concerns over equity relate to more than just the choice of social welfare function, and there may well be rationale for a government to care about different individuals differently. In a complicated world, individuals who find decisions easy consistently do well, and so may have a lower marginal utility of income. As a result, making optimizers better off may not contribute as much to social welfare as imposing a default

which helps individuals who find decisions difficult. In a laissez-faire society, a group of individuals who find decisions difficult will end up taking the default in almost all situations. That same group, therefore, will consistently do less well than an identical one that finds decisions easier. A paternalistic state may feel a responsibility to help individuals who are consistently disadvantaged in society, regardless of the impact on individuals who find decisions easy.

How do these policy implications compare to the literature? The introduction of a formal model allows for a systematic cost-benefit analysis of behavioural policy, something frequently not done in applied policy design. Many treatments of the problem, and to an even greater degree much policy design, fail to account for the extra benefits from encouraging people to optimize. In the case of simplification, this is a difference of degree, not of sign, but for choice intervals and defaults, omitting the composition effect can actually reverse the net welfare effect of the policy. The models highlight the need to closely consider who is affected by a nudge and how they will react, instead of hoping that consumers will naturally behave in the expected manner. Nudges can impose distributional effects as serious as standard economic policy, and equity, as well as efficiency, can arise as a serious concern in their design. Particularly challenging is the fact that potentially efficient policies, such as forced choice or worsening a default, may have particularly costly equity implications, meaning the value of such policies will depend on the goals of the policymaker.

## 2.6 FIRMS

Firms are also in a position to affect the decisions of consumers. In particular, they are often able to reduce the decision complexity faced by consumers, either by restricting choice intervals, or more directly by providing clearer labeling or giving purchasing advice. I find that monopolists can actually be efficiency improving as they have an incentive to simplify consumption decisions, while more competitive market structures can exploit decision costs as switching costs, potentially increasing profits but reducing efficiency. This may explain the observed level of complexity consumers often face, from breakfast cereal to colours of paint.



default. For simplicity, I also assume that decision costs are constant across consumers: heterogeneity in decision costs is reflected in the variation in  $V(x)$ . This implies that a symmetrical range of agents will decide to take the default, represented as the range  $[-y, y]$ . Consumers decide whether conditional on the prices set by the monopolist, they are better to take the default or optimize.

If consumers choose to optimize, they incur the decision cost,  $\alpha c$ , and pay the expected price of their optimal technology. If they take the default, they incur no decision costs, but pay the price of the default. It is thus optimal for them to optimize instead of take the default if

$$E[V(x)] - \lambda \alpha c - E[p_x] \geq -p^D \quad (23)$$

where  $E[p_x]$  is their expected price of their optimal technology,  $\lambda$  is how much they overweight the immediate decision cost, and  $p^D$  is the price of their default technology. Note that there is likely considerable uncertainty over  $p_x$  and  $V(x)$  from the perspective of the consumer, since by definition they do not know before optimizing what their optimal technology is: it is therefore their expectations of the value and price of their optimal technology that affects their decision. For simplicity of notation, however, I omit the expected value term in what follows.

First, each consumer is assigned an optimal technology from the interval  $x$ . Second, the monopolist sets prices for the default and non-default products: third, each consumer forms expectations about the prices they will face if they optimize; and finally, consumers decide whether it would be optimal for them to optimize or take the default.

Under the standard assumption that monopolists cannot charge consumers more than their willingness to pay, and since the value of the default to the consumer is zero, it follows that the monopolist sets  $p^D = 0$ . The monopolist's revenue is therefore

$$R = \int_a^{-y} p_x f(x) dx + \int_y^b p_x f(x) dx \quad (24)$$

Since the monopolist gets no revenue from consumers who take the default, it has an incentive to maximize the number of consumers who optimize. We can now compare monopolist behaviour to the standard case of pooling and separating equilibria.

In a pooling equilibrium, the monopolist charges a flat price  $\bar{p}$  to all consumers. A marginal increase in  $\bar{p}$  increases the revenue per unit of technology sold, but reduces the number of consumers in the market. As  $\bar{p}$  increases, more agents drop out of the market, but the firm receives a higher price from the remaining consumers. The monopolist sets the marginal revenue from a price increase equal to the marginal cost of lost consumers, and so as in the standard case the elasticity of demand will determine the optimal price for the monopolist.

Here, any consumer who has  $V(x) < \lambda\alpha c + \bar{p}$  does not optimize. Hence, at  $y$ ,  $V(x) = \lambda\alpha c + \bar{p}$ . If we assume  $V(x) = |x|$ , for example, since consumers are distributed symmetrically the monopolist's optimal price satisfies

$$\frac{dR}{d\bar{p}} = F(b) - F(\bar{p} + \lambda\alpha c) - \bar{p} f(\bar{p} + \lambda\alpha c) = 0 \quad (25)$$

The monopolist balances the increase in revenue from increasing the price against the marginal loss of revenue from consumers who stop optimizing,  $\bar{p} f(\bar{p} + \lambda\alpha c)$ . Consumer surplus is

$$2 \int_{\bar{p} + \lambda\alpha c}^b (x - \alpha c - \bar{p}) f(x) dx \geq 0 \quad (26)$$

Consumer surplus is weakly positive because monopolists cannot extract their entire surplus in a pooling equilibrium unless they charge a price equal to the willingness to pay of the consumer with the highest value of optimizing, in which case all other consumers drop out of the market.

The monopolist may do better by price discriminating between products and creating a separating equilibrium. Since consumers vary in their optimal choice along the interval, a firm can offer a menu of choices increasing in price and value to the consumer, and each consumer identifies their type by the technology they choose to adopt.<sup>8</sup>

Consumers optimize if the net value of optimizing is greater than the price they pay at their optimum,  $p_x$ . In a perfectly separating equilibrium, the monopolist therefore sets their price equal to the expected surplus of the consumer:

<sup>8</sup> Price discrimination requires firms be able to rank products by their value to customers, which may not always be possible

$$p_x = V(x) - \lambda\alpha c \quad (27)$$

For all  $V(x) \geq \lambda\alpha c$ . At all other points, the benefit from optimizing is too small to convince the consumer to optimize even if the product was free, and so they take the default.

Note, however, that consumer surplus is based on realized utility, not expected utility. The monopolist must offer each consumer a return of  $\lambda\alpha c$  in order to convince them to optimize, but the consumer only incurs costs of  $\alpha c$ . The consumer surplus to each optimizing consumer is therefore  $\lambda\alpha c - \alpha c$ . A consumer who chooses to optimize gets a strictly positive consumer surplus due to their overestimation of the decision costs involved in optimizing. Monopolists cannot extract this surplus, because if they increase the price sufficiently to extract a consumer's entire surplus, then because the consumer overestimates decision costs they will choose to take the default. The monopolist must therefore share some of the surplus in order to convince consumers to optimize.

If we again assume  $V(x) = |x|$ , then total consumer surplus is

$$CS = 2 \int_{\lambda\alpha c}^b (\lambda\alpha c - \alpha c) f(x) dx \geq 0 \quad (28)$$

The scale of the shared surplus depends on the value of  $\lambda$ ,  $\alpha$ , and  $c$ ; when these are low, little surplus must be left to consumers to convince them to optimize, but when they are high the monopolist may be able to extract relatively little from consumers. Consumer surplus is therefore potentially increasing in  $\lambda$ : increased irrationality can be welfare enhancing for some consumers, because it increases the amount that must be left to convince them to optimize. If  $\lambda\alpha c > V(x)$ , however, the consumer takes the default and gets zero surplus, and so if  $\lambda$  increases sufficiently, all consumers take the default and get zero surplus.

Price discrimination increases social surplus by allowing the monopolist to charge each consumer's willingness to pay and eliminating deadweight loss. Unlike in the standard model, however, even with considerable information about consumer preferences the monopolist still cannot extract all consumer surplus, because consumers must retain some surplus to convince them to optimize. Whether total consumer surplus is higher in the separating as opposed to pooling equilibrium, however, is ambiguous. Consumers

who do not optimize or who dropped out of the market in the pooling equilibrium are always weakly better off in the separating equilibrium, but optimizers who previously paid  $\bar{p}$  and captured the rest of their surplus may now pay a higher price  $p_x$ , and keep only what is necessary to convince them to optimize.

### *Decision Complexity*

We have so far assumed that the monopolist takes decision costs as given. In some cases though the monopolist may be in a position to change decision costs, by offering advice or decision tools to consumers.

If the monopolist cannot identify consumers, then as before it must charge the same price to all consumers. Since it makes no revenue from consumers who take the default, it wishes to maximize the number of consumers who optimize. It is therefore intuitive that the monopolist would prefer decision costs to be as low as possible, to increase the number of optimizers and maximize their surplus from optimizing.

What if the monopolist can price discriminate? One feasible strategy for the monopolist is to maximize decision costs: all agents then choose to take the default. In this case, the monopolist returns to a pooling equilibrium in which they are able to extract the entire consumer surplus from taking the default. They do so, however, at the cost of reducing the total surplus available: they can only extract the value of the default choice from consumers, since all of them take the default.

What if the monopolist sets  $c = 0$ ? Agents make decisions optimally and receive the full value of their choice,  $V(x)$ . If the monopolist is able to price discriminate, however, it can then extract this surplus. Revenue for the monopolist is therefore

$$R = \int_a^b V(x) f(x) dx \quad (29)$$

Unlike in the previous case, since the consumer pays no decision cost, the monopolist need not share surplus with them to convince them to optimize. The monopolist can therefore extract all the consumer surplus. Eliminating decision costs is also the socially optimal choice, since if decision costs are zero, all consumers get the maximum utility from optimizing, and also do not overweight the cost of making the decision. Monopo-

lists thus both maximize total social surplus, and capture all of it; they no longer face the trade-off they did when unable to change decision costs. When monopolists are able to price discriminate and control decision costs, in other words, they achieve the efficient solution, but make consumers worse off.

Monopolists able to control decision costs are thus efficiency enhancing under some conditions, though this may not make consumers better off as they can also extract considerable surplus from them. When monopolists cannot control decision costs, or when a decision is inherently difficult despite the efforts of the monopolist, they must leave some surplus to consumers to incentive them to optimize, such that optimizers have a positive consumer surplus: when monopolists can freely control the decision cost, they set  $c = 0$ , maximizing social efficiency but reducing consumer surplus to zero.

### 2.6.2 *Competitive Firms*

Consider a market structure with two firms engaged in Bertrand competition with heterogeneous products, so that  $p = mc$  for both firms. One firm offers the default technology, while the other offers a variety of possible optimal technologies for consumers who choose to optimize. One firm, for example, may be a simple home improvement firm that always insulates lofts with the same brand of insulation and to the same level, without asking consumers for input: the other is a more exclusive custom insulation firm, which leaves all decisions in the hands of the consumer.

In this market, decision costs effectively serve as a switching cost, preventing consumers from easily moving between firms. They thus allow a firm to maintain its market share even if it is relatively uncompetitive, and can also increase profits for firms (see, for example, Beggs and Klemperer (1992)).

Consumers choose to buy from the default firm if  $p^o - p^d > V(x) - \lambda\alpha c$ ; as before, the price difference, decision cost, and benefits of optimizing relative to taking the default determine behaviour. Similar to the problem facing policymakers, because consumers overestimate the difficulty of making the decision by  $\lambda$ , in equilibrium inefficiently few of them will choose to optimize.

In the standard switching costs model, once a consumer purchases from a firm, that firm becomes their default. Firms therefore make good offers to new consumers, and then exploit their stickiness from switching costs to extract more surplus in later periods. Here, the overestimation of decision costs creates a similar effect. A consumer that buys the good once will settle on it as a default, and then persistently overestimate how difficult it would be to optimize. From the perspective of firms, this increases consumer switching costs, and so increases the ability of firms to extract surplus. Competitive firms thus have little incentive to reduce decision costs.

A new entrant has an incentive to offer a simplified product, in order to attract consumers. Such entrants face a fundamental challenge, however. In entering the market, they increase the complexity faced by consumers, reducing the likelihood of any given consumer switching to their product. Consumers previously had only one pairwise comparison to make, their current choice against the offering of the other firm. With a new entrant they are forced to decide which of the two alternative firms they prefer before deciding whether to switch. Even offering a relatively simple product might be insufficient to overcome the increased complexity created by their entry into the market. In addition, even if they offer significant incentives to consumers to switch, their efforts will have an inefficiently small effect, due to consumers' overestimation of decision costs.

Decision costs in a competitive market provide some insight as to how switching costs arise as well as the potential challenges facing a new entrant hoping to acquire market share. In comparison to a monopolist, however, competitive firms have less incentive to reduce decisions costs, leading to an inefficient equilibrium and making consumers worse off despite the presence of competition in the market. Decision costs provide a motivation for preferring a monopoly to a competitive market structure for efficiency reasons, because of their incentive to simplify the decision faced by consumers while still potentially being forced to leave them some surplus.

## 2.7 CONCLUSIONS

To determine when and how behavioural policies should be used, clear modeling is essential. This paper has addressed that gap in the literature through the use of decision-making costs, drawing from the model various implications and conclusions.

My fundamental conclusion is that behavioural policies lead to not just changes in the welfare of groups in the population, the payoff effect, but also to the composition of those groups, the composition effect. In a perfect world, a chosen default would reduce the loss from taking the default only of individuals who should be taking the default based on a rational trade-off, rather than inefficiently affecting intergroup movement. This avoids incentivizing individuals to stop optimizing, while still helping those who take the default: an unadulterated benefit. Even when this is not possible, defaults and choice intervals should be designed to mimic the effect as closely as possible, targeting only the preferences of the desired population.

In the absence of perfectly targeted defaults, behavioural economics does not provide a silver bullet for dealing with decision-making costs. Instead, a policymaker must carefully balance the improvement in welfare for some against the costs to others. In doing so, he or she is helped by the availability of more than one tool: controlling decision costs, the choice interval, and defaults all present slightly different trade-offs, and so are best used in different circumstances or to complement each other within the same context. In addition, behavioural policies may still provide advantages over traditional taxes and subsidies by allowing for relatively better targeting of policies and reduced deadweight loss.

Reducing the costs of a decision provides the clearest gains because it both encourages people to optimize and makes optimizers better off, leading to complementary effects. Given it is generally a low-cost policy option, providing reliable and simple advice to consumers is likely the first-best policy for governments to adopt.

Reducing the choice interval requires considerable information about the preferences of the population, and may thus be complex. To some extent, however, limiting the choice interval may be optimal, because it encourages agents to optimize. This policy option should however be used with caution: though valuable when the policymaker has sufficient information about preferences, in other situations restricting choice efficiently may be difficult.

Finally, policymakers should not assume that always improving a default is optimal, because it can also increase the number of agents who are discouraged from optimizing, lowering social welfare. A negative default or forced choice policy, however, may not be optimal from an equity perspective, particularly if too much of the population is likely

to adhere to the default in any case. It may also lead to a group of consistent losers within society, those who find decisions difficult and always take the default.

Overall, a benign policymaker can draw three broad conclusions from the model. First, policies should be designed so that compliance is simple, in order to encourage individuals to optimize, rather than take the default. Second, policymakers should not assume that an increase in choices or improving a default will always increase welfare: both policies have their costs and benefits, most notably through the payoff and composition effects. Whether their use is efficient will depend on the relative number of consumers taking the default and optimizing, and how many will decide to switch when the policy is implemented. Finally, targeting is of critical importance in any use of behavioural policy: the better a policy can be designed to suit the needs of one population while not affecting the choices of another, the better will be its impact on welfare. If the policymaker does not have full information about the population targeted, behavioural policies may be inefficient or worse yet actually lower welfare.

As the world becomes more complex, there are indeed good reasons to employ defaults and behavioural nudges, but to believe they are a perfect solution to problems is disingenuous at best. They are a policy tool like any other, and must be critically evaluated and well-designed if they are to have a positive impact. This is particularly true given their potentially conflicting payoff and encouragement effects, and their potential for distributional consequences and equity-efficiency trade-offs. That said, they also provide a useful complement to more traditional policies, and can have an effect on and deal with challenges that other policies may be ill-equipped to solve.

*“Pay attention to your thoughts, because they become your words.  
 Pay attention to your words, because they become actions.  
 Pay attention to your actions, because they become habits.  
 Pay attention to your habits, because they become your character.  
 Pay attention to your character, because it is your fate.”*  
 - Talmud

### 3.1 INTRODUCTION

In standard economic models, preferences determine behaviour. In practice, however, preferences can also be affected by behaviour: what we do may have implications for what we want. For consumers, this is particularly relevant when the behaviour at issue is the consumption of tempting goods: succumbing to temptation, by this reading, is not just a matter of consumption but also of preferences. If the consumer does indulge, they potentially increase the appeal of further indulgence: if they resist, they may make it easier for themselves to resist further (Muraven, Baumeister and Tice, 1999; Oaten and Cheng, 2006). Those who resist temptation can do better as their preferences improve, reaping the rewards of self-control, while those who do not get trapped in a self-reinforcing cycle of indulgent behaviour. The welfare differences between these two equilibria are large, and consumers with self-control perform better on almost any measure, including aggression, anxiety, alcoholism, self-esteem, grades, income, and even dental care.<sup>1</sup> Understanding how self-control evolves over time is valuable to designing policy to fight obesity, increase savings, and deal with other policy challenges that depend on long-term patterns of behaviour.

This paper studies how consumers behave when their preferences – particularly their self-control – are affected by their behaviour. I introduce a model of quasi-hyperbolic discounting in which present bias, the temptation to consume today, is plastic and can change over time. When consumers resist the temptation to consume today and allocate

---

<sup>1</sup> The literature on these benefits is extensive. See, for example, Mischel, Ebbesen and Zeiss, 1972; Tangney, Baumeister and Boone, 2004; and Baumeister and Tierney, 2011, for studies on the benefits of self-control; and Duckworth and Seligman, 2005, for work on school results.

more consumption the future, they train themselves to resist further temptation, and so reduce their present bias.

I find that in equilibrium a consumer may either consumption smooth and reach a prudent equilibrium in which she cares about future selves and leaves resources for them, or binge, bringing consumption forward and inducing an imprudent equilibrium in which she cares little about the welfare of future selves and so leaves little for them. For some preferences, multiple consumption paths are locally optimal, and so multiple equilibria may exist. In essence, consumer behaviour can be self-reinforcing, as behaving imprudently induces preferences that encourage further imprudent behaviour, and similarly for prudence. An agent who chooses to exercise can develop a preference for going to the gym and for resisting temptation more generally; another might choose to indulge in alcohol, anticipating that it will induce myopia and make them insensible to the future costs of their behaviour. In many cases, it is an optimum for consumers to consume all their resources immediately, because doing so makes them oblivious to the future costs of their behaviour.

Differences in an agent's preferences and expectations about the future, both about the conditions he or she will face as well as about the behaviour of future selves, can determine whether the consumer self-improves or binges. The welfare differences between an agent who chooses to self-improve and one who habitually indulges in order to become insensible to the future costs of their behaviour can be wide. Consumers must therefore take care when succumbing to temptation to avoid beginning a more enduring cycle of indulgence, and may benefit from the ability to commit their future selves to specific consumption paths.

Models of individual choice generally treat preferences as fixed: one lives with the preferences one is given, and there is no incentive to change them since by definition they represent what the consumer wants. The body of literature on non-stationary preferences was begun by Strotz (1955-1956), who introduced the idea of multiple selves in a formal setting similar to that used today. These models typically focus on present-biased preferences: preference parameters are fixed, but since which period is the present changes, how consumers weight present and future benefits also changes. Preferences are thus in some sense endogenous to the period, with consumers 'tempted' to consume in the present. Consumers can resist that temptation through precommitment (Gul and

Pesendorfer, 2001), or self-control, such as Bénabou and Tirole's model of personal rules or Dhar and Wertenbroch's model of self-signalling (2004; 2012). An alternative, however, has been to allow preferences to be endogenous to other factors, such as the economic context (Bowles, 1998), or, more popularly, to the past decisions of the consumer (Becker and Murphy, 1988; O'Donoghue and Rabin, 2002).

This paper adopts the second approach, making preferences depend on the decisions of the consumer. It considers a relatively narrow case: consumers whose behaviour affects their preferences immediately. In doing so, it draws on a literature from psychology that suggests consumer willpower functions similar to a muscle, strengthening or atrophying depending on use (Muraven and Baumeister, 2000; Muraven, Baumeister and Tice, 1999), and several papers in economics that suggest that a planner or long-run self might be able to train or otherwise affect the preferences of a doer self (Fudenberg and Levine, 2006; Thaler and Shefrin, 1981). The closest work, however, is by Becker and Mulligan (1997), who consider consumer behaviour when consumers can invest resources in order to change their exponential discount rate. There, however, consumers invest resources directly into changing their discount rates, rather than it being a result of consumption, and they assume time-consistent behaviour. Some empirical evidence for such an effect also exists: Oaten and Cheng (2006), for example, find that students who are taught self-control exercises demonstrate better self-control in a number of areas, including caloric intake, spending, and emotional restraint.

Unlike typical models of habits, such as Becker and Murphy (1988), which suggest that with a low endowment, consumers consume small amounts and so form good habits, preference change in this model is based on the amount of self-control employed in a period, not the absolute level of consumption, and so consumers can develop or weaken their willpower independent of their endowment. More similarly to a habit model, but in contrast to the standard model, instead of precautionary saving consumers may respond to anticipated negative future shocks by choosing to increase immediate consumption, in order to become oblivious to those future costs. When consumers do endure a shock that forces them to consume a higher or lower fraction of their wealth, it may alter their equilibrium consumption, leading to persistence in the effect of the shock. In standard addiction models, consumers who are very present biased are more

likely to become addicted: this model shows the opposite, that consumers who believe they are likely to become addicted are likely to have higher levels of present bias.

My results explain observed behaviour as well as provide several lessons for policy. Fundamentally, policymakers may struggle to improve long-run lifestyles and behaviours if they are reinforced by preferences: punishing consumers for bad preferences is unlikely to be good policy. If agents can be shifted from bingeing to self-improvement, however, the potential increase in welfare is dramatic. In addition, negative shocks may induce imprudent behaviour in anticipation of their occurrence and, if they affect preferences, shift the consumer between equilibria. If a consumer expects to be forced by poverty to consume a high a share of their endowment, for example, then in the present they may choose to also consume more in order to become indifferent to their future disutility, making it harder for them to save and emerge from poverty. Policymakers can benefit from minimizing or compensating for the effect of such shocks.

I begin in Section 3.2 with a discussion of the cognitive foundations for the model as well as the relevant literature. In Section 3.3, I introduce the model, defining a new class of agents, ultrasophisticates, who maximize utility anticipating endogenous preference change. Section 3.4 examines the implications of the model for consumer behaviour and contrasts it to that of a standard, fixed-preference agent, highlighting the importance of agents' own preferences as well as their expectations about future conditions and selves in determining the equilibrium. Section 3.5 discusses some welfare and policy aspects of the model, while Section 3.6 concludes.

### 3.2 LITERATURE

Strotz (1955-1956) formally introduced non-stationary preferences, and his work serves as the basis for almost all more recent studies. He introduced both the idea of sophisticates, who anticipate future preferences, and naïfs, who expect future preferences to be identical to today's (which he refers to as "thrifty" and "spendthrift", respectively), and formalized the utility of the agents in much the same way as we do today:

$$\int_0^T \lambda(t - \tau) u(C(t), t) dt \quad (30)$$

Agents have utility  $u(C(t), t)$ , weighted by a term whose value depends on the distance between the current period,  $\tau$ , and the specified future (or past) date  $t$ .

Further work focused on the sophisticate case, for which Strotz had incorrectly derived an optimal consumption path based only on the rate of change in the discount function at the present time (Goldman, 1980; Peleg and Yaari, 1973; Pollak, 1968), and Phelps and Pollak (1968) introduced the hyperbolic model for intergenerational accounting. In 2008, Tyson fully characterized the naïf consumption path behaviour, under a series of progressively stronger restrictions (Tyson, 2008). Quasi-hyperbolic discounting was introduced by Laibson's 1997 work in the context of savings, and it is this formalization that I use below.

Though these models reflect non-stationarity of preferences, preferences are endogenous to the period, and are not affected by the choices of consumers. There is also considerable support, however, for the idea that preferences are affected by decisions that have been made. I outline three effects – habits, willpower, and self-signalling – which could lead behaviour to affect preferences in the way captured by the model. Broadly speaking, all three reflect the possibility for neural plasticity, that our brain and preferences can actually shift over time (Lashley, 1923; Tricomi, Balleine and O'Doherty, 2009).

### *Habits*

Habits play an important role in behaviour. At a simple level, Knowlton, Mangels and Squire (1996) find that patients with amnesia can learn over time to predict which of two outcomes will result based on cues from repeated exposure, even absent any conscious memory of past experiences or the outcomes themselves.

Equally, habits can be a powerful driver of undesirable behaviour even if no physically addictive substance is involved. Work with rats has found that if a habit of lever pressing and receiving food is formed, the rats will trigger the lever and eat the food, even when the food is poisoned and they would refuse the food in conditions when the habit is not triggered (a devaluation test; see Hilario and Costa, 2008 for details). It is this kind of behaviour that is often observed in studies of obesity or smoking, in

which individuals are capable of resisting temptation in most circumstances, but are particularly vulnerable in contexts which cue their undesirable habit (Laibson, 2001).

From a more theoretical perspective, Becker and Murphy (1988) introduced a model of rational addiction, while O'Donoghue and Rabin (2002) introduced a behavioural aspect to addiction with present bias and the presence of naïfs and sophisticates. Finally, Lally et al. (2009) fit empirically observed habit dynamics to a theoretical model, and find that it takes 18 to 254 days for participants to reach 95% of their asymptote of automaticity.

As consumers make choices, therefore, they can over time affect their preferences, making a particular behaviour more or less likely. In my model, agents can form not just habits related to a specific behaviour, but also affect their perception of future costs and benefits generally. To understand why, I turn to studies of willpower, a more general resource.

### *Willpower*

Willpower is fundamental to resisting temptation, and has become increasingly popular among economists. Recent literature has tended to focus on two stylized facts. First, that in the short-term willpower is an exhaustible resource, a situation often called ego-depletion, as coined by Baumeister et al., 1998. This is supported by considerable experimental evidence, including experiments forcing people to choose radishes over cookies, and then measuring the time they spent attempting an unsolvable puzzle, or asking individuals to suppress an emotion, then spend time solving anagrams. Second, there is also evidence suggesting that over time, individuals can build up willpower through practice, strengthening their ability to resist temptation as if it were a muscle (Baumeister and Tierney, 2011; Muraven, Baumeister and Tice, 1999; Muraven and Baumeister, 2000).

More recently, experimental work has demonstrated that repeatedly using willpower in one context makes it easier to resist temptation in other settings. In a study by Oaten and Cheng (2006), individuals suffering from a depletion of self-control due to exam stress saw marked improvements in self-regulatory abilities when taught to practice self-control exercises, not just in study habits but also in other activities. Students smoked and drank less and consumed less caffeine after the intervention, while reporting in-

creases in healthy eating, emotional control, household chores, spending control, and fulfillment of commitments.

These results motivate the effects observed in the model. When consumers resist temptation, they affect how appealing they find other temptations. Since willpower is a domain-general resource, they can use it not just to resist the temptation they just experience, but other temptations as well. This contrasts to more typical models in the literature, such as Gul and Pesendorfer (2001), who introduce temptation and self-control by giving agents a choice over menus of options in the first period, before the actual consumption choice in the second. In my model, temptation is not just an unavoidable conflict between multiple selves, but also a potential source of preference change.

### *Self-Signaling*

Similar to the self-reputation work of Bénabou and Tirole (2004), some studies suggest that decisions can serve as a way for agents to learn their own preferences or alter what they believe their true preferences to be, often thought of as self-signaling or cognitive dissonance.

The seminal work on cognitive dissonance was done by Leon Festinger in his study of how cult members adapted when reality clashed with their predictions (Festinger, 1962). Cognitive dissonance has also proven popular in economics, as for example in work by Akerlof and Dickens (1982). In essence, the theory proposes that individuals experience discomfort when holding conflicting opinions, and as a result may change their beliefs in order to reconcile their two opinions. In order to make their expectations meet reality, in other words, they change not their expectations but their perception of what is reality.

In this context, cognitive dissonance might suggest that when the consumer behaves differently from how they would have expected themselves to do, instead of believing they have made an error and done something they did not want, they change what they want. Similarly, in the self-signalling literature (Bénabou and Tirole, 2002; Dhar and Wertenbroch, 2012), a choice of a given action in the present can help convince future selves what their preferences truly are (“I went to the gym last week, so I must be the kind of person who goes to the gym”), helping to alter future outcomes. By this

reading, in my model consumers are learning about or convincing themselves of their own willpower type as they choose to indulge or resist, a process that gradually actually affects how tempted they are to consume in the present.

### 3.3 MODEL

The following analysis develops a theory of endogenous preferences in an otherwise typical infinite-horizon Euler cake-eating environment, as introduced in Gale (1967). This allows for a focus on how consumers delay rewards over time, a fundamental aspect of temptation. A consumer has endowment  $W_0$  in period zero that she must allocate as consumption across an infinite horizon. The consumer finds the urge to consume in the present tempting: this temptation is measured by present bias,  $\gamma$ . How well she can resist the temptation to consume in the present is a function of how much she indulges, and so it can vary over time.  $\gamma$  is assumed to be smooth and monotonically decreasing in the ratio of current consumption to subsequent consumption, so that consuming more in the present period increases present bias.<sup>2</sup> The consumer also discounts the future exponentially, as measured by a discount rate  $\delta$ . She therefore has a quasi-hyperbolic discount function as introduced in Laibson (1997), with present bias  $\gamma$  dependent on her chosen consumption in the current and subsequent period.

$$U_0 = u(c_0) + \gamma \left( \frac{c_0}{c_1} \right) \sum_{t=1}^{\infty} \delta^t u(c_t) \quad (31)$$

As is standard, this representation imposes additive separability of each period's utility, as well as multiplicative separability of the discount factor,  $\gamma\delta$ , and utility. For simplicity, uncertainty is omitted; including it would introduce risk aversion as an important parameter in choice. Utility is increasing and concave in consumption.  $W_0$ , the endowment, can also be thought of as the present value of all future income in the presence of complete lending markets. An interest rate of  $r$  is paid on savings, and so consumption is subject to a budget constraint of

<sup>2</sup> I focus on the case where reducing consumption increases  $\gamma$  (reduces present bias). The opposite is plausibly the case in the short run: there, willpower has been observed to be a finite resource, and so resisting one temptation might make another harder to resist. Equally, consuming more today could reduce present bias, if going without glucose reduces willpower, for example (Wang and Dvorak, 2010). Such depletion effects appear to be limited to the short run, however, and they also appear to depend on the beliefs of the consumer (Job, Dweck and Walton, 2010).

$$W_{t+1} = (W_t - c_t)(1 + r) \quad (32)$$

Disposable income equals  $W_0$  in the first period, after which it will be a function of disposable income in the last period (which determines consumption in that period) and the interest rate. In addition, the standard transversality condition is imposed,  $\lim_{t \rightarrow \infty} \gamma \left( \frac{c_0}{c_t} \right) \delta^t u'(c_t) W_t = 0$ : the discounted value of the terminal capital stock must approach zero. The fundamental departure from the traditional model is that present bias is a function of consumption.

I assume that present bias is a function of the ratio of current consumption to expected future consumption:  $\gamma \equiv \gamma \left( \frac{c_t}{c_{t+1}} \right)$ , so that the more the consumer resists consumption in the present and saves, the lower their present bias. At the extreme,  $\gamma(0) = 1$ , and  $\gamma(W_t) = 0$ ; consuming nothing today eliminates present bias, and consuming everything today maximizes present bias. This form reflects the work suggesting that consumers can develop the ability to resist temptation by doing so (Muraven, Baumeister and Tice, 1999; Muraven and Baumeister, 2000; Oaten and Cheng, 2006). For consumers, giving in to temptation does not just affect their consumption, but also affects their relative weighting of current and future benefits. This might be exacerbated by other psychological effects, such as cognitive dissonance.

Importantly, this functional form imposes two constraints on the evolution of preferences. The first is that consumption decisions affect preferences in the current period. This increases tractability by making the consumer decision a trade-off between immediate and later effects. Since a period is of arbitrary length, they can be defined as the smallest unit of time required to affect preferences via a decision. Secondly, this imposes scale independence in preference change. The wealth of the consumer does not affect their preferences directly, and consuming more in the present period need not imply a greater change in preferences: what matters is the ratio of consumption now against consumption allocated to the future, an implicit measure of the consumer's displayed future preference rather than their available resources.

Consumers thus express their preference for the present over the future by allocating consumption between the current and future periods. The more the consumer chooses to consume today compared to how much they plan to consume tomorrow, the higher their present bias will be. By choosing a level of consumption, they also implicitly choose

a level of present bias. This effect of consumption on present bias is central to the preference endogeneity of the paper: one's actions become one's preferences.

The degree of present bias is thus a function of consumption, which is in turn a function of disposable income. In addition to the standard naïfs and sophisticates of the literature (O'Donoghue and Rabin, 1999), therefore, I introduce *ultrasophisticates*, who are aware that their preferences will change based on their decisions. The model thus adds endogenous present bias to a hyperbolic discounting framework. Setting  $\gamma \left( \frac{c_0}{c_1} \right) = \beta$ , the model nests the standard results of naïf and sophisticate behaviour as a benchmark, as discussed in Harris and Laibson (2001) and Tyson (2008).

This gives us three distinct types of agents:

1. Naïfs: Fixed present bias, assume preferences in future are identical to their own
2. Sophisticates: Fixed present bias, are aware future selves will have different preferences
3. Ultrasophisticates: Variable present bias, are aware future selves will have different preferences

I briefly discuss the behaviour of the first two, before considering ultrasophisticate behaviour in more depth.

### 3.3.1 Fixed Preference Types

#### *Naïfs*

Naïfs do not anticipate future deviations from their optimal consumption path as predicted in period  $t$ , and so expect their future selves to abide by their current self's decisions. In this case, in periods  $t + 1$  and after, they expect to follow the standard exponential discounting Euler Equation,

$$\frac{u'(c_t)}{u'(c_{t+1})} = \delta(1 + r) \quad (33)$$

In the present period, however, they have an extra discount factor of  $\beta_t$ , and so will distort their first period consumption upwards.

$$\frac{u'(c_0)}{u'(c_1)} = \beta\delta(1+r) \quad (34)$$

In essence, they plan to over-consume in the first period, and then follow the optimal path for an agent with no present bias after that period. Clearly, this is not an equilibrium between the various selves, as each self will in turn experience present bias and distort consumption towards their own period. Instead, in each period they will consume as in equation 34 and so in a simple setting will be less well off than the other types. Their behaviour is further characterized in Tyson (2008).

Naïfs place no value on being offered commitment devices, because they do not anticipate any time inconsistency. As a result, there is no reason to force future selves to behave in a particular way, because they expect those future selves to want identical things to the current self.

#### *Sophisticates*

Sophisticates anticipate their future divergence from their optimal consumption path as predicted in period zero, and so adapt their decision to account for future deviations.

Following Harris and Laibson (2001), the equilibrium consumption path of a sophisticate is

$$u'(c_t) = (1+r) [c'_{t+1}(W_{t+1})\beta\delta + (1 - c'_{t+1}(W_{t+1}))\delta] u'(c_{t+1}) \quad (35)$$

The effective discount factor is now a weighted average of the short-run discount factor,  $\beta\delta$ , and  $\delta$ , the long-run discount factor.<sup>3</sup> Note that somewhat counter-intuitively, sophisticates typically consume more than naïfs in a given period, because they foresee distortions to consumption in future periods and so anticipate less utility from leaving resources to be consumed by future selves.

Sophisticates do place value on being able to commit to specific behaviours. Anticipating time-inconsistency and that their future selves may want different things than

<sup>3</sup> As a technical condition, this holds only when consumption is Lipschitz continuous, a property that holds in the neighbourhood of  $\beta = 1$ . Harris and Laibson go on to derive a generalized Euler Equation that does not require Lipschitz continuity. Further detail can be found in their 2001 paper.

their current self, being able to commit their future selves is welfare-enhancing for the present self.

### 3.3.2 *Ultrasophisticates*

I now turn to a new type of agent: ultrasophisticates, who have endogenous preferences.

In brief, the solution proceeds as follows. I construct a value function for the ultrasophisticate which is a function of their current and future consumption levels, and where present bias depends on the ratio between the two. Consumers choose the fraction of their wealth they wish to consume in the current period, and in equilibrium, the consumer consumes the same fraction of wealth in each period. I therefore derive a first order condition that defines optimal consumption for an ultrasophisticate.

Ultrasophisticate agents are able to alter their preferences through their consumption decisions, and so include the benefits of being able to alter their preferences in their calculus. The present bias in their optimization decision is therefore a function of how much the agent chooses to consume - the less they choose to consume, the more they value the future. Future selves have the opportunity to over-consume or under-consume based on those new preferences, potentially undoing the positive habits formed by this period's agent, but in equilibrium the fraction of wealth consumed will be time consistent.

To increase tractability and allow a focus on equilibrium behaviour, ultrasophisticates have isoelastic preferences, with utility as in equation 36. Non-isoelastic preferences are discussed briefly below.

$$u = \frac{c^{1-\rho}}{1-\rho} \quad (36)$$

Agents allocate their wealth  $W_0$ , the present value of future income flows, as consumption across infinite periods, and so agents face a standard cake-eating problem with the additional constraint that the consumption level determines present bias.  $\gamma\left(\frac{c_t}{c_{t+1}}\right)$  is assumed to be homogenous of degree zero, in order to introduce scale independence in the preference change; what matters is the relative consumption levels of each period.

3.3.2.1 *Decision-making*

Since present bias only affects the weight of utility in the present against later utility, not the relative weights of later utilities, the agent weighs the relative utility from periods  $t$  and  $t + 1$  differently from  $t + 1$  and  $t + 2$ . The agent's value function is therefore

$$V(W_t) = \arg \max_c u(c_t) + \gamma \left( \frac{c_t}{c_{t+1}} \right) \delta \tilde{V}(W_{t+1}) \quad (37)$$

where  $\tilde{V}(W_t)$  is the value of the income stream in future periods unweighted by present bias:

$$\tilde{V}(W_t) = u(c_t) + \delta \tilde{V}(W_{t+1}) \quad (38)$$

Note that  $\tilde{V}(W_t)$  is a contraction mapping and so admits a unique value function, and  $V(W_t)$  is then a function of this unique value function.<sup>4</sup> As is standard with isoelastic preferences, I focus on a solution in which the ratio of consumption to wealth is constant:  $c_t = \bar{\alpha}W_t$  (for further discussion, see Arrow (1970)).

For simplicity, I treat the proportion of wealth consumed in the current period,  $\alpha$ , as the choice variable, since in equilibrium  $\alpha = \bar{\alpha}$  in all periods. In the case of non-isoelastic utility, this leap to equilibrium need not hold: consumption might instead trend toward the equilibrium level over time, as agents gradually improve their preferences and each period reduce consumption slightly. Isoelastic utility allows us to focus on equilibrium behaviour.

I first rewrite the value functions in terms of fraction of wealth consumed, using the isoelastic utility function;

$$V(W_t) = \arg \max_c \frac{(\alpha W_t)^{1-\rho}}{1-\rho} + \gamma \left( \frac{1}{(1-\alpha)(1+r)} \right) \delta \tilde{V}(W_t (1-\alpha)(1+r)) \quad (39)$$

and

$$\tilde{V}(W_t) = \frac{(\alpha W_t)^{1-\rho}}{1-\rho} + \delta \tilde{V}(W_t (1-\alpha)(1+r)) \quad (40)$$

<sup>4</sup> A value function is unique if the Contraction Mapping Theorem is satisfied. For details, see Stokey and Lucas (1989).

Note that  $\frac{c_t}{c_{t+1}} = \frac{1}{(1-\alpha)(1+r)}$ , which is therefore the argument of the present bias term expressed in  $\alpha$ . For simplicity of notation, I temporarily omit writing this argument in the analysis that follows. Differentiating the welfare function 39 with respect to the fraction of wealth consumed gives the first order condition:

$$W_t^{1-\rho} \alpha^{-\rho} + \left[ \frac{\gamma'}{(1-\alpha)^2 (1+r)} \right] \delta \tilde{V}(W_t (1-\alpha)(1+r)) - W_t \gamma \delta (1+r) \tilde{V}'(W_t (1-\alpha)(1+r)) = 0 \quad (41)$$

Where  $\gamma'$  refers to the derivative of  $\gamma\left(\frac{c_t}{c_{t+1}}\right)$  with respect to  $\frac{c_t}{c_{t+1}}$ . This first order condition represents an agent's optimal consumption  $\alpha$  given their ability to form habits,  $\gamma'$ , and their wealth,  $W_t$ .

I now posit forms for the value functions, using the form familiar from the standard isoelastic case.

$$V(W_t) = \frac{\chi W_t^{1-\rho}}{1-\rho} \quad \tilde{V}(W_t) = \frac{\gamma W_t^{1-\rho}}{1-\rho} \quad (42)$$

Substituting these into the value functions 39 and 40, I find the value functions as a function of the consumer's preferences and their chosen level of consumption:

$$\tilde{V}(W_t) = \left[ \frac{\alpha^{1-\rho}}{1-\delta(1-\alpha)^{1-\rho}(1+r)^{1-\rho}} \right] \frac{W_t^{1-\rho}}{1-\rho} \quad (43)$$

This is the value to the consumer of having wealth  $W_t$  if they only discount exponentially, or equivalently how they weight the value of having  $W$  in a period that is not the present. The value of having wealth  $W_t$  in the present period is then

$$V(W_t) = \left[ \alpha^{1-\rho} \frac{1 - (1-\gamma)\delta(1-\alpha)^{1-\rho}(1+r)^{1-\rho}}{1-\delta(1-\alpha)^{1-\rho}(1+r)^{1-\rho}} \right] \frac{W_t^{1-\rho}}{1-\rho} \quad (44)$$

Finally, we can combine equations 43 and 41 to establish the first order condition that determines a consumer's optimal fraction of wealth to consume,  $\alpha^*$ , as a function of their preferences and their ability to change their preferences,  $\gamma'$ . Defining  $A = \delta(1-\alpha)^{1-\rho}(1+r)^{1-\rho}$  for simplicity of notation,

$$\alpha^{-\rho} - \gamma \frac{\delta \alpha^{1-\rho} (1+r)^{1-\rho}}{(1-A)(1-\alpha)^\rho} + \gamma' \frac{\delta \alpha^{1-\rho}}{(1-A)(1-\rho)(1+r)^\rho (1-\alpha)^{1+\rho}} = 0 \quad (45)$$

The first order condition is independent of wealth, and so the consumer's optimal fraction of wealth is not affected by wealth  $W_t$ . This independence confirms that consumption as a fixed share of wealth is an optimal consumption rule.

When a consumer chooses what share of their wealth to devote to consumption, they must balance several competing incentives. First, they get direct utility from consumption. In the current period, they get marginal utility of  $\alpha^{-\rho}$  from consuming, the first term of equation 45. This return is modified by two other factors, however. First, when a consumer adjusts how much of their income they consume they affect their wealth in future periods, a change that is weighted by their present bias and discount rate. This dynamic is captured by the second term. If the consumer increases  $\alpha$ , for example, they consume a higher share of their resources each period. In all subsequent periods, however, this means their wealth will be relatively lower.

In addition, the consumer must account for the fact that their consumption will affect their preferences. The third term in equation 45 reflects this effect: how much the agent decides to consume determines how much they care about future consumption. If they increase the share of resources they consume, for example, they become more indifferent to future consumption, making immediate utility relatively more important.

If  $\gamma' = 0$  and  $\gamma = k$ , then the first order condition simplifies to

$$\alpha^{-\rho} - k\delta \frac{\alpha^{1-\rho} (1+r)^{1-\rho}}{(1-A)(1-\alpha)^\rho} = 0 \quad (46)$$

This is identical to the condition determining optimal consumption for a consumer who faces  $V(W_t) = u(\alpha W_t) + k\delta V(W_{t+1})$ , or in other words a consumer who has fixed preferences. The model therefore nests the standard case when isoelasticity is imposed. In addition, we can rewrite the first order condition as

$$\alpha^{-\rho} (1-\alpha)^\rho - \alpha^{-\rho} (1+r)^{1-\rho} (\delta(1-\alpha) + k\delta\alpha) = 0 \quad (47)$$

Consumption is a function of a weighted average of the effective rates of short and long-run bias,  $\delta$  and  $k\delta$ , as in the general case for present biased consumers (Harris

and Laibson, 2001).  $k\delta$  is weighted by the amount consumed in a given period, and  $\delta$  is weighted by the amount saved, reflecting the relative trade-off between choosing to consume now, or allocating consumption to future periods. Recall however that I have assumed isoelastic utility: further restrictions apply to the solution method in the general case.

### 3.3.3 Uniqueness

A consumer thus selects a level of consumption that maximizes their value of having a given level of wealth. One important difference from the standard model is the number of local optima. In particular, for many parametrizations, multiple locally optimal consumption choices may exist, potentially including a corner solution.

Whether there will be a unique optima depends on the concavity of the problem, as determined by the second order condition (SOC). When the second order condition is strictly negative, the problem is strictly concave, and the level of consumption is a unique maximum. Thus, when:

$$\begin{aligned}
 & -\rho\alpha^{-\rho-1} + \frac{\delta\alpha^{1-\rho}}{(1-A)(1-\alpha)^{2+\rho}(1-\rho)(1+r)^\rho} \left( \gamma'\rho + \frac{\gamma''}{(1-\alpha)(1+r)} \right) \\
 & + \delta \left[ \frac{(1-\alpha)^\rho (1-\rho + \frac{\alpha\rho}{1-\alpha}) - \delta(1+r)^{1-\rho}(1-\rho + \alpha\rho)}{\alpha^\rho(1-\alpha)^{2\rho}(1-A)^2(1+r)^\rho} \right] \\
 & \cdot \left( \frac{\gamma'}{(1-\rho)(1-\alpha)} - \gamma(1+r) \right) < 0 \quad (48)
 \end{aligned}$$

the problem has a unique maximum, and the consumer has a single globally optimal equilibrium consumption level.

The concavity of the problem is a function of the consumer's preferences, the interest rate, their consumption, and the nature of preference change,  $\gamma'$  and  $\gamma''$ . The plasticity of preferences thus presents an important empirical question: intuitively, consumers might find it easier to self-improve once they have already improved their preferences to some extent, or conversely find the final stages of preference change more difficult than the initial steps.

The possible existence of multiple optima reflects two conflicting paths consumption can take. If the consumer increases their consumption in the present, they get more utility immediately. Their present bias also increases, and so they place a higher value on that immediate utility in comparison to future utility. The incentive to consume more therefore increases. This harms future selves by leaving less consumption available for the future, but the consumer's increased present bias means they place less weight on this cost. Each self then makes the same decision, reducing consumption for future selves while having consumption reduced for them by their past selves.

At an extreme, if the function is well-behaved there is generally an optimum where the consumer indulges completely and consumes their entire stock. This occurs because if an agent is indifferent to the future, then it is rational to consume everything at once. Postponing consumption brings future benefits, but the agent places no weight on these future benefits due to his infinite present bias. This solution ceases to exist only when the benefits from moving a marginal amount of consumption to the next period improves preferences sufficiently to make the agent value that marginal amount of consumption plus interest more than the foregone consumption in this period.

If the consumer reduces their consumption in the present, in contrast, they get less immediate utility but anticipate more utility in the future. This reduces their present bias, increasing the value they place on that future utility. They thus consume less in the present period, but leave more for their future selves, who they care about more because of the increased weight they place on future utility.

As a result, both behaviours, consuming more and consuming less, can be self-reinforcing. Two types of optima therefore exist: *self-improvement optima*, in which a consumer reduces consumption as a share of wealth, improves their preferences, and consumption smooths; and *binge optima*, in which the consumer chooses to increase consumption, becomes more present biased, and front-loads consumption. Because in an isoelastic setting consumption is independent of wealth, the consumer consumes a fixed proportion of their wealth in both cases.<sup>5</sup>

If multiple local optima exist, a consumer who optimizes at the margin may reach either point in equilibrium. If the consumer can choose from all local optima, they will select the globally optimal one: as we shall see below, however, this globally optimal

<sup>5</sup> Equilibria in which  $\alpha$  jumps discretely between different locally optimal equilibria of equal utility remain possible. These equilibria are of measure zero, however, and so are disregarded.

consumption path may be unstable. Which optimum is selected in equilibrium then depends on the agents' expectations of behaviour by their future selves, as well as upon initial conditions. If an agent believes future selves will indulge in the high-consumption, high-present-bias, equilibrium, then it is rational for him to do the same; if he is optimistic about the behaviour of future selves, on the other hand, then he will also behave prudently. This is discussed further in Section 3.4.2.

We can thus return to the possible existence of multiple optimal consumption levels. Welfare analysis is necessarily complicated when preferences are not fixed, but it is intuitive that in all periods except the present one, consumers would prefer a self-improvement equilibrium, which makes all future selves better off, to a locally optimal binge equilibrium that makes all future selves worse off. Whether this is optimal for a given self, however, depends on their preferences and what they believe their future selves will do.

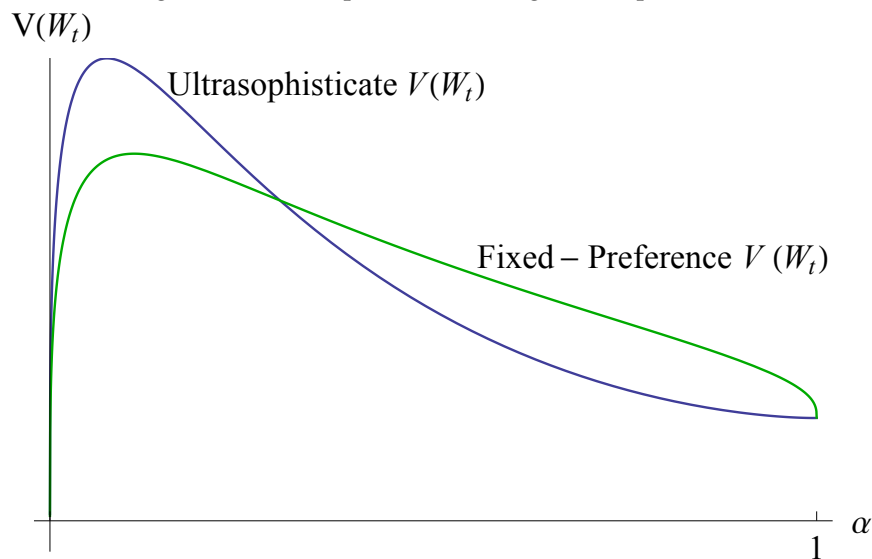
### 3.4 CONSUMPTION PATTERNS

Ultrasophisticated agents therefore account for the effect on their preferences that their consumption decision will have. To illustrate the implications of the model, I consider a parametrized example, with  $r = .1$ ,  $\delta = .9$ , and  $\gamma(x) = 1 - x$ . I first examine ultrasophisticate behaviour relative to that of a benchmark fixed-preference agent, before discussing their incentives for consumption more generally.

#### 3.4.1 *Ultrasophisticates and Fixed-Preference Agents*

Consider two agents, one ultrasophisticated and able to change their preferences, and another with a fixed level of present bias, but otherwise identical preferences. Their optimal consumption paths are therefore determined by equations 45 and 46. Relative to the fixed-preference agent, the ultrasophisticate may prefer to either consumption smooth or front-load consumption, depending on the parameters of their optimization

Figure 6: Consumption Smoothing Ultrasophisticate



The ultrasophisticate's optimum is at a lower level of consumption than a fixed preference agent.

problem. To see this, consider each type's value of wealth  $V(W_t)$  as a function of their consumption level  $\alpha$ . First, assuming  $\rho < 1$ :

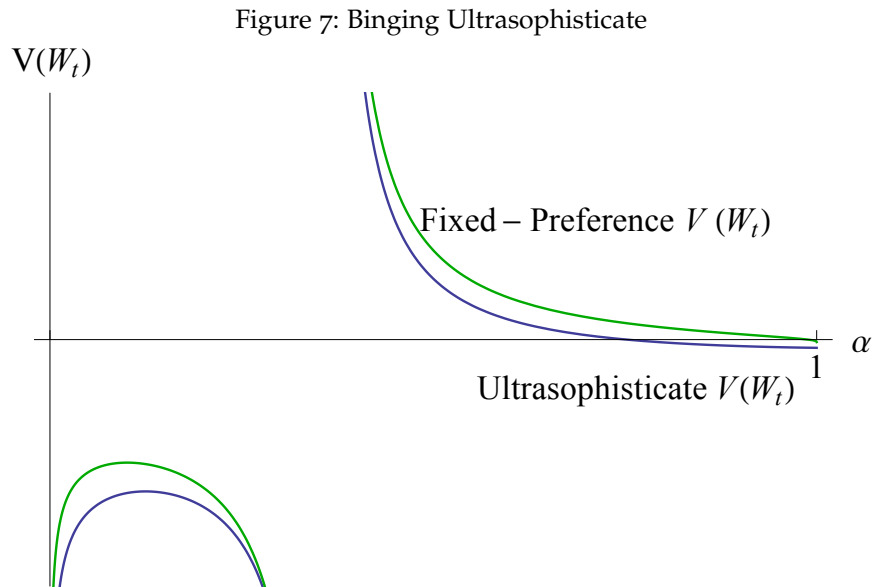
There is a unique interior optimal consumption point for both types of consumer.<sup>6</sup> The ultrasophisticate consumes relatively less at the interior optimum than the fixed preference agent, however – their interior optimum is at a lower value of  $\alpha$ .

Note that the ultrasophisticate also has a local optimum at the corner solution of  $\alpha = 1$ , while the fixed-preference agent does not. If the ultrasophisticate chooses to consume everything, then they become indifferent to the future costs of doing so, and so consuming everything is their best response to being indifferent to the future. If the fixed-preference agent consumes everything, in contrast, they still care about the future harms that result from their behaviour, and so consuming everything is a local minimum. This reflects the discussion in Section 3.3.3.

Since wealth is bound by the budget constraint, if an ultrasophisticate chooses to self-improve and reduce the share of wealth they devote to consumption, in early periods they will consume less in absolute terms than a fixed-preference agent. In later periods, however, their wealth is greater than that of a fixed-preference agent, and so even though they consume a lower share, their absolute consumption may be higher: the ultrasophisticate consumption-smooths.

<sup>6</sup> Estimates of the rate of intertemporal substitution vary widely. A meta-analysis by Havranek et al. of 169 published studies finds significant variation between studies and countries, with a mean elasticity of 0.5 (2013).

If, in contrast,  $\rho > 1$ , the consumer places a higher value on future consumption, *ceteris paribus*. Relative to the fixed-preference agent, the ultrasophisticate may therefore choose to bring consumption forward. This is shown in Figure 7.



The ultrasophisticate's stable optimum is at a higher level of consumption than a fixed preference agent.


Two interior locally optimal consumption points exist for both types (one of which is asymptotic and unstable, but is the global optimum). At the unstable optimum, a small perturbation that lowers  $\alpha$  could incentivize the consumer to move to the lower, stable optimum. Agents have an incentive to select their global consumption optimum in equilibrium, but its instability may make this difficult.

At the stable interior optimum, the ultrasophisticate consumes more than the fixed-preference agent. Increasing their consumption gives the ultrasophisticate more immediate utility from consumption, and reduces the weight they place on future benefits, but also reduces their future wealth. Doing so is optimal if the increased indifference to the future costs of their behaviour compensates for the increased costs.

If the ultrasophisticate does choose to bring consumption forward and consumes at the stable optimum, in early periods ultrasophisticate consumption may be higher than that of a fixed-preference agent, but as their wealth falls in later periods their consumption may also become lower than that of the other agent.

Figure 8 shows how consumption changes over time for an ultrasophisticate and a fixed preference agent.

Figure 8: Consumption over Time



Over time, an ultrasophisticate may consumption smooth or bring consumption forward relative to a fixed-preference agent

If the ultrasophisticate chooses to self-improve, their consumption begins below that of a fixed-preference agent, but eventually overtakes it because their wealth is higher: they consumption smooth. If they choose to binge, in contrast, then their consumption is initially high, but as their wealth falls so does their consumption. At an extreme, the corner solution, an ultrasophisticate may consume their entire endowment in the first period, and consume nothing in all subsequent periods.

Agents balance the cost of distorted consumption in future periods with the loss of utility from reduced consumption in the present period. Depending on preferences and their plasticity, an ultrasophisticate may have one or multiple optimal consumption paths, reflecting self-improvement or binge effects.

Consumption dynamics in the model therefore depend crucially on the nature of the incentives faced by ultrasophisticates. Though increased awareness of their own preferences and how they change might be assumed to always reduce consumption, under some conditions ultrasophisticates will bring consumption forward, enduring worse preferences in order to increase current consumption. This result complements work by O'Donoghue and Rabin (2002), who found that in an addiction context, sophistication could either increase or reduce consumption relative to a naïf depending on incentives.

Ultrasophisticates face a similar choice. They can improve their preferences by reducing their consumption, which reinforces the value they place on the future rewards from doing so: a self-improvement optimum. A binge may also be optimal, however, encouraging them to consume large amounts today in order to reduce how much they care about the future and so be indifferent to the wants of their future selves. Which they choose depends on their incentives for each behaviour, which I discuss next.

### 3.4.2 *Ultrasophisticate Behaviour*

As shown by equation 45, the ultrasophisticate's rate of intertemporal substitution, discount rate, interest rate, and plasticity of preference-change all affect their incentives for consumption. To demonstrate this, let us consider a consumer with  $0 < \rho < 1$ , such that there is only one interior optimum consumption path (though a solution in which the agent consumes their entire wealth may still exist).

As  $\rho \rightarrow 1$ , similar to the standard model, consumers place increasing weight on future benefits, in the limit approaching log preferences, and agents consume increasingly less because they care more about their future self. As  $\rho \rightarrow 0$ , agents become more indifferent to future benefits, and ultrasophisticates consume an increasingly large fraction of their wealth.

Similarly, if a consumer has  $\delta = 0$ , so that they entirely discount the future, then they consume everything immediately, and their optimum is  $\alpha = 1$ : this lowers their future utility, but they are indifferent to this, as in the standard model. As the discount rate increases, the agent cares more about the future, and so they have an incentive to reduce their present bias and consume less: as  $\delta \rightarrow 1$ ,  $\alpha \rightarrow 0$ . Unlike in the standard model, present bias and discount rates are therefore correlated. Consumers with a high exponential discount rate will have a high rate of present bias in equilibrium, and those with a low exponential discount rate will have a low rate of present bias. Differences in future discounting can therefore be self-reinforcing. In terms of welfare, this suggests that consumers that exponentially discount the future highly may have the effect compounded by also being present biased, a potential concern for policymakers.

High interest rates increase the return to saving, and so increase the incentive to reduce consumption. As interest rates increase, therefore, the fraction of wealth devoted to consumption falls. As Tyson (2008) discusses, an increase in interest rates also increases total wealth, potentially affecting a consumer's absolute consumption. Here, however, since consumption is independent of wealth, this second effect does not change the fraction of wealth consumers devote to consumption.

Specific to this model is the role of the plasticity of preferences in equilibrium. If we assume the form  $\gamma = 1 - A\alpha$ , as  $A \rightarrow 0$ , consumers must sacrifice more immediate consumption in order to improve their preferences, or equally do not significantly affect their preferences when they do consume more. As  $A$  increases, small changes in consumption can have a large effect on preferences. If  $A = 0$ , then preferences are fixed.

As  $A$  increases, the incentives for more extreme behaviour increases. For a consumer that self-improves, the optimal interior level of consumption is decreasing in  $A$ : as  $A$  increases, it is easier for the consumer to change their preferences, and so the incentive to abstain from consumption and self-improve further increases. For a consumer who is binging, as  $A$  increases they are more indifferent to the costs of their behaviour for a given level of overconsumption, and so they overconsume further. More generally, the form of  $\gamma$  will help determine whether there are multiple locally optimal values of  $\alpha$  and whether solutions are interior.

Sacrificing present consumption to reduce present bias is therefore welfare-enhancing when:

1. The consumer cares more about future consumption ( $\delta, \rho$  high)
2. There is a higher return to saving ( $r$  high)
3. Changing consumption has a large effect on present bias ( $\gamma'$  is large and negative)

Conditions 1 and 2 affect the marginal benefit of a reduction in  $\gamma$ , and Condition 3 affects the marginal cost of improving preferences in terms of consumption. Put differently, these conditions exacerbate the utility cost of distorted decision-making or reduce the cost of eliminating the distortion, encouraging an agent to sacrifice some present consumption in order to reduce those distortions. Ultrasophisticate agents must balance the

benefits of consumption today against the cost in terms of consumption tomorrow and the effect on their preferences.

This provides some insight into the behaviour of agents with non-isoelastic utility. Agents might have increasing  $\rho$  over time, for example (which in the presence of risk would imply CARA utility). As a result, consumption will generally be falling over time. Intuitively, think of a young man who inherits a vast fortune. Initially, he spends large amounts of money, and adopts an extremely high-consumption lifestyle. As his fortune dwindles, however, he starts worrying about the possibility of running out entirely, and so begins to reduce his consumption and acclimatize to a low-consumption lifestyle, building better preferences.


Implicit in the discussion above are the ultrasophisticate's expectations about the future. If the consumer expects behaviour from their future selves that is different from what they would choose, then their incentives change. To see this, consider ultrasophisticate behaviour if in all subsequent periods the consumer is forced to consume a fixed proportion of their wealth,  $\bar{\alpha}$ . The consumer's first order condition is then

$$\alpha^{-\rho} - \gamma \frac{\delta \bar{\alpha}^{1-\rho} (1+r)^{1-\rho}}{\left(1 - \delta (1+r)^{1-\rho} (1-\bar{\alpha})^{1-\rho}\right) (1-\alpha)^{\rho}} + \gamma' \frac{\delta \bar{\alpha}^{1-\rho}}{\left(1 - \delta (1+r)^{1-\rho} (1-\bar{\alpha})^{1-\rho}\right) (1-\rho) (1+r)^{\rho} (1-\alpha)^{1+\rho}} = 0 \quad (49)$$

Now, the current self's consumption decision affects how many resources will be left for the subsequent selves to consume, but the subsequent selves consume a fixed proportion of wealth in all future periods. We can thus graph  $\alpha(\bar{\alpha})$ ; the consumer's optimal consumption as a function of the anticipated consumption of their future selves.

As anticipated future consumption increases, the ultrasophisticate has an incentive to also increase their consumption: anticipating imprudent behaviour, the ultrasophisticate also behaves imprudently. This reduces the resources to be used inefficiently in the future and reduces the value they place on the future, further reducing the expected cost of the bad behaviour. For very low consumption levels in the future, the ultrasophisticate has an incentive to increase their consumption: because future selves are being too prudent, resources are going to waste, and the current self increases consumption.

Figure 9: Optimal Consumption and Expected Consumption



If an ultrasophisticate anticipates their future selves consuming too little or too much, their own incentive to consume increases.

Agents expecting good conditions and well-behaved selves in the future have an incentive to be well-behaved in the present, while agents expecting low interest rates or high exponential discounting in the future will expect their future selves to binge, incentivizing them to do the same. Agents thus also gain value from being able to commit their future selves to a specific level of consumption, since their own consumption is affected by what their future selves will do.

In the standard model, a consumer expecting a negative shock in the future has an incentive for precautionary saving, in order to consumption smooth. Here, there remains an incentive to save in order to consumption smooth, but consumers also have an incentive to increase consumption in order to forget about the future suffering. If, for example, the consumer expected negative utility in future periods, then in the standard model consuming everything immediately would lower their future utility even further. Here, however, though their future utility would still be lowered, they would be indifferent to this because of their increased present bias, making consuming everything an optimum. If the expected future utility were sufficiently negative, consuming everything immediately might be the unique optimum.

Ultrasophisticates thus face similar incentives to the standard, fixed-preference consumer, with the addition of the importance of the plasticity of preferences in determining their behaviour and how they respond to future events. The introduction of preference endogeneity provides an incentive for their behaviour to deviate from that of a fixed preference agent.

### 3.4.3 *Pseudorationality*

If the level of present bias varies between agents in equilibrium, policymakers may have an interest in identifying these types. Identifying agents' present bias from consumption may be difficult, however. Agents that choose to have non-zero levels of present bias are *pseudorational*: they act as if they were an agent that only exponentially discounts the future, but who had a higher exponential discount rate.

Let us compare two agents, one of whom is not present biased and the other of whom is. I define  $\delta_{\gamma=1}$  as the exponential discount rate for a non-present biased consumer, and  $\delta_{\gamma<1}$  as the discount rate of a consumer with an interior level of present bias. When the first order conditions of both agents are the same, the two types consume identically. This occurs when

$$\frac{\delta_{\gamma=1}}{\delta_{\gamma<1}} = \frac{\gamma'(\alpha^*) - \gamma(1+r)(1-\alpha)(1-\rho)}{\gamma'(1) - (1+r)(1-\alpha)(1-\rho)}$$

So optimal consumption is the same for both types if they have levels of present bias  $\delta_{\gamma=1}$  and  $\delta_{\gamma<1}$ . Whether  $\delta_{\gamma=1}$  is less than or greater than  $\delta_{\gamma<1}$  will depend on the concavity of the preference change. The agent's present bias,  $\gamma$ , distorts his behaviour only insofar as he will act as a rational agent with different preferences. This pseudorationality need not imply that these agents are entirely inseparable, however: eliciting discount rates in situations when only  $\delta$  is relevant, through experiments or by offering different future income flows or levels of salience, may allow for identification of an agent's type.

Given endogeneity in other variables, consumers might have an increased incentive for the complete elimination of present bias. If an agent expected a higher rate of interest in the future and lowered his present bias in response, for example, such a decision

could actually increase future rates of interest, creating a self-fulfilling prophecy, while agents who pessimistically didn't bother to self-improve could actually induce the low rates of return they expected.

### 3.5 WELFARE AND POLICY IMPLICATIONS

A challenge with welfare analysis given non-stationary preferences is selecting which preferences agents should be evaluated by, since multiple alternatives exist. I do not propose to give a full treatment to these issues here, though they are substantial: for a longer discussion, see Bernheim and Rangel (2009) or von Weizsacker (2013).

An intuitive welfare measure to compare fixed-preference agents and ultrasophisticates is  $\tilde{V}(W)$ , which represents how an agent would weight the utility from period  $t$  when it is not the present. Since agents give the same weight to the utility in all preceding periods,  $\tilde{V}(W)$  lends a stability to the analysis. I do not introduce this measure as a definitive utility comparison, but merely as one possible source of comparisons between fixed-preference and ultrasophisticate utility.

By the standard of  $\tilde{V}(W)$ , an agent is best off when she consumes as if she is not present biased. This implies that agents are better off when they choose to improve their preferences and reduce consumption in the current period, and worse off when they succumb to a binge and choose to overconsume in the present period. This distinction is intuitively plausible, given the popularity of self-help books on how to resist temptation.

It follows that agents are better off with endogenous preferences only when that ability encourages them to self-improve. If, in contrast, it leads to overconsumption in the present period, they are worse off. Being more sophisticated is only a blessing when it encourages virtue: otherwise, agents would be better off unable to control their preferences, and having less agency. Expectations about the behaviour of future selves, and expectations about the parameters that drive that behaviour, are crucial in determining equilibrium behaviour, and welfare, for ultrasophisticates. Given the importance of expectations, commitment devices present a potentially useful option. If a consumer decides to binge because they anticipate irresponsible future selves, then being provided with the opportunity to commit their future selves to a given behaviour could help encourage them to self-improve instead.

Three basic areas of policy interest relate to the model. First, because multiple equilibria may exist, policies that change which equilibrium is reached can make a large difference to welfare. In addition, anticipated negative shocks to consumption may also incentivize negative behaviour in the present. If a consumer expects to be forced by sudden poverty to consume their entire endowment with large costs to welfare, for example, then they may choose to do so immediately in order to reduce the weight they place on those future costs, making it even harder for them to save and emerge from poverty. Early interventions to ensure agents get into the self-improvement consumption path may therefore be a promising area of research, as may reducing or compensating for shocks.

Second, beliefs about whether preferences are fixed or variable (reflecting whether a consumer is ultrasophisticated) can change incentives. This implication parallels work done by Mangels et al. (2006) on the difference between a 'fixed' and 'growth' mindset; children can believe that success stems from either innate ability or hard work and learning, and their beliefs play a role in determining their success in school. Policymakers may be able to improve outcomes by encouraging students to adopt the growth or variable preferences mindset, but only if the conditions are such that they reach the self-improvement equilibrium if they do so.

Finally, though not directly demonstrated by the model, it is intuitive that because agents are forward looking, policies risk having a crowding out effect on incentives: students might be worse off, for example, if paying them for grades or giving them a self-commitment device removes the self-improvement benefits. This effect may explain the relatively high drop-out rates of graduates of demanding charter schools, like the Knowledge Is Power Program (KIPP, 2011), but it is not clear coercion eliminates the possibility of preference change. More work is needed to better understand how to use these ideas for policy.

### 3.6 DISCUSSION AND CONCLUSION

This paper proposes a model of human behaviour founded in the psychology literature, in which preferences are fluid and depend on the allocation of consumption across

time made by the consumer. I consider the implications for the behaviour and strategies of agents, as well as equilibrium behaviour.

Two types of optima emerge, self-improvement optima and binge optima, due to opposing incentives. Some agents may choose to defy their own preferences in order to self-improve and consume more efficiently, while others indulge in the present in the knowledge that they become increasingly indifferent about their damaged future. Consuming can begin a cycle of indulgence, while resisting temptation can compound into virtuous behaviour.

In equilibrium, consumers may reach either a low-consumption or a high-consumption equilibrium – including one in which all resources are consumed immediately and the agent becomes indifferent to future consequences – which are local optima. Which equilibrium is selected depends upon preferences and the agents' beliefs about their behaviour and conditions in future periods. Agents who anticipate their future selves choosing a high-consumption equilibrium will do the same, even if they would otherwise prefer all selves choose the low-consumption equilibrium. Agents can either choose to consumption smooth and self-improve, or front-load consumption and potentially become trapped in a negative pattern of behaviour where they find it hard to care about the future, as we see in dependency on alcohol, drugs, or even video games. Endogenous preferences may thus provide some explanation as to why we observe attempts to resist temptation, when in a simplistic model agents would simply act to satisfy their preferences - they are aware they are better off if they can resist indulging in temptation because doing so improves their preferences.

These effects provides a potential role for policy interventions. Early opportunities to exercise self-control, for example during education, may help determine the equilibrium path and so affect welfare. Shocks to behaviour, whether from policy or exogenous, may also create persistent patterns of behaviour by changing the optimum which the consumer achieves in equilibrium. In addition, consumer expectations about the future play a role in determining their present behaviour, and so also provide a useful target for policy and a potential need for commitment devices. More broadly, modern policy interventions, which increasingly focus on long-term lifestyle changes such as reducing obesity or increasing savings rates, may need to consider how to change both behaviour and preferences.

These results underscore the difficulty policymakers face in designing policies when preferences are not fixed. As von Weizsacker (1971) has suggested, in some situations we might believe preferences to be flawed from myopia or lack of imagination, among other possibilities. The obvious question, however, is who decides whether preferences are inadequate. This paper presents a model of self-improvement, in which agents may choose to improve their own preferences.

The inclusion of endogenous preferences can add explanatory power to models, both in predicting individual behaviour and describing optimal behaviour. If we accept that preferences are not exogenously fixed, however, we also provoke an important discussion: when economics discusses maximization, what exactly should agents be maximizing? This paper has not addressed that issue directly, yet it is an issue of significant concern. Our preferences give us goals and direction, without which agents are aimless, yet the question of how those preferences are established is critical. It is not enough to ask what people want: we must also ask why they want what they want, and whether it could or should be changed.

## BET YOU CAN'T EAT JUST ONE: CONSUMPTION COMPLEMENTARITY AND 'SELF'-CONTROL

---

### 4.1 INTRODUCTION

The idea of temptation implicitly suggests the existence of multiple selves: that one part of the consumer is struggling against the desires of another. Since temptation lies at the heart of many of the decisions consumers face, whether saving for retirement, pursuing education, or trying to lose weight, understanding how consumers can resist and account for it presents an important policy challenge. Temptation also has the potential to predict and explain consumer behaviour in a variety of economic settings, from personal rules for resisting temptation to free samples in supermarkets, and can present an important microfoundation for long-term behaviour in the form of habits or addictions.

This paper explores the idea that consumption of a tempting good can temporarily induce a 'hot' state in a consumer. In a hot state, consumers behave as if they had different preferences from when they were in a cold state: consumption is thus a game between hot and cold selves, where each self is the consumer in a different state. Choosing to indulge can induce a hot state that then makes subsequent indulgences more likely, and prior decisions to consume can have a causal effect on future ones, in contrast to the typical prediction of diminishing marginal returns. Eating a single square of chocolate from the bar – or taking a single drink, or cheating just a little bit – might well be optimal from a cold-state perspective, but if doing so enters the consumer in a hot state, eating only a single square may not be feasible: instead, the consumer is forced to choose between eating no chocolate or the entire bar. To understand the effect, I first introduce a model of its implications, and then test for its existence in the lab. Reflecting that the consumer has different preferences in the two states, I will refer to the consumer when in a cold state as the cold self, and when in the hot state as a hot self.

I first consider behaviour when both selves are naïve or when only the cold self is sophisticated, before extending the model to consider behaviour when both selves are sophisticated. I find that naïve consumers who do not anticipate a hot state overconsume

relative to their cold-self optimum: the glutton who is repeatedly surprised by how much they eat, for example. Sophisticates, however, may adopt personal rules – strict rules restricting individual choice (Ainslie, 1992) – that change their optimal strategy for consumption because they anticipate the costs that result from entering a different state. The cold self adopts rules restricting consumption: in response, the hot self adopts rules increasing consumption, and so periods of underconsumption can be followed by periods of binging once the consumer enters a hot state, both with the goal of avoiding a state-change. Multiple equilibria can exist – one in which selves expect the other to deviate widely from their optimum, and so also deviate widely, and another where both selves expect only minor deviations – with implications for welfare. If sophistication can vary between selves, then a sophisticated self would prefer the other self be naïve. The larger the shift in preferences between selves, and the more weight consumers place on the future, the greater the costs to inducing the other state, and so personal rules are particularly valuable. Consumers who expect more indulgent hot selves thus have a larger incentive to restrict behaviour while in the cold state, but will also binge more once in the hot state.

This motivates several hypotheses about behaviour in the lab, and I conduct an experiment testing for the existence of a short-term causal effect of consumption on indulgence. In particular, I test one of the predictions of the theoretical model: whether inducing subjects to consume a tempting good – a piece of chocolate – can increase the total amount of the tempting good consumed. Eating a piece of chocolate once, the theory suggests, might induce a hot state that increases the likelihood that consumers will succumb to temptation later on. I find no evidence for such an effect generally, but do find some support for an effect on consumers on a diet, who the theory suggests should be particularly likely to demonstrate such an effect: cold selves who anticipate particularly indulgent hot selves have more incentive to restrict their own consumption, and so to diet. This effect exists even when the initial indulgence was imposed externally, rather than the result of choice. Though the focus of the experiment is to test for the existence of such an effect, I also briefly discuss possible explanations for it, several of which are ruled out by the conditions of the experiment.

I thus build on the literature of hot and cold states – in which hot states are an affected state characterized by strong emotions, hunger, exhaustion, and other factors,

while a cold state is a more emotionally detached, ‘rational’ state (Loewenstein, 2005) – to consider consumer behaviour when consumption of a tempting good enters them into an affected state. This could reflect either a change in preferences between the two states, or equivalently the existence of a fixed cost to consumption of the tempting good – either material or psychological, such as a loss of self-image – which diminishes once the consumer indulges. Either of these effects could serve as a mechanism for the transition between a hot and cold state. Consumption as an affect may also help describe and explain other more specific affects, such as drunkenness or anger, which result from consumption of particular types of goods.

Several models of multiple selves exist. In 1981, Thaler and Shefrin introduced a planner-doer model, in which a long-run self with rational preferences best responds to the short-run doer self (Thaler and Shefrin, 1981). More recently, Loewenstein and O’Donoghue discuss deliberative and affective systems and Fudenberg and Levine consider long-run and short-run selves, while Bernheim and Rangel link hot and cold selves to cues (Bernheim and Rangel, 2004; Fudenberg and Levine, 2006; Loewenstein and O’Donoghue, 2004). Most of the literature, however, focuses on a rational long-run self that responds to a naïve short-run self that always seeks to indulge. Here, I allow both selves to best respond to the behaviour of the other self, instead of assuming that one self is entirely naïve or ignores the other’s behaviour. Though it is possible to rate the preferences of the cold self over those of the hot self, it is not obvious *ex-ante* that this is always reasonable, and so my model also does not inherently prioritize one self over the other. I also allow disagreement between the two selves that may reflect matters of taste, as well as potentially preferences about the future.

The existence of hot and cold selves can both explain economic behaviour and provide important lessons for policy design. For the individual, personal rules that restrict consumption even when immediate marginal utility is positive can be welfare-improving for the cold self by reducing the likelihood of entry into a hot state. Consumers who anticipate the change in preferences between hot and cold selves have an incentive to resist temptation while in the cold state because they see indulging once as part of a larger choice bundle that is associated with being in a hot state.

Individuals also benefit from optimism about their own behaviour. Because multiple equilibria may exist, selves who expect the other self to deviate widely from their

optimum will in turn also deviate widely, reducing welfare for both selves. A more optimistic consumer, in contrast, can achieve the more moderate equilibrium, increasing welfare. For firms, the potential for hot selves also increases the value of offering free samples or trials of products that consumers find tempting.

For policymakers, the existence of multiple equilibria means there may be a welfare-improving role for them to play. If a policymaker chooses to elevate the preferences of one self over another, there may be value in the provision of commitment devices, so that the self can achieve a time-consistent equilibrium. In addition, the experiment highlights the fact that imposing a behaviour on consumers can affect their future decisions. Policymakers must therefore take care that when policies or circumstances force a behaviour, they do not lead to negative consequences in the future. If a consumer is forced to indulge in a tempting behaviour due to circumstance – buy fast food because of how cheap it is on a per calorie basis, for example – then that behaviour may induce a hot self and snowball into a larger effect, with a significant impact on welfare. Policies that promote abstention and help consumers resist early temptations, in fields from work with at-risk youth to obesity and education, may help consumers avoid entering hot states altogether.

In Section 4.2, I introduce a theoretical analysis of the interaction between hot and cold selves, considering behaviour both when the hot self is naïve (Section 4.2.2) and when both selves are sophisticated (Section 4.2.3). I then discuss some experimental work testing for the transition between states as a result of indulgence in Section 4.3. Section 4.4 concludes.

## 4.2 THEORY

Economic models of consumption and saving often focus on the long term: periods represent working and retirement, or years in a life, across which a rational consumer must allocate his or her lifetime endowment (Modigliani and Brumberg, 1954). Consumption persistence – referred to in a general sense as habit formation – is frequently observed in this setting, and used to explain phenomena in both a micro- and macroeconomic

context, including addiction, equity premiums, the link between economic growth and savings, and aggregate spending patterns, among others.<sup>1</sup> In the short-run, however, consumption persistence is less well understood.

Loewenstein (2005) highlights the importance of affective states in decisions: that is, the effect emotions and psychological states have on our mental functioning and choices. In 'hot' states, when we are hungry, tired, angry, or otherwise affected, we have a tendency to behave myopically or otherwise irrationally. In 'cold' states, people are better able to act according to the traditional model of rationality, but may not appreciate how being in a 'hot' state will affect their preferences and behaviour. This model of behaviour draws on a perspective from psychology which originally introduced hot and cold selves, as in Metcalfe and Mischel (1999).

This effect has been shown in a range of domains, from thirst (Van Boven and Lowenstein, 2003) to addiction (Badger et al., 2007) and sexual arousal (Ariely and Loewenstein, 2006). I suggest a new possible affect: consumption of some types of goods. Indulging in such a good induces intertemporal complementarity in consumption, in contrast to the standard, cold-state substitutability in consumption that is typically predicted in the long-run. Though typically this is thought of as a change in preferences between selves, it could equivalently be thought of as a type of fixed cost to indulging in temptation, either material, such as getting unpleasantly sticky hands from eating candy, or psychological, such as a loss of self-image. Before indulging, the consumer faces a high fixed cost, but once they have indulged further indulgences impose a lower cost and so the net benefit of indulging again temporarily increases. Either mechanism, preferences or fixed costs, could thus serve as a mechanism for the model. In either case, inducing a period of complementarity in consumption can then create what I will refer to as a *temptation effect*: indulging once in a good can have a short-term causal effect on future consumption.

I introduce a model of short-term consumption persistence based on hot and cold states. When a sufficiently large temptation strikes a consumer, this induces consumption and so begins a hot state, in which the marginal utility of consumption increases. In subsequent periods the consumer is therefore more likely to consume again. This affected state is temporary: after the hot state passes, the consumer's preferences re-

---

<sup>1</sup> See, for example, Becker and Murphy (1988); Carroll, Overland and Weil (2000); Constantinides (1990); Crawford and Polisson (2014); Fuhrer (2000)

turn to their cold levels. This could reflect a change in preferences, as the consumer musters their mental resources to resist temptation again, or equivalently the return of a fixed cost to indulgence. This captures the frequently observed cycle of attempts to resist temptation, a ‘fall off the wagon,’ and then another attempt to resist.

Extreme shock models describe systems that face repeated shocks of varying magnitude, and which break down when they receive a shock of magnitude above a given threshold.<sup>2</sup> In my model, consumers indulge when a shock surpasses a threshold that reflects their desire to resist temptation. I endogenize the threshold based on their utility maximization problem: depending on their ability to foresee the consequences of their decision, consumers may be more or less motivated to resist temptation. The model is thus related to studies of unemployment, for example, in which consumers optimize across employed and unemployed states (Diamond, 1982).

#### 4.2.1 *Model*

Consider a consumer facing a binary consumption decision – eating a square of chocolate, for example – over a series of periods. Choosing to consume gets them net utility of  $u_t$  beyond what they would get if they abstained or consumed some alternative. This net utility can be positive or negative: if they abstain, they get zero utility, and so they only want to consume if indulging gives them more utility than abstaining. The utility from consumption varies over time, as indexed by  $t$ : each period, the consumer receives an idiosyncratic taste shock that determines their value of  $u_t$ .

In addition, utility from consumption depends on the state of the consumer. Each period, either a cold self or hot self is in control, depending on the consumer’s decision in the previous period. For the hot self, choosing to consume gives utility of  $u_t + h$ ; for the cold self, it gives only  $u_t$ . For simplicity, I assume this is an additive effect: since the consumer faces a binary consumption choice, this increases their marginal utility of consumption.

---

<sup>2</sup> Examples in an economic context include the random walk models used to describe financial and stock market phenomena, and models of gambler’s ruin (Bachelier, 1900; Fama, 1965; Grinstead and Snell, 2003)

The consumer's state is endogenous to their choices. In particular, choosing to consume a strictly positive amount in period  $t$  induces a hot state in period  $t + 1$ . In the two-period model, this is assumed to persist for one subsequent period. After the hot state ends, the cold self once again predominates.<sup>3</sup> In the infinite-horizon game between two sophisticated selves, a hot state persists until the hot self chooses to abstain, and a cold state persists until the cold self chooses to indulge.

For a cold self with discount rate  $\delta$  and facing an infinite horizon, utility is therefore

$$u_t = \sum_{t=0}^{\infty} \delta^t [c_t \cdot (u_t)] \quad (50)$$

while for the hot self, utility is

$$u_t = \sum_{t=0}^{\infty} \delta^t [c_t \cdot (u_t + h_t)] \quad (51)$$

where  $c_t$  is a binary variable of zero or one, representing the choice of whether to indulge or abstain. The idiosyncratic taste shock represents the changing appeal of the good over time. Intuitively, we could consider waves of temptation of varying strength arriving every few minutes to a consumer's willpower and ability to resist, as in Figure 10.<sup>4</sup> Each shock is an independent random variable drawn from a known distribution: individuals do not know how tempting they will find a good in subsequent periods, but they know the distribution of taste shocks. Each period is conceptualized as relatively short, lasting only a few minutes.

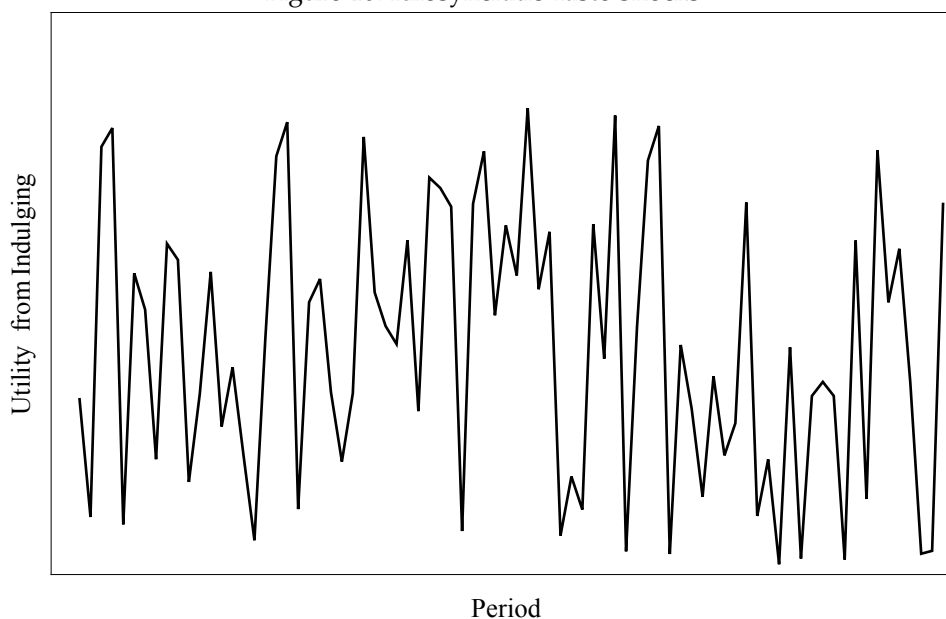
Each period,  $u$  is drawn from a known distribution with density  $p(u)$  and cumulative distribution function of  $P(u)$ . The net utility of consumption may be positive or negative.

Each self is either sophisticated or naïve. A sophisticated self makes a full optimization decision across all periods: a naïve self does not take into account the possibility that the decision in period  $t$  will affect period  $t + 1$ , and so indulges if the marginal util-

<sup>3</sup> Persistence for one subsequent period is chosen for simplicity; the analysis extends straightforwardly to a hot state that endures for multiple periods.

<sup>4</sup> This shock reflects both fixed traits about an individual as well as the environment or context in which the decision is made. Some consumers might find a product more tempting, or believe themselves less able to resist a given temptation, for example (Job, Dweck and Walton, 2010). Such consumers receive shocks that first order stochastically dominate the shocks received by consumers who do not find a good tempting, or believe they are able to resist the temptation.

Figure 10: Idiosyncratic Taste Shocks



ity of doing so is greater than the marginal cost in the current period.<sup>5</sup> I first consider a simple two-period model in which the cold self is sophisticated, and the hot self is naïve. I relax this assumption and explore the implications if hot selves also optimize in Section 4.2.3.

Other factors could exacerbate or moderate the effect of a hot state. Diminishing marginal returns, for example, might gradually reduce the size of the taste shocks as the good loses its appeal. Asking a consumer to perform a willpower-demanding task beforehand or distracting the consumer while he or she chose might make it harder for consumers to resist temptation, while teaching techniques to resist temptation might make it easier (Baumeister et al., 1998). In order to focus on the implications of entry into a hot state directly, however, I abstract away from these factors.

The consumer thus faces an intertemporal game between the hot self and cold self. A consumer in a hot state is identical to one in a cold state, except that they now place a different weight on the value of consumption:  $u_t + h$  instead of  $u_t$ . If the self chooses to consume, they get utility from consuming; if they abstain, they get nothing. If they are sophisticated, each self cares about the utility they get from being in that state, and the utility they expect from being in the other state as judged by the preferences of the

<sup>5</sup> Loewenstein (2005) argues that naïvety on the part of selves is relatively common: consumers have a hot-cold empathy gap, and so find it difficult to predict the preferences of the other self.

current state. For the cold self, indulging provides positive utility but may also induce a hot state in which their future decisions are no longer optimal based on their cold preferences.

The model as presented assumes that entering a hot state is a binary effect: the consumer enters a hot state with the initial indulgence, but  $h$  does not increase further. An alternative would be that  $h$  increases each time a consumer indulges (the state gets 'hotter'). This second type could then lead to a cascade effect, as each indulgence makes a subsequent indulgence more likely, which in turns makes further indulgences even more likely. Even a binary hot and cold state has implications for total consumption and the likelihood of indulgence, however.

#### 4.2.2 *Two Periods, Naïve Hot Selves*

I consider first a two period setting, where the hot self consumes if the marginal utility of doing so in that period is greater than the marginal cost. The hot state persists for one period subsequent to an indulgence by the cold self, and then ends, regardless of the choice of the hot self.

The cold self maximizes utility over both periods (equation 50), and so chooses to indulge when the anticipated benefits of doing so are greater than the anticipated costs. Their decision will therefore be a function of the realized taste shock in that period, which affects the utility of consumption. For some level of taste shock, the cold self is indifferent between consuming and indulging. For the cold self, I denote this level as  $\bar{u}$ : for the hot self, it is  $\underline{u}$ . For simplicity, I model the consumer as choosing this threshold directly, but the threshold reflects an implicit binary choice to consume or abstain on the part of the consumer in each period.

Consider first the hot self's decision. Since hot selves act naïvely, they consume in period  $t$  if the immediate marginal benefit is greater than the marginal cost, which holds if

$$u + h > 0 \tag{52}$$

They are therefore indifferent to consuming or abstaining when

$$u_t = \underline{u} = -h \quad (53)$$

The probability of the hot self consuming in period 1 is therefore  $\Pr(u_1 > -h) = 1 - P(-h)$ .

The cold choice depends on the relative sophistication of the cold self. A naïve cold self does not anticipate that their behaviour will affect their future preferences, and so simply maximizes utility in the stage game. A sophisticated self, in contrast, anticipates that they will have different preferences in a different state, and optimizes accordingly.

#### *Naïve Cold Self Behaviour*

Because naïve cold selves do not anticipate a change in their preferences, they decide to consume in a similar manner to the hot self: if the idiosyncratic taste shock is sufficiently large to make the utility from consuming greater than the cost of doing so. Since their decision is discrete, they consume whenever the utility they get from consuming is greater than the utility they get from abstaining, and face an identical problem in every period in which they are in a cold state. Naïve cold selves are therefore indifferent between consuming and abstaining when:

$$u = \bar{u}^N = 0 \quad (54)$$

The cold self abstains if the taste shock is sufficiently small that the utility they get from consuming is less than  $\bar{u}^N$ . So, for a naïve cold self the probability of consuming in a period is the probability of a shock in excess of a threshold of 0,  $\Pr(u > 0) = 1 - P(0)$ .

Naïve choices are independent of their discount rate,  $\delta$ , or the scale of preference difference between the two states,  $h$ , because they do not foresee future consequences to their decision. Naïve cold selves may therefore find themselves surprised by how much they indulge, because they do not anticipate that indulging once will induce hot-self preferences.

### *Sophisticated Cold Self Behaviour*

A sophisticated consumer, in contrast, anticipates entry into a hot state and includes this in their decision. They therefore weight their decision by the anticipated effect on future welfare. Because being in a hot state increases marginal utility, the hot self might consume when it is not optimal for the cold self to do so, meaning inducing a hot state is costly for the cold self state. Since there are only two periods, a sophisticated cold self chooses identically to a naïf in the second period, because the hot state cannot persist beyond the final period, and so inducing one imposes no costs.

First, let us define  $v_t^S$  as the expected value of being in state  $S$  in period  $t$  to the cold self. Since the cold self chooses as a naïf in the second period, it follows that:

$$v_2^C = \Pr(u_2 > 0) E(u_2 | u_2 > 0) \quad (55)$$

Similarly,

$$v_2^H = \Pr(u_2 > -h) E(u_2 | u_2 > -h) \quad (56)$$

Which can be rewritten as:

$$v_2^H = \Pr(-h < u_2 \leq 0) E(u_2 | -h < u_2 \leq 0) + \Pr(u_2 > 0) E(u_2 | u_2 > 0) \quad (57)$$

As is intuitive, it follows that  $v_2^C > v_2^H$  if the two selves disagree ( $h \neq 0$ ) and the probability that a shock could occur that would make the hot self consume when it would not be optimal for the cold self to do so is non-zero ( $\Pr(-h < u_t \leq h) \neq 0$ ).

Expected utility for a sophisticated cold self is thus the combination of their utility given a shock is sufficiently small to convince them to abstain – in which case they get nothing immediately, and remain in the cold state in the subsequent period– and their utility of a shock is large enough to convince them to indulge – so that they get  $u$  immediately, and enter a hot state.

$$U = P(\bar{u}) (0 + \delta v_2^C) + \int_{\bar{u}}^{\infty} (u + \delta v_2^H) p(u) du \quad (58)$$

At the cold self's optimum,  $\frac{dU}{d\bar{u}} = 0$ . It follows they are indifferent between indulging and abstaining when:

$$\bar{u} = \delta (v_2^C - v_2^H) \quad (59)$$

Since:

$$v_2^C - v_2^H = (P(0) - P(-h)) (E(u_2 \mid -h < u_2 \leq 0)) \quad (60)$$

it follows that:

$$\bar{u} = \delta \int_{-h}^0 u p(u) du \quad (61)$$

$\bar{u}$  measures the incentive for a sophisticated cold self to resist consuming relative to a naïf. The total cost of consuming in period 1 increases from nothing to  $\delta (v_2^C - v_2^H)$ , because the cost of consuming includes both the direct cost as well as the expected cost of inducing a hot state as opposed to staying in a cold state.

A few implications emerge from this simple model. First, sophisticates and naïfs consume identically when  $v_2^C = v_2^H$ , or equivalently when  $h = 0$ : the consumer has identical preferences in both states. They may also consume identically in a corner solution, such as if both selves always prefer to abstain. Otherwise, since  $v_2^C > v_2^H$  a sophisticated cold self requires a higher idiosyncratic shock than a naïf in order to decide to consume, because they account for the effect of being in a hot state in the second period. A sophisticated cold self may be willing to sacrifice considerable consumption to avoid inducing a hot state.

This willingness to sacrifice present consumption in order to resist temptations in the subsequent period is increasing in  $\delta$ : the more the cold self cares about the future, the more they will resist temptation. In the standard model, consumers resist tempting goods because they anticipate long-term consequences from indulgent behaviour such as obesity or addiction. Here, caring about the future reduces the appeal of consumption because it increases the cost of anticipated irresponsible behaviour in the hot state. If the consumer is totally indifferent to the future, in contrast, there is no reason not to consume identically to a naïf.

Since for  $h > 0$ ,  $\frac{d\bar{u}}{dh} > 0$ , the amount of consumption a cold self sacrifices to avoid inducing the hot self is also increasing in  $h$ : the more the two selves disagree, the greater the cost to inducing a hot self. An increase in  $h$  increases the range of possible shocks for which the two selves disagree, and so increases time inconsistency and the cost of inducing a hot state. This range over which the two selves disagree,  $(-h, 0]$ , determines the scale of the time-inconsistency of preferences.

The likelihood of a naïve cold self consuming in period 1 is therefore  $1 - P(0)$ , while the likelihood of a sophisticate consuming is  $1 - P(\delta(v_2^C - v_2^H))$ : sophisticates are less likely to consume. A sophisticated consumer is better off than a naïf by the principle of revealed preference, since sophisticates could choose to consume identically to a naïf, but do not. Both types, however, are worse off than they would be if they were able to consume without inducing a hot state.

#### 4.2.3 *Infinite Horizon, Sophisticated Hot Selves*

So far, we have assumed that hot selves do not optimize, but instead consume if the immediate utility of doing so given their preferences is positive: that they are naïve. The cold self is therefore able to best respond to the hot self, while the hot self's behaviour is independent of the cold self's choice. If both selves are sophisticated, however, then they each best respond to the behaviour of the other self.

In order to examine behaviour when both selves are sophisticated, let us consider a consumer facing an infinite horizon, so that the two selves interact repeatedly. As before, in each period the consumer makes a binary decision on whether to consume or abstain. When the cold self indulges, they induce the hot self. Now, however, the hot self retains control until they choose to abstain, which induces the cold self. The cold self resists temptation if a shock is below a critical threshold of  $u$ ,  $\bar{u}$ ; the hot self resists temptation if the shock is below their critical threshold of  $u$ ,  $\underline{u}$ . Each self is indifferent between consuming and abstaining at their threshold.

If both selves are naïve, then optimal behaviour is identical to the finite horizon: consumers adopt naïve strategies of  $\underline{u} = -h$  and  $\bar{u} = 0$ . If only one self is sophisticated, that self can best respond to the naïve behaviour of the other self. If both are sophisticated, however, they each account for the behaviour of the other self when they optimize. As

in the two-period game, the cold self must account for the fact that if they choose to indulge, they induce a hot state in which they may not behave optimally based on their cold preferences. In addition, however, the hot self accounts for the fact that if they choose to abstain, they induce a cold self that may not behave optimally based on their hot preferences.

### *Cold State Utility*

I denote the expected cold self's utility from being in state  $S$  in period  $t$  as  $V_t^S$  and the hot self's utility from being in state  $S$  as  $\hat{V}_t^S$ . Since the cumulative distribution function of  $u$  is  $P(u)$ , the probability a taste shock is below  $X$  is  $P(X)$ . The value of being in a cold state to a cold self at the beginning of period  $t$  is therefore a combination of the probability a shock will be sufficiently low that they abstain, in which case they get no utility this period and remain in the cold state for the subsequent period, and the probability a shock will be sufficiently large that it convinces them to consume, in which case they get utility of  $u$  immediately and are in a hot state for the subsequent period.

$$V_t^C = P(\bar{u}) (0 + \delta V_{t+1}^C) + \int_{\bar{u}}^{\infty} [u + \delta V_{t+1}^H] p(u) du \quad (62)$$

Note that if a consumer always chose to abstain,  $V_t^C = 0$ , because there is no direct utility from abstaining. Similar to their value of being in a cold state, the cold self values being in a hot state at the beginning of period  $t$  as a combination of the utility they get from abstaining and the utility they get from indulging, weighted by the probability of each.

$$V_t^H = P(\underline{u}) (0 + \delta V_{t+1}^C) + \int_{\underline{u}}^{\infty} [u + \delta V_{t+1}^H] p(u) du \quad (63)$$

Since the consumer faces an infinite horizon, the only difference in each period is the realized idiosyncratic shock, and so the value of being in a given state at the beginning of each period (before the shock is realized) is equal:  $V_t^S = V_{t+1}^S = V^S$ .

The cold self maximizes their utility with respect to  $\bar{u}$ , selecting a level of indulgence given the value they assign to being in a hot state. At their optimum, they therefore set  $\frac{dV_t^C}{d\bar{u}} = 0$ . Differentiating equation 62, it follows that at their optimal threshold,

$$\frac{dV^C}{d\bar{u}} = p(\bar{u}) \delta V^C - p(\bar{u}) (\bar{u} + \delta V^H) + \delta (1 - P(\bar{u})) \frac{\partial V^H}{\partial \bar{u}} = 0 \quad (64)$$

To find how the value of being in a hot state to the cold self changes as the cold self changes their consumption threshold, I differentiate equation 63:

$$\frac{\partial V^H}{\partial \bar{u}} = 0 \quad (65)$$

Substituting equation 65 into equation 64, we can thus derive a first order condition that determines the cold self's optimal threshold

$$\bar{u}^* = \delta (V^C - V^H) \quad (66)$$

Similar to the two period case, the cold self's consumption threshold is a function of the value they get from being in the cold state and the value they get from being in the hot state. Each of these values, in turn, is a function of the hot self's behaviour, particularly their threshold for indulging in consumption,  $\underline{u}$ . Let us define the conditional mean of  $u$  given that it is greater than  $X$  as

$$G(X) = \frac{\int_X^\infty u p(u) du}{\int_X^\infty p(u) du} \quad (67)$$

Substituting this into equations 62 and 63, we can rewrite  $V^C - V^H$  as a function of  $\bar{u}$  and  $\underline{u}$ . Finally, we can substitute this into equation 66 to get an implicit equation defining the cold self's optimal threshold as a function of the behaviour of the hot state,  $\underline{u}$ ; their preferences; and the distribution of  $u$ , which determines both the cumulative distribution function and the conditional expected value,  $P(X)$  and  $G(X)$ .

$$\bar{u} = \delta (V^C - V^H) = \delta \frac{G(\bar{u}) (1 - P(\bar{u})) - G(\underline{u}) (1 - P(\underline{u}))}{1 - \delta (P(\bar{u}) - P(\underline{u}))} \quad (68)$$

When  $u > \bar{u}$ , the cold self indulges: when  $u \leq \bar{u}$ , the cold state self resists.

Each self's best response is a function of the best response of the other player. The Nash equilibrium will depend on the optimal strategy of the other self, and on the distribution of  $u$ . I therefore now consider the optimal behaviour of the hot self.

### Hot State Utility

Inversely to the cold state, the hot state persists when the hot self consumes, and ends when the hot self abstains. As a result, it is abstaining that induces a change in preferences and a time-inconsistency problem. The hot self's expected value of being in a hot state at the beginning of period  $t$  is therefore based on the probability the taste shock will convince them to abstain, in which case they get nothing in the current period and begin the next in a cold state; and the probability the taste shock will induce them to consume, in which case they get their utility from consumption and remain in a hot state in the subsequent period. As before, since the consumer faces an infinite horizon and the taste shocks are independent and identically distributed,  $\hat{V}_t^S = \hat{V}_{t+1}^S = \hat{V}^S$ .

$$\hat{V}^H = P(\underline{u}) (0 + \delta \hat{V}^C) + \int_{\underline{u}}^{\infty} [u + h + \delta \hat{V}^H] p(u) du \quad (69)$$

The hot self anticipates value from being in the cold state of:

$$\hat{V}^C = P(\bar{u}) (0 + \delta \hat{V}^C) + \int_{\bar{u}}^{\infty} [u + h + \delta \hat{V}^H] p(u) du \quad (70)$$

The hot self optimizes by choosing a value of  $\underline{u}$  that maximizes their value function. At the optimum,  $\frac{d\hat{V}^H}{d\underline{u}} = 0$ . By an analogous process to that of the cold self, we can differentiate equations 69 and 70, and derive a first order condition that determines the hot self's optimal threshold for consumption:

$$\underline{u} = \delta (\hat{V}^C - \hat{V}^H) - h \quad (71)$$

Finally, rewriting equations 69 and 70 and substituting into equation 71, we get an equation for the hot self's consumption threshold as a function of the cold self's behaviour and the distribution of  $u$

$$\underline{u} = \delta (\hat{V}^C - \hat{V}^H) - h = \frac{\delta G(\bar{u})(1 - P(\bar{u})) - \delta G(\underline{u})(1 - P(\underline{u})) - h}{1 - \delta (P(\bar{u}) - P(\underline{u}))} \quad (72)$$

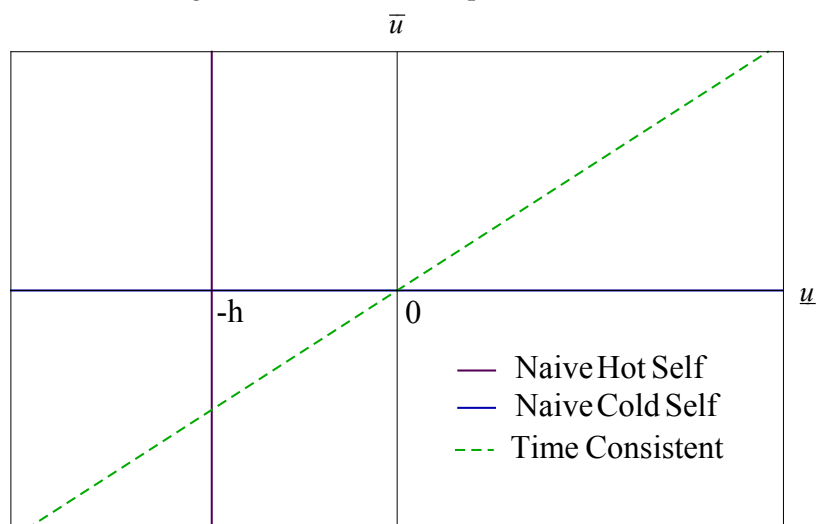
When  $u > \underline{u}$ , the hot self indulges: when  $u \leq \underline{u}$ , the hot self abstains. The hot self's optimal threshold parallels that of the cold self, but accounts for the extra utility they get from consumption,  $h$ .

### Equilibrium Behaviour

We now have implicit equations  $\bar{u}(\underline{u})$  and  $\underline{u}(\bar{u})$ , which define each self's best response as a function of the other self's behaviour. How each self behaves depends on their own preferences, how the other self will behave, and the expected distribution of shocks. In the Nash equilibrium the two selves play equilibrium values of  $\{\bar{u}^*, \underline{u}^*\}$ , where  $\bar{u}^* = \bar{u}(\underline{u}^*)$  and  $\underline{u}^* = \underline{u}(\bar{u}^*)$ : they best respond to the behaviour of the other self, conditional on their expectations about the shocks they expect in future periods.

A naïve self does not anticipate that their choices will affect their future behaviour. If  $\delta = 0$ , the consumer places no value on the future, and so though a self may anticipate future consequences from their behaviour, they are indifferent to them. In either case, they then act as if they faced a repeated stage game that is identical in each period, and consume whenever the marginal utility of doing so in the stage game is greater than zero:  $\bar{u} = 0$  and  $\underline{u} = -h$ . Figure 11 shows the best response functions for naïve selves.

Figure 11: Naïve Best Response Functions



Naïve selves behave independently of the decisions of the other self, because they do not anticipate that their future preferences will be different.

If both selves are naïve, each plays their stage-game utility-maximizing strategy, independent of the behaviour of the other self since they do not foresee a future change in preferences. The equilibrium is time-inconsistent, since because the two selves have different preferences, their naïve strategies are different.

The cold self's optimum point is at the origin, and so their best response function is along the x-axis. The hot self's optimum point is at  $\{-h, -h\}$ , and so they consume along

the line  $\underline{u} = -h$ . Relative to their optima, the hot self would be better off if the cold self consumed more often, and the cold self would be better off if the hot self consumed less often.

Sophisticated selves account for the fact that their behaviour may induce another self which has different preferences from their own. Inducing another self – for the hot self, choosing to abstain, and for the cold self, choosing to indulge – may therefore impose a cost on the current self in terms of time-inconsistency.

Both selves would prefer to impose behaviour on their other self and achieve a time-consistent consumption path. In that case, since each self's naïve strategy maximizes their utility in the stage game, with perfect self-control<sup>6</sup> the self in control adopts their naïve strategy in all periods. They therefore consume where their naïve best response function crosses the 45° line.<sup>7</sup>

Absent perfect self-control, time consistency might occur at a corner solution. If  $u > 0$  for all  $u$ , for example, both selves always consume, and so there is a time consistent equilibrium of complete indulgence. Equally, if  $u < -h$  in all periods, both selves always abstain.

To consider the behaviour of sophisticates in more detail, let us assume that  $u \sim U[a, b]$ . It follows that the best response curves of each self are

$$\bar{u} = \delta \left[ \frac{\frac{1}{2} (\underline{u}^2 - \bar{u}^2)}{b - a - \delta (\bar{u} - \underline{u})} \right] ; \quad \underline{u} = \frac{\delta \left( \frac{1}{2} \right) (\underline{u}^2 - \bar{u}^2) - (b - a) h}{b - a - \delta (\bar{u} - \underline{u})} \quad (73)$$

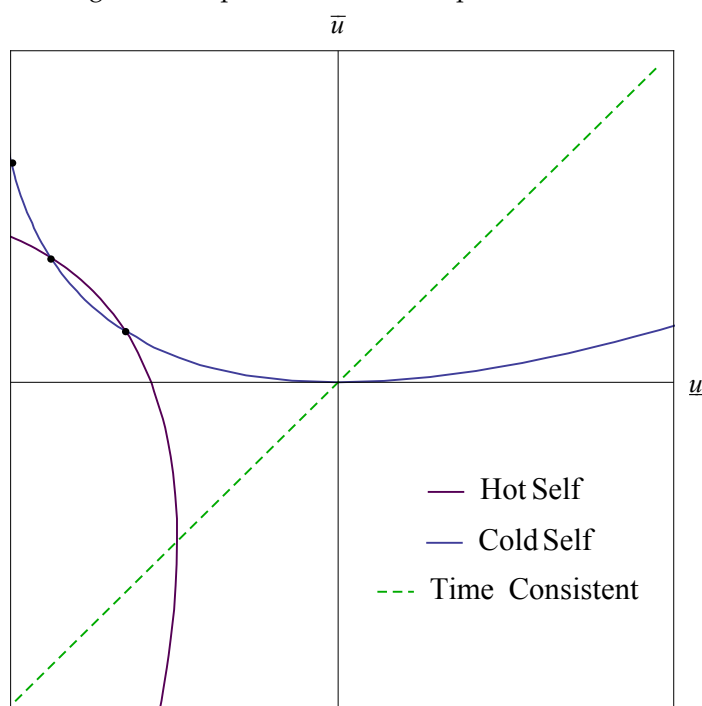
If, for example, we assume  $a = -1$ ,  $b = 1$ , then we can graph the best responses of each self, as in Figure 12.

Because the two selves have time-inconsistent preferences, both selves move away from their naïve behaviour, and behaviour is time-inconsistent: the best response curves do not intersect on the 45° line. Recall that if there was only one self, that self would welfare maximize in each stage game by adopting the naïf strategy for their type. Relative to that benchmark, here the cold self would prefer the hot self consume less often, and

<sup>6</sup> Literally 'self'-control, in this case - the ability to control the behaviour of the other self

<sup>7</sup> Consumers might act as if they had cold or hot preferences with perfect self control, depending on which self was able to impose its choices over the other – typically, cold preferences are treated as the 'real' preferences of the consumer (as, for example, in Loewenstein (2005), who distinguishes between hot affected states and cold unaffected ones).

Figure 12: Sophisticated Best Response Functions



Each self best responds to the behaviour of the other self. As the behaviour of the other self departs from the self's optimum, the cost of inducing another state increases, and the self abstains or indulges more often.

the hot self would prefer the cold self consume more often. Absent coercive ability over the other self (whether self-control or an external mechanism), each self best-responds to the behaviour of the other.

Conditional on the hot self behaving non-optimally by the standards of the cold self, the cold self has an incentive to consume less often to avoid inducing the misbehaving hot self. Similarly, conditional on the cold self behaving non-optimally by the standards of the hot self, the hot self has an incentive to consume more often, to avoid inducing the misbehaving cold self.

Note that multiple equilibria may exist. In Figure 12, there are three: two interior equilibria (one of which is unstable), and one corner solution where the hot state always consumes, and the cold state often – though not always – abstains. The worse the behaviour each self expects of the other (by their standards), the larger their incentive to also adopt extreme behaviour relative to their naïf choice, in order to reduce the chance of entering the other state. If they expect only moderate deviation by their other self, however, then their best response is also a relatively moderate deviation, and so the equilibrium is closer to their naïf preference.

Only one of the two interior equilibria – the more moderate one – is stable. The corner solution need not involve extreme strategies by both selves, however. In Figure 12, the cold self's best response to a hot state that always indulges ( $\underline{u} = -1$ ) is to heavily restrict consumption, but not to abstain completely ( $\bar{u} < 1$ ). Since abstaining provides no utility, even when the hot self always consumes, the cold self still prefers to consume for a sufficiently high taste shock, despite the permanent loss of control that entails. In response to this behaviour by the cold state, the hot state's best response is to indulge as much as possible, and so it is an equilibrium. For other distributions of shocks, however, the corner solution might involve both selves adopting a threshold at the edge of the distribution.

Since welfare falls as behaviour becomes more extreme, both selves would prefer to be in a more moderate equilibrium, since it allows them both to consume more closely to their optimum. Depending on their expectations, however, this may or may not be the equilibrium result. In all cases, however, the cold self consumes less than they would as a naïf, while the hot self consumes more.

The more the two selves disagree, the larger the absolute value of sum of their optimal strategies: from equations 73,  $\underline{u} + \bar{u} = -h$ . The difference between the optimal strategies for each self, in contrast, need not be unique: rearranging equations 73,  $\delta (\underline{u} - \bar{u})^2 + (b - a) (\underline{u} - \bar{u}) + (b - a) h = 0$ . This provides a measure of the total sacrifices each self makes in order to best respond to the other self's behaviour – informally, we could consider this a measure of willpower. It is a function of their discount rates, the difference between the two utility functions, and the distribution of taste shocks.

Implicitly differentiating the cold self's best response function, we can find how their best response changes in response to a change in the behaviour of the other self:

$$\frac{d\bar{u}}{d\underline{u}} = \frac{\delta (\underline{u} - \bar{u})}{b - a - \delta (\bar{u} - \underline{u})} \quad (74)$$

and, for the hot self:

$$\frac{d\underline{u}}{d\bar{u}} = \frac{\delta (\underline{u} - \bar{u})}{b - a - \delta (\bar{u} - \underline{u})} \quad (75)$$

The cold self's slope is greater than that of the hot self in the relevant region if  $\bar{u} - \underline{u} < \frac{b-a}{2\delta}$ . For sufficiently low  $\delta$ , there is a unique interior equilibrium: if selves do not place

much weight on the future, the incentive to alter their behaviour to prevent the other self from taking control is muted, and so within the range of possible values of  $u$ , only one equilibrium exists.

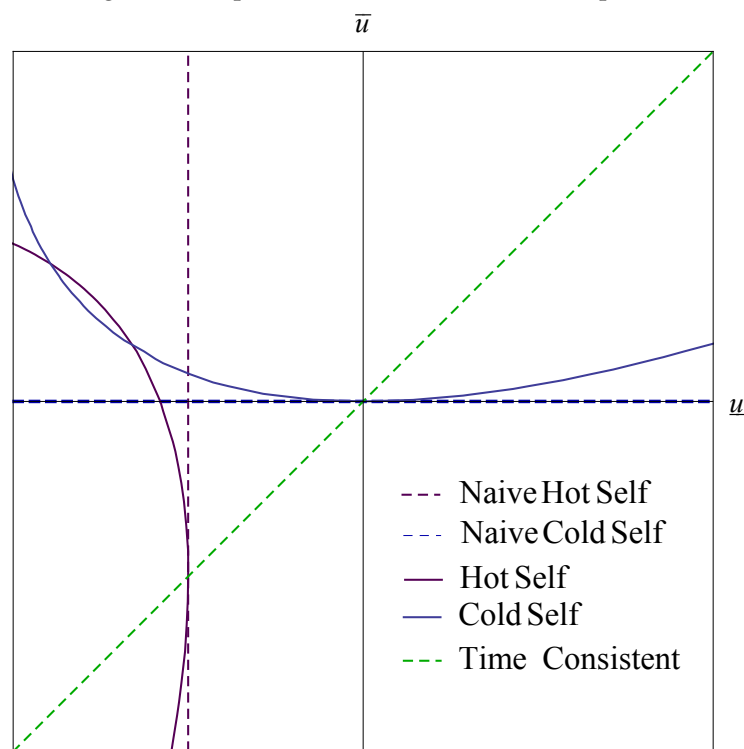
As the parameters of the problem change, the best response curves of each self shift. As  $\delta$  falls, for example, each self places less value on the future and so cares less about the costs of being in the other state. The best response curves of each self flatten, and sophisticated behaviour approaches naïve behaviour. As  $h$  increases, in contrast, the difference in preferences between the hot and cold self increases, increasing the cost of inducing the other state. For  $h = 0$ , the naïve and sophisticated best response functions for both selves cross at the origin, and so the naïve and sophisticated equilibrium thresholds are time consistent (and identical) because the two selves have identical preferences. As  $h$  increases, the two selves increasingly disagree, and the other self's behaviour is so undesirable that both selves adopt extreme behaviour to prevent it.

Also important is the distribution of  $u$ . The cost the cold self anticipates from inducing a hot self is determined by the likelihood the taste shock is in the range  $-h < u_i \leq 0$ . If  $u \leq -h$ , then both the hot and cold selves agree they would prefer to resist temptation. If  $u > 0$ , both the hot and cold selves prefer to indulge in temptation. There is therefore no disagreement. For  $-h < u \leq 0$  however, the hot self gains positive utility from indulging while the cold self does not. The probability that the taste shock falls in that range, and the utility when it does so, therefore determines the scale of the time inconsistency.

The range of the distribution affects the expected return for each self from consuming as well as the likelihood that the two selves will disagree about the optimal behaviour. Straightforwardly, for some distributions the selves agree on the optimum course of action in every period: since when the taste shock is below  $-h$  or above  $0$  both selves agree on the optimal choice, if the distribution of taste shocks is entirely in those regions, the two selves always agree. For a fixed disagreement  $h$  and a uniform distribution of shocks, as  $a \rightarrow -\infty$  and  $b \rightarrow \infty$  the equilibrium goes to  $\{-h, 0\}$ ; the probability of a shock occurring in the range that provokes disagreement goes to zero, and so both selves play their naïve strategy. As the range increases and the probability of a disagreement falls, the incentive for each self to alter their behaviour in order to prevent a change of state falls, and so they consume closer to the naïve optimum.

The relationship between naïve and sophisticated behaviour is shown in Figure 13. Sophisticated selves anticipate future utility or disutility from their choices, encouraging them to consume more or less as appropriate. Each self would prefer, however, to impose behaviour on the other self so they can achieve their naïve optimum.

Figure 13: Sophisticated and Naïve Best Responses



Each self deviates from their naïf optimum based on the behaviour of the other self.

As  $\underline{u} \rightarrow 1$ , the hot self always abstains, even more than the cold self would find optimal. The cost of inducing a hot state for the cold self therefore increases, and the cold self's best response deviates increasingly far from their naïve strategy. If the hot self always abstains, however, then the cold self regains control immediately after they abstain once, and so the cost of inducing a hot state is limited: consuming means that the cold self will get no utility in the subsequent period, and then be in a cold state again. As  $\underline{u} \rightarrow -1$ , the cold self's best response is again to deviate further and further from their naïve strategy. Now, however, if the cold self induces a hot state it may persist for many periods (for  $\underline{u} = -1$ , it is permanent), and so at the limit the cold self must weigh the utility they will get from consuming in a period against the costs of enduring a permanent hot state in which they consume in all subsequent periods.

The hot self faces similar incentives. Their naïve optimum is to set  $\underline{u} = -h$ . As  $\bar{u} \rightarrow -1$  or  $\bar{u} \rightarrow 1$ , the cost of inducing a cold state increases, and so the incentive to consume

and perpetuate the hot state increases, encouraging deviation from their naïve strategy. This is particularly costly as  $\bar{u} \rightarrow 1$ , because this implies that if the hot self abstains, the cold state is anticipated to endure for a large number of periods. As  $\bar{u} \rightarrow -1$ , in contrast, inducing a cold state guarantees consumption in the subsequent period, after which the hot self regains control.

If the hot self is naïve, then a sophisticated cold self faces a vertical best-response function, as in Figure 13. Similarly, if it were possible for the hot self to be sophisticated while the cold self was naïve (as described by Jekyll and Hyde, perhaps), then the hot self would face a horizontal best response function. In either case, the other self is then able to best respond to the naïve behaviour, and the sophisticated self is able to achieve a consumption path closer to their own time-consistent optimum. It is not clear whether one self could be sophisticated while another was naïve, however, since both make up the same consumer. Whether the consumer is better off to have both selves naïve or both sophisticated will depend on relative welfare for at  $\{-h, 0\}$  and  $\{\underline{u}^*, \bar{u}^*\}$ .

Similarly to the two period case, we observe that consumers adopt strategies modifying their consumption from what they would choose as a naïf or if they had perfect self-control. The hot self increases consumption in order to reduce the frequency of inducing a cold state: the cold self resists temptation more often in order to avoid inducing a hot state. With a sophisticated hot self, however, multiple equilibria are possible, with relatively moderate or extreme deviations from optimal behaviour. The equilibrium also depends on the preferences of the consumer – particularly the size of time inconsistency between the two selves as measured by  $h$ , the discount rate  $\delta$ , and the scale of the temptation.

### *Discussion*

$\bar{u}$  is a personal rule for a sophisticated cold self. They are aware that the marginal utility of consuming a square of chocolate is higher than the marginal cost in period  $t$ , but because they anticipate that indulging will make them more likely to overindulge in the subsequent period, they resist temptation. The higher the expected costs from being in a hot state, the higher  $\bar{u}$  will be, until in some cases the cold self might prefer never to consume.

Personal rules thus reflect the relative sophistication of consumers: anticipating the change in preferences from being in a hot state, they commit to a personal rule that restricts consumption, because doing so reduces their time inconsistency. A cold self who can stay in that state reduces their time inconsistency, and so has an incentive to follow the rule. Similarly, a hot self that can stay 'hot' reduces their time-inconsistency.

Ainslie (1992) introduced personal rules as a potential strategic response by agents facing temptation, suggesting that people adopt a strategy of grouping decisions today into sets with decisions tomorrow, so that each choice became a precedent predicting future choices in the set. Even if a consumer prefers to indulge in the immediate moment – is happy to eat a single chocolate in exchange for the 100 calories it contains – he does not prefer to consume the entire bundle – the rest of the chocolate box. The cost to making any single bad decision increases, and people are better able to resist the initial temptation. Consumers thus create their own activation cost for a behaviour by making a choice to indulge in it also a choice to indulge in later, undesired, behaviours.

In order for a personal rule to be effective, however, the consumer must be able to commit to it. Otherwise, each time a rule is violated, it is optimal for a consumer to create a new rule that groups all future violations of the rule together, and exempts the preceding indulgence. In other words, if a consumer creates a personal rule not to eat chocolate, and then does so, there is an incentive to find a reason doing so doesn't break the rule, or begin a new rule that states that *from that point forward*, the consumer will eat no chocolate. Indeed, it is easy to picture consumers doing exactly this. Anticipating this time-inconsistency, however, a consumer gains no value from establishing a personal rule: they are aware that they can always ignore the rule after it is violated, and so a personal rule has no extra benefit beyond their original desire to resist temptation. Bratman (1999) worries that abstaining or indulging today may give no practical reason to abstain or indulge tomorrow, rendering personal rules impractical because there is no element of commitment.

The model introduced above suggests an alternate perspective. If indulging once induces a hot state, then a personal rule is a sensible strategy on the part of the consumer. Knowing that they are more likely to succumb in subsequent periods if they indulge in one, a temptation must be correspondingly stronger to be worth succumbing to. Personal rules are thus a result of sophistication on the part of the cold self, because they

anticipate that indulging once will begin a hot state. In an extreme, a rule might even forbid consuming a type of good entirely.

The effect of personal rules, however, is determined by which self is adopting the rule. A consumer deciding while in a cold state how much chocolate to eat has an incentive to adopt a rule that reduces their likelihood of consumption. If the temptation becomes overwhelming, however, or if something else intervenes to make them succumb, then their hot self takes over and may also employ a personal rule: realizing that if they return to a cold state, they might not get a chance to eat more chocolate, the hot self consumes more than even they would want.

#### 4.2.4 Long-Run Effects

I now briefly extend the model to consider consumers who choose a continuous amount of consumption, rather than whether to indulge or abstain, in a setting where only a subset of goods induce a hot state. For simplicity, I use a deterministic version of the above model. This could also be seen as a longer-term version of the model, in which the short periods described above aggregate into longer, macro-periods.

Let us consider a consumer facing a decision on how to allocate an endowment  $W$  across an infinite horizon. The consumer can choose between two types of goods:  $c$  (as in chocolate), which induces a hot state with respect to consumption of that good, and  $b$  (as in bread), which does not. If the consumer chooses to consume a strictly positive amount of  $c$ , then their marginal utility of consuming  $c$  increases. Utility is separable in  $c$  and  $b$ . In the cold state, utility is thus

$$u_t^C = f(u(c_t), u(b_t)) \quad (76)$$

If the consumer indulges in period  $t$ , they enter a hot state. In a hot state, the consumer places a higher value on immediate consumption of good  $c$ , so utility is

$$u_t^H = f(\theta u(c_t), u(b_t)); \quad (77)$$

Note that the effect of entering a hot state is now multiplicative, not additive, in order to preserve the increase in marginal utility. Hot states persist only for the immediate

period, and hot selves optimize as if they received a higher marginal utility from consumption in that period. Total utility for the cold self is therefore

$$U^C = \sum_{t=0}^{\infty} \delta^t u_t^C \quad (78)$$

while total utility for the hot self is

$$U^H = u_t^H + \sum_{t=1}^{\infty} \delta^t u_t^C \quad (79)$$

The price of good  $c$  is  $p_c$ , and the price of good  $b$  is normalized to 1: the consumer is a price taker in both goods, and faces constant prices across time. The optimization decision is therefore subject to the standard budget constraint,  $\sum_{t=1}^{\infty} p_c c + b \leq W$ .

A naïve cold self does not anticipate this change in preferences, and so plans to consume their optimal amount each period. As in the standard model, they set the marginal utility of consuming good  $c$  equal to the marginal utility of consuming good  $b$ , and the marginal utility of increasing consumption equal across all periods. This consumption plan, however, is not enforceable. Conditional on consuming  $c > 0$ , a hot state is induced, in which the marginal utility of consuming  $c$  is increased. Each period, the naïf ends up consuming more  $c$  than they planned, and is worse off than they expected: the glutton who is constantly surprised by their overconsumption. The utility the naïve cold self receives from consumption is therefore lower than he would have gotten had he been able to commit to his optimal consumption pattern, because the hot state preferences affect his behaviour but not his cold state utility.

A sophisticated cold self, in contrast, faces a kinked utility curve. Marginal utility at  $c = 0$  is different from utility at  $c > 0$ , and so a corner solution is particularly likely. Since the cold self is unable to commit to their desired consumption path, in other words, they can choose either to consume nothing or to begin consuming and so consume as if they had hot preferences.

The global optimum for the cold self might be to consume a moderate amount of chocolate. The cold state self faces a limited choice set, however, because conditional on consuming a strictly positive amount, they cannot control how much they consume. They are therefore forced to choose between a local optimum of consuming nothing, or the hot self's optimum consumption point. Defining the hot-self optimal consumption

bundle as  $(c^H, b^H)$  and the cold self's optimal consumption of  $b$  if they abstain entirely from  $c$  as  $\hat{b}$ , if

$$u^C(c^H, b^H) \leq u^C(0, \hat{b}) \quad (80)$$

then it is optimal for the cold self to consume nothing of good  $c$ , and consume exclusively good  $b$ : in other words, to implement a personal rule forbidding consumption of a good that induces a hot state. When condition 80 holds, in other words, the consumer prefers to eat no chocolate at all than to eat the entire bar, even though eating a single square may be better than both of those options.

This is likely to hold when the hot state is particularly distortionary to preferences (when  $\theta$  is large), or when good  $c$  has rapidly diminishing or negative marginal returns over parts of its range. In that case, a consumer may be willing to give up a lot to resist temptation and avoid overconsuming relative to their optimum. This may explain why personal rules are often discussed in the context of goods that are best consumed in moderation, such as chocolate or alcohol: the health impacts of both when consumed in excess mean that their returns are sharply diminishing. Otherwise, the cold self may accept the best they can do is to allow the hot self to indulge: they would ideally like only a single square of chocolate, but accepts that their second best option is to allow the hot self to eat the entire bar.

The potential for hot states thus predicts we might observe some consumers abstaining entirely from consuming some types of goods, particularly if such goods induce very strong hot states, or if overconsumption has particularly deleterious effects, even if moderate consumption has no harmful effects at all.

In the long run, permanent shifts in marginal utility might result from repeated entry into the hot state. This is what is captured in studies of addiction or long-term habits: consumers find themselves permanently affected by their behaviour, so much so that even after a significant amount of time, they are more vulnerable to a temptation than consumers who never indulged in the first place (Nestler, 2001).

This long-run effect would increase the incentive to use a personal rule. The classic model of long-term habits, Becker and Murphy (1988), suggests consumers are characterized by the creation of 'consumption capital': as they consume more of the addictive good  $c$ , their stock of consumption capital  $S_t$  increases, which increases the marginal

utility they receive from consuming more of it. Here, since entering a hot state increases marginal utility, hot selves accumulate more consumption capital than the cold self would deem optimal, and may begin a path of addiction. Naïf cold selves then risk accidental addiction. Sophisticates, aware of the risk, would then employ correspondingly stronger personal rules, further restricting consumption.

Personal rules denying consumption of an entire class of good are thus a rational strategy for sophisticates to adopt in some contexts, particularly when the hot self has very different preferences, or when  $c$  has rapidly diminishing or negative marginal returns over parts of its range, such as health implications from overconsumption. If there is the potential for long-term habits to form, then naïve consumers could also become addicted by accident, with significant welfare implications. In light of Loewenstein's work on hot-cold empathy gaps (Loewenstein, 2005), which suggests that cold selves often struggle to predict the preferences of their hot selves, naïve accidental addiction could be a serious concern. In addition, the potential for addiction increases the incentive for sophisticated consumers to use stronger personal rules, particularly when the potential temptation effect is large. We might therefore expect consumers who are particularly concerned about the health implications of chocolate to demonstrate the largest temptation effects in the lab, for example.

### 4.3 EXPERIMENT

Having considered a model of hot and cold states, I test for their existence in the lab. In particular, the model suggests consumers behave differently if they are in different states. In light of the theoretical analysis, I propose two testable implications of the model. First, and the most basic prediction, is that early indulgence in a consumption good increases total consumption, and makes it less likely a consumer will abstain completely. Second, the theory also suggests that the effect should be particularly pronounced for subjects who are tempted by the consumption good or who are motivated to use personal rules.

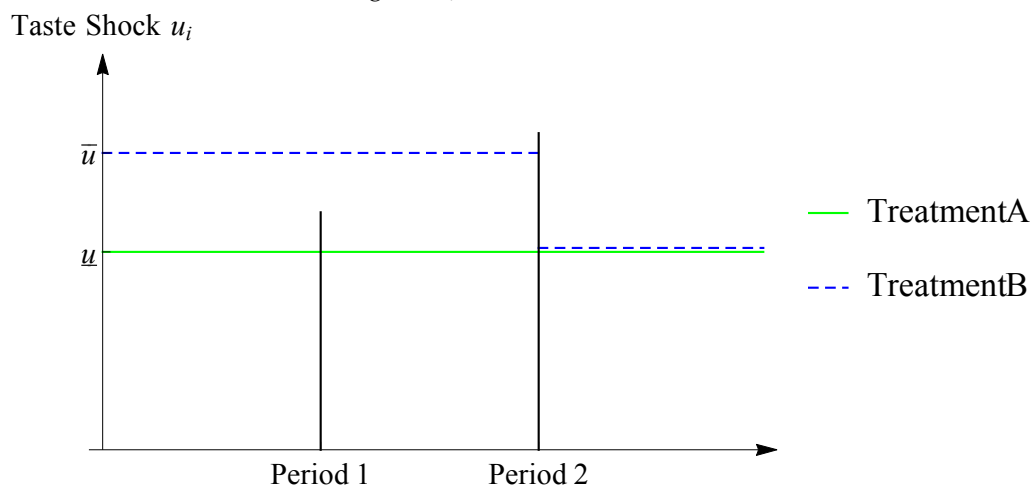
The initial state of the consumer affects expected consumption. To see this, consider behaviour if in period zero nature determines whether the consumer begins in a hot or cold state. Naïve cold selves consume less than hot cold selves: if the two selves are

sophisticated, this difference is exacerbated. The increased likelihood a consumer who begins in a cold state will abstain reduces their consumption directly and reduces the probability of the consumer entering a hot state, further reducing their expected consumption. A consumer beginning in a hot state therefore consumes more in expectation and is less likely to abstain entirely from consumption.

**Hypothesis 1: Consumers who indulge in a good once (and so are in a ‘hot’ state) will consume more in total over subsequent periods, and be less likely to abstain entirely from consumption.**

In the experiment, participants in Treatment A are asked to eat a piece of chocolate immediately, while participants in Treatment B – the control – are not and so begin in a cold state. Figure 14 illustrates this setup in light of the model, with a particular draw of taste shocks  $u$  over two periods. Participants in Treatment A indulge in both periods, while those in Treatment B manage to resist in the first period, and only indulge in the second, after which their threshold drops to the same level as participants in A. Consumers in Treatment A would thus be observed to consume 50% more chocolate in total, assuming each indulgence is of equivalent size.

Figure 14: Treatments A and B



In the experiment below, some participants are asked to indulge immediately in chocolate (Treatment A), while others are not (Treatment B). The subjects in Treatment A, who begin in a hot state, indulge in both periods, while those in Treatment B resist the shock in period 1 and only consume and enter a hot state in period 2.

An important question is whether having an initial behaviour exogenously imposed, rather than the result of free choice, could affect the change in the temptation threshold. The literature is mixed as to whether this affects behaviour. Some theories, such as self-signalling (Bénabou and Tirole, 2004) or the potential for willpower to be used up as consumers resist temptation (Baumeister et al., 1998), rely on an initially free choice by consumers, while others, such as a physiological or information effect, do not.

Second, the motivation for consumption to induce a hot state was psychological evidence that suggests complementarity in consumption of tempting goods. Since consumers vary in how tempting they find a given product, we might reasonably expect consumers who find something tempting but are attempting to resist that appeal to demonstrate a larger temptation effect. Consumers on a diet may find chocolate particularly tempting, since they are attempting to restrict their consumption of it. In addition, the theory predicts that consumers with particularly extreme hot states ( $h$  high) will be particularly reluctant to indulge while in a cold state. These consumers are therefore more likely to diet while in the cold state, restricting their consumption of tempting foods so as to avoid inducing the hot state. This selection effect implies that consumers who have strong personal rules while in the cold state – in other words, who are on a diet – should demonstrate a particularly pronounced change in preferences when they enter a hot state:

**Hypothesis 2: Consumers who find a good more tempting – such as those on diets, in the case of food, or attempting to quit, for cigarettes – are particularly likely to demonstrate a temptation effect, or to demonstrate a stronger temptation effect.**

To test these hypotheses, I conduct a lab experiment.

#### 4.3.1 *Design*

The experiment was conducted in order to test for the existence of a temptation effect: would subjects who had already had a piece of chocolate, and thus potentially entered a

hot state, consume more chocolate in total? 149 students from the University of Oxford participated in eight lab sessions, all scheduled at 3pm in order to minimize the variation in hunger between sessions.<sup>8</sup> One subject was unable to take part due to a food allergy, and so the total sample was 148. Based on the theory, subjects on diets were of particular interest, but to avoid increasing the salience of diet and temptation to potential subjects, this was not mentioned in recruitment. The experiment took roughly one hour. Genders were balanced, with 74 women and 74 men. Subjects were paid just over £10.50 on average, including a £4 show-up fee. Students were recruited through the Centre for Experimental Social Sciences at the University of Oxford.

As participants arrived at the lab, they were randomly allocated to either a treatment or control group. All subjects were taken individually into a small room and asked about any potential allergies. Group A – the treatment group – was given an unopened (but previously weighed) container of chocolates and asked to try one of the them. The participant was told they could take the chocolates with them to the workstation and have as many as they wanted. They were then led to the workstations. This was intended to induce a hot state of consumption in these participants.

Group B – the control group – was not asked to try a piece of the chocolate, and was instead just given the chocolate container and told to have as many as they wanted, then led to the workstations. For both groups, how long they would have the chocolates and whether the amount of chocolate they consumed would be measured was left unspecified. Subjects were all told to have as many as they wanted at the workstation, to minimize the difference in any implicit permissions from the experimenter.

The chocolates had a single exterior container (see Appendix A) to minimize the sound of consumption and potential spillover effects, as well as to make it difficult for participants to take chocolates with them. Chocolates also had a relatively small unit size, about the size of a 5 pence piece or American dime and weighing less than a third of a gram, to minimize the possibility that the initial piece consumed by participants in the treatment group would fully satiate participants.

Once all participants were seated, they were given a series of tasks to perform on their workstations in order to measure their discount rate and their executive function. If the treatment affected discount rates or executive function, this might lend some insight

---

<sup>8</sup> The experiment was programmed and conducted with the experiment software z-Tree (Fischbacher, 2007)

into the channel by which any effect on chocolate consumption might occur. Discount rates were elicited through a series of decisions between receiving four pounds today or a different amount in one week (ranging from £4 to £7.60 in forty pence increments – for comparison, the minimum hourly wage in the UK is £6.50), based on the choice-based approach in the literature (Hardisty et al., 2013). If subjects chose the reward in one week, they were asked to return to the lab to pick it up. Executive function was measured with a Stroop test, in which subjects were presented for several moments with the name of a colour written in a potentially different colour, and then asked to either select the colour named by the text or the colour of the text (Stroop, 1935; Duckworth and Kern, 2011).

After subjects had completed these tasks, the chocolate containers were collected from the workstations. Finally, subjects were asked questions to elicit their beliefs about the nature of willpower, including whether it was a limited resource exhausted by use or self-sustaining, as well as their demographics. Previous work has shown that beliefs about the nature of willpower can affect performance on some tasks (Job, Dweck and Walton, 2010). These questions were asked after the chocolate was collected to avoid revealing the relevance of willpower to the experiment. Subjects were then paid. Details of the experimental procedure can be seen in Appendix B.

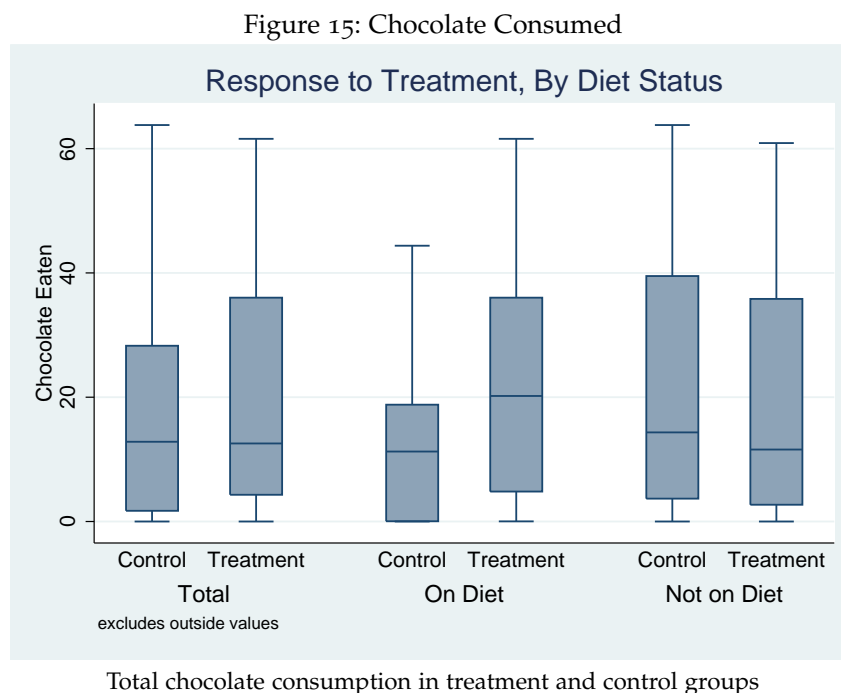
An important aspect of the experiment is that the initial behaviour was to some extent exogenously imposed on the consumer, rather than a result of free choice. This rules out several otherwise possible explanations for the effect, including the possibility for consumers to improve their ability to resist temptation through practice (Oaten and Cheng, 2006), or to learn about their ability to resist temptation (Bénabou and Tirole, 2004). The experiment is therefore a relatively strong test for consumption persistence, since it rules out explanations that require personal choice in the initial indulgence. The advantage of doing so is that subjects who choose to indulge a first time are clearly not a random sample.

#### 4.3.2 *Results*

I study the effect of the treatment on how much chocolate was consumed out of the approximately 65 grams of chocolate provided. Based on the theoretical analysis

(Section 4.2), I was particularly interested in how groups who might find chocolate tempting but be trying to resist the temptation would respond to the treatment. This criterion potentially applied to two groups: consumers on diets ( $n = 52$ ), who might be expected to find chocolate more tempting because they were denying it to themselves, and women ( $n = 74$ ), who anecdotally might have been expected to find chocolate more tempting. Because some participants ( $n = 12$ ) consumed all the chocolate given to them, I use a one-sided Tobit model (Table 1).

The treatment is not a significant predictor of chocolate eaten for participants as a whole ( $n = 148$ ,  $p = 0.633$ ).<sup>9</sup> The effect is also not significant for women ( $p = 0.508$ ). There is, however, evidence for an effect for consumers on a diet, as the theory would have ex-ante suggested. This may suggest the importance of attempting to resist temptation in inducing a temptation effect, as discussed in Section 4.2.1.



As can be seen in Table 1, absent controls the effect of the treatment on participants on a diet is only weakly significant ( $p = 0.082$ ). A non-parametric test, the Mann-Whitney  $U$  test, finds a significant relationship ( $p = 0.0286$ ). Controlling for the time they spend at the workstation, gender, and how much they like chocolate in general, the treatment is a significant predictor of how much chocolate consumers on a diet ate, based on the Tobit regression ( $p = 0.023$ ). Consumers on a diet ate a significant 8.92 grams more chocolate

<sup>9</sup> All figures reflect robust standard errors.

if they were in the treatment group – about half the weight of a Reese’s Peanut Butter Cup – in comparison to the control group. This difference is about 50% of the mean total chocolate consumed, 19.08 grams.

Because an interaction term between being on a diet and the treatment is not significant, this provides only weak evidence for the temptation effect: subgroup analysis has been found to be potentially misleading in some cases (Brookes et al., 2001). The theoretical motivation for focusing on consumers who are restricting their diet, and the existence of a significant non-parametric relationship, makes this result of particular interest, but it is more suggestive of an effect than conclusive. As can be seen in Figure 15, consumers on a diet do appear to eat more chocolate when in the treatment group in comparison to the those not on a diet and overall, but the variance in chocolate consumed is wide.

To some extent, therefore, the experiment provides support for the idea that consumers will consume more if asked to consume a chocolate beforehand (Hypotheses 1), but only for consumers on a diet (Hypothesis 3). Hypothesis 1 also suggests that consumers in a hot state are less likely to abstain from indulging entirely. Seven subjects ate no chocolate at all while at the workstations: six of them were in the control group, the cold state condition, and only one was in Treatment A, the hot condition. This is supportive of Hypothesis 2, but the small sample size means no conclusion can be drawn.

In the first specification, I consider the treatment effect on chocolate eaten absent all controls. In the second, I control for the order in which subjects sat at the workstation as a measure of how long they spent at the workstation (which affects the time the participant had to eat the chocolate), gender, and how much they said they liked chocolate in general.<sup>10</sup> The results are presented in Table 1. Figures reported include all chocolate consumed by the subjects, including both the treatment and at the workstation. This was done because the delay between the treatment and being at the workstation was small, suggesting consumers might treat the decision as a single choice, rather than two discrete opportunities to consume. Due to the small size of individual chocolates, however, examining the amount of chocolate consumed solely at the workstation does not

---

<sup>10</sup> As a check for bad controls, I also look at whether the treatment affects the subjects’ rating of how much they like chocolate. There is not a significant difference in the rating of the chocolate between the treatment and control groups.

change the significance of any of the results. Interestingly, whether a subject was on a diet, measures of their willpower, and gender were not significant in explaining total chocolate consumption.

Table 1: Effect of Treatment on Chocolate Eaten  
Dependent Variable: Amount of Chocolate Eaten while at Workstation

	Full Sample, No Controls	Full Sample	Restricted Diet, No Controls	Restricted Diet
Treatment	1.68 (3.51)	2.63 (3.97)	9.17* (5.16)	11.27** (4.80)
Time		-0.47* (0.25)		-0.32 (.34)
Gender		3.19 (3.33)		1.60 (4.85)
Chocolate Rating		5.29*** (1.07)		6.42*** (1.70)
Diet × Treatment		3.967 (5.05)		— —

Robust Standard Errors are in parentheses. \*Significant at 10%. \*\*Significant at 5%. \*\*\*Significant at 1%.  
Regression is a one-sided Tobit on the right tail.

A secondary question of interest was how the treatment affected consumer discount rates, since such a channel might explain the treatment effect. If the measure of discount rates is included in the above regression, however, the effect on consumers remains significant ( $p = 0.021$ ), suggesting a change in discount rates cannot explain a change in chocolate consumption. In fact, the measured discount rates are not a significant predictor of chocolate consumed ( $p = 0.27$ ). The treatment does affect discount rates directly, however. I run a Tobit regression examining how participants' discount rates are affected by the treatment, finding that the treatment does have a significant effect on how much participants discounted the future ( $p = 0.018$ ). A Tobit model is used because some subjects choose the immediate reward at all opportunities. Controlling for the same factors as above, the effect remains significant ( $p = 0.013$ ). Table 2 shows the results.

Being asked to taste one of the chocolates beforehand is associated with a subject choosing to take the reward in one week as opposed to immediately for an average of

Table 2: Effect of Treatment on Discount Rates  
 Dependent Variable: Observed Discount Rates

	No Controls	Controls
Treatment	-0.98** (0.41)	-1.03** (0.41)
Time		-0.02 (0.03)
Gender		-0.38 (.40)
Chocolate Rating		-0.06 (0.15)

Robust Standard Errors are in parentheses. \*Significant at 10%. \*\*Significant at 5%. \*\*\*Significant at 1%. Regression is a one-sided Tobit, since some students always chose the immediate reward.

£0.38 less. The treatment increases how much consumers value the future, as measured by long they are prepared to wait for a reward. A possible explanation for this effect may be the earlier consumption of glucose in the treatment: other work has suggested that glucose helps agents resist temptation in contexts like emotional regulation and thought suppression (Gailliot et al., 2007; Wang and Dvorak, 2010).

The treatment effect observed for consumers on a diet also does not operate through an effect on executive function: if performance on the Stroop test is included in the original Tobit regression, the treatment effect remains significant ( $p = 0.023$ ). Interestingly, an OLS regression examining the effect of the treatment on the performance of participants on the Stroop test shows no effect ( $p = 0.763$ ). This suggests the treatment has no effect on executive function. Since the Stroop test is commonly also used as a measure of willpower (Stroop, 1935), this suggests there is a difference in how the treatment interacts with a consumer's ability to resist chocolate, one form of willpower, and with their ability to exercise executive function, another form.

#### 4.3.3 Hot/Cold State Transition Mechanisms

The experimental results thus broadly conform to the model's prediction that consumers on a diet who indulged once would consume more in total. In some circumstances, diminishing marginal returns is potentially insufficient to describe short-term consumption patterns: instead, consumers can display complementarity in consumption. Though

the focus of the experiment was to test for the existence of such an effect, we can also go some way in understanding the mechanism underlying the hot/cold state transition.

Above, I discussed the possibility that hot state behaviour could be explained either through a change in preferences, or the existence of a fixed cost associated with succumbing to temptation that temporarily fell when the consumer indulged. A possible confound would also be if consumers were acquiring information from consumption, either about their type or how much they like the good.

Bénabou and Tirole (2004) suggest that consumers are unaware of their willpower type, and learn it by attempting to resist temptation. If they find they are a low willpower type, then they are more likely to fail to resist temptation the next time, inducing persistence in consumption of a tempting good. In the experiment, however, subjects are asked to eat the initial chocolate by the experimenter instead of succumbing on their own. The existence of an exogenous reason to have the chocolate should minimize any self-reputational effects for the participants, because their initial consumption of chocolate did not reflect a failure of willpower on their part. The use of a relatively common consumption good, chocolate, also limits the amount of information subjects are able to get from trying the chocolate. Indeed, when asked, only seven subjects said they were unsure as to whether they liked chocolate. A purely informational mechanism is thus unlikely to fully explain the observed temptation effect.

In contrast, a change in consumers' preferences or their knowledge of their preferences could explain the observed hot/cold state transition. If consumers forget how much they like chocolate, trying a good once could provide a reminder of their true, chocolate-loving preferences, corresponding to the hot state. Alternatively, consuming a good might induce a change in preferences directly. From a biological perspective, this could arise from a physiological effect due to the ingredients of chocolate, which would then correspond to the hot state. The experiment does not find an effect through a change in discount rates or executive function, but a change in the taste for chocolate remains possible. Biological work suggests that a physiological effect should be general to the human nervous system (Nassera et al., 2011), but consumers who experience a strong physiological effect from chocolate might be more likely to go on a diet.

The existence of a fixed cost to indulgence would also explain the observed state transition. When the fixed cost is unpaid, the consumer is in the cold state: once they

have paid the fixed cost, and the marginal cost of consumption has fallen, they are in the hot state. In a general setting, these costs could be either material or psychological: eating candy could require purchasing and opening a candy packet, and might make the hands of the consumer sticky. In that case, the first piece of chocolate could have a relatively large cost to indulgence, after which the cost of consumption would fall. In the experiment, however, these costs are minimized: the chocolate was provided to consumers and beginning consumption was relatively easy. Some material costs might still exist, but they appear to be relatively small.

A psychological fixed cost, in contrast, could explain the state transition described in the theory and is potentially present in the experiment. Indulging in a tempting good might temporarily sap willpower: having indulged once, it becomes harder to resist again, corresponding to the hot state. If consumers have psychologically committed to a personal rule as suggested by Ainslie (1992) (in this case, a diet), then a violation of that rule would also reduce the cost of further violations. As discussed in section 4.2.1, however, it is not clear how consumers would commit to such a strategy: they face a time-consistency problem in which it would be rational to create a new rule each time a rule was broken.

The experiment thus limits the number of possible explanations for the existence of a hot state, but does not isolate a single explanation: both the presence of fixed costs to indulgence and a temporary shift in preferences could explain the hot/cold state transition observed. The experiment does, however, highlight the existence of such an effect for some consumers.

#### 4.4 DISCUSSION AND CONCLUSIONS

We consistently observe persistence in macroeconomic patterns of consumption. Short-term patterns of consumption, however, also present an important factor in consumer behaviour. The assumption of diminishing marginal returns, though reasonable in many contexts, may not be sufficient to describe behaviour when indulgence can induce a hot self with different preferences.

This paper introduces a formal model of hot and cold selves in consumption. Whether consumers anticipate their transition between selves, as captured by their sophistication, is an important predictor of behaviour, and naïfs may typically overconsume while sophisticated cold selves adopt personal rules to restrict their consumption because they anticipate the future costs of consuming due to being in the other state. The experimental work also finds some support for a temptation effect on consumers on diets, even when consumers sample a good only once, and that sample is imposed on them by an external force.

For individuals, the results highlight the value of Ainslie-style personal rules as a welfare-improving strategy. Knowing that they risk entering a hot state if they succumb to temptation, it is optimal for cold selves to adopt rules limiting indulgence, even though it would not be for a rational consumer without multiple selves. This sophistication may not be common among consumers, however: other work has suggested that the cold self often naïvely finds it difficult to anticipate the preferences and choices of the hot self, reducing the appeal of personal rules as a self-control technique (Loewenstein, 2005). If, in addition, an exogenous effect of some sort imposes an indulgence despite the personal rule, or if consumers face such a strong temptation they are unable to resist it, then after entering the hot state consumers might be vulnerable to binge-like behaviour, as the hot self increases consumption in order to avoid re-entering a cold state.

Policymakers can also derive useful lessons. First, interventions to stop early temptations are an important policy tool. Helping consumers resist an initial temptation can help them avoid the snowball effect of entry into a hot state, while creating an incentive system allowing for or encouraging minor transgressions may actually lower welfare if it induces a hot state. Particularly if such behaviours eventually culminate in a long-run habit, optimizing conditions to help consumers avoid a hot state is valuable.

Second, the experiment provides some insight into the kinds of policy tools that may be effective. Behaviours that are imposed by the policymaker can still affect future decisions and even preferences about the future. This means policy interventions may be a meaningful way of beginning desirable behaviour patterns. By the same token, however, it means that if the government imposes 'bad' behaviour on consumers as a temporary

measure – encouraging overspending as a stimulus plan, for example – such behaviours may induce a hot state that promotes negative behaviour.

Extending the standard model to incorporate consumption persistence can also help us understand economic behaviour. Most obviously, we frequently observe marketers giving away free samples or trial runs of their product. This strategy has the potential to be an effective one, particularly when the product is something people find tempting but are attempting to resist purchasing. Such free trials, such as the freemium model of game distribution which allows a player to begin playing but then requires a payment for an enhanced version of play, provide consumers with information but may also induce a hot state that increases their willingness to pay.

Finally, the experiment suggests several lines of further inquiry. In particular, a number of possible mechanisms for the effect exist, and more work is needed to confirm which – if any – of these potential channels serves to explain the temptation effect observed in the lab. Field experiments present a promising method to study these ideas in other contexts and in a more long-run setting. Path dependence can pose a significant challenge to consumer behaviour and policy design, and both policymakers and consumers can benefit from considering it in a sophisticated way as they plan their behaviour.

## CONCLUSION

---

How preferences evolve over time and how policymakers can improve long-term outcomes represent important areas of research. This paper has aimed to contribute to that work in several ways.

First, given the popularity of behavioural nudges with policymakers, I take a systematic look at the welfare implications of nudges to change behaviour. Though considerable work has been done on nudging, much of it focuses on the effectiveness of various policy tools in changing behaviour, and not on the potential welfare implications. This paper therefore contributes to the ability of policymakers to decide when nudging is useful. Though nudges can be a powerful tool when used well, I also find reason for caution: because nudges can discourage some consumers from optimizing, they can actually reduce welfare in some circumstances. A well-informed policymaker can avoid this pitfall, but doing so may require substantial information about consumer preferences. Policymakers are welcome to continue using nudges as a policy tool, but they are neither a key for all locks nor are they appropriate without sufficient cost-benefit analysis to ensure they will increase, and not reduce, welfare.

Second, I look at endogenous preferences. Much of economics assumes stationarity of preferences: even models that allow for non-stationarity tend to assume preferences change in a predictable, exogenous fashion. If consumers can affect their preferences by their behaviour, however, then their incentives may vary. In particular, consumers may risk being trapped in a negative pattern of behaviour that becomes self-reinforcing, as it induces preferences that minimize the perceived costs of such behaviour. At an extreme, consumers may even consume all their resources immediately. This has useful implications for how policymakers should deal with consumers in long-run harmful patterns of behaviour, such as obesity, carbon emissions, or low educational achievement, and policymakers can benefit from attempting to shift consumer's behaviour between locally optimal consumption paths to reach the global optimum.

Finally, I consider the possibility of a game between multiple selves through the induction of a short-term hot state of consumption complementarity. Such short-term pat-

terns can describe much behaviour of economic and individual interest, and may also have implications for behaviour over the long term. Sophisticated consumers have an incentive to abstain from some goods more than is otherwise optimal in order to avoid entering a hot state: once in a hot state, the consumer has an incentive to binge beyond what is otherwise optimal, to avoid returning to the cold state. My experimental work serves to provide new evidence for such an effect.

This thesis is therefore intended to contribute to the literature in several ways. It focuses on a relatively understudied area of consumer behaviour, time-inconsistent preferences, and introduces formal models of several related phenomena in order to develop testable predictions and ideas that will allow for further study. It also aims to derive lessons for policymakers, both on how to use current policy tools in an optimal manner and introducing areas where policy might have the potential to be welfare-improving by being designed around consumer preferences. Finally, it uses experimental work to begin the testing process of these theories, and study their implications for behaviour empirically.

It also suggests a number of lines for further inquiry. The lab presents a valuable place to test the assumptions that underlie economic models, but all three chapters present ideas and implications that can also be tested in the field and with data. Other work has suggested preferences might be affected by external stimuli, such as economic institutions (Bowles, 1998): understanding how such exogenous effects might interact with individual optimization presents an interesting area for further research, both theoretically and in the field. One particularly interesting case, for example, would be how student preferences change with their field of academic study. Other significant life decisions, such as home-ownership or a change in employment status, might also be reasonably expected to affect preferences, effects that could be tested for. In addition, several more theoretical areas remain for further research. In Chapter 3, for example, the assumption of isoelasticity provides tractability but limits the scope of questions the model can address: further analysis is needed of how consumers behave when their decisions affect their future preferences, which would allow them to commit their future selves to types of behaviour and so use the endogeneity of preferences as a form of internal commitment device.

Economics should not ignore the idea that preferences may change over time and with stimuli, and the implications this has for behaviour. It also presents, however, a fundamental challenge to welfare analysis and to making normative recommendations, because there is no longer one set of preferences by which outcomes can be judged. In many circumstances, this challenge may be relatively minor, but when confronting issues of public health, addiction, saving, environmental policy, and other fields that require long-term planning and judgment, it is frequently fundamental. Better understanding of it represents an important and still-emerging field.

## BIBLIOGRAPHY

---

- Ainslie, George.** 1975. "Specious reward: A behavioral theory of impulsiveness and impulse control." *Psychological Bulletin*, 82(4): 463–496.
- Ainslie, George.** 1992. *Picoeconomics*. Cambridge University Press.
- Akerlof, George, and Rachel Kranton.** 2010. *Identity Economics*. Princeton University Press.
- Akerlof, George, and William Dickens.** 1982. "The Economic Consequences of Cognitive Dissonance." *The American Economic Review*, 72(3): 307–19.
- Akst, Daniel.** 2011. *We Have Met the Enemy: Self-Control in an Age of Excess*. The Penguin Press.
- Allcott, Hunt.** 2014. "Paternalism and Energy Efficiency: An Overview." *NBER Working Paper*, , (20363).
- Allcott, Hunt, Sendhil Mullainathan, and Dmitry Taubinsky.** 2014. "Energy Policy with Externalities and Internalities." *Journal of Public Economics*, 112: 72–88.
- Altmann, Steffen, Armin Falk, Paul Heidhues, and Rajshri Jayaraman.** 2014. "Defaults and Donations: Evidence from a Field Experiment." *IZA Discussion Paper No. 8680*.
- Ariely, Dan, and George Lowenstein.** 2006. "The Heat of the Moment: The Effect of Sexual Arousal on Sexual Decision Making." *Journal of Behavioral Decision Making*, 19: 87–98.
- Arrow, Kenneth J.** 1970. "The Theory of Risk Aversion." In *Essays in the Theory of Risk-Bearing*. North-Holland Publishing Company.
- Bachelier, Louis.** 1900. "The Theory of Speculation." *PhD Thesis*.
- Badger, Gary J., Warren K. Bickel, Louis A. Giordano, Eric A. Jacobs, George Loewenstein, and Lisa Marsch.** 2007. "Altered states: The impact of immediate craving on the valuation of current and future opioids." *Journal of Health Economics*, 26: 865–876.
- Baicker, Katherine, Sendhil Mullainathan, and Joshua Schwartzstein.** 2015. "Behavioral Hazard in Health Insurance." Dartmouth College.
- Baumeister, Roy, and John Tierney.** 2011. *Willpower: rediscovering our greatest strength*. The Penguin Press.
- Baumeister, Roy F., Ellen Bratslavsky, Mark Muraven, and Dianne M. Tice.** 1998. "Ego Depletion: Is the Active Self a Limited Resource?" *Journal of Personality and Social Psychology*, 74(5): 1252–65.
- Becker, Gary, and Kevin Murphy.** 1988. "A Theory of Rational Addiction." *Journal of Political Economy*, 96(4): 675–700.
- Becker, Gary S., and Casey B. Mulligan.** 1997. "The Endogenous Determination of Time Preference." *The Quarterly Journal of Economics*, 112(3): 729–58.

- Beggs, Alan, and Paul Klemperer.** 1992. "Multi-Period Competition with Switching Costs." *Econometrica*, 60(3): 651–666.
- Benartzi, Shlomo, and Richard H. Thaler.** 2001. "Naive Diversification Strategies in Defined Contribution Saving Plans." *The American Economic Review*, 91(1): 79–98.
- Benhabib, Jess, Alberto Bisin, and Andrew Schotter.** 2010. "Present-Bias, Quasi-Hyperbolic Discounting, and Fixed Costs." *Games and Economic Behaviour*, 69(2): 205–223.
- Bentham, Jeremy.** 1789. *An Introduction to the Principles of Morals and Legislation*. London: T. Payne.
- Bernheim, B. Douglas, and Antonio Rangel.** 2004. "Addiction and Cue-Triggered Decision Processes." *American Economic Review*, 94(5): 1558–1590.
- Bernheim, B. Douglas, and Antonio Rangel.** 2009. "Beyond Revealed Preference: Choice Theoretic Foundations for Behavioral Welfare Economics." *Quarterly Journal of Economics*, 124(1): 51–104.
- BIT.** 2011-2012. "Behavioural Insights Team Annual Update 2011-2012."
- Bénabou, Roland, and Jean Tirole.** 2002. "Self Confidence and Personal Motivation." *The Quarterly Journal of Economics*, 117(3): 871–915.
- Bénabou, Roland, and Jean Tirole.** 2004. "Willpower and Personal Rules." *Journal of Political Economy*, 112(4): 848–887.
- Bowles, Samuel.** 1998. "Endogenous Preferences: The Cultural Consequences of Markets and other Economic Institutions." *Journal of Economic Literature*, 36: 75–111.
- Bratman, Michael.** 1999. *Faces of Intention: Selected Essays on Intention and Agency*. Cambridge University Press.
- Brookes, ST, E. Whitley, TJ Peters, PA Mulheran, M Egger, and Davey Smith G.** 2001. "Subgroup analyses in randomised controlled trials: quantifying the risks of false-positives and false-negatives." *Health Technology Assessment*, 5(33): 1–56.
- Carroll, Christopher D., Jody Overland, and David N. Weil.** 2000. "Saving and growth with habit formation." *American Economic Review*, 90: 341–355.
- Chetty, Raj, Adam Looney, and Kory Krof.** 2009. "Salience and Taxation: Theory and Evidence." *American Economic Review*, 99(4): 1145–1177.
- Choi, James, David Laibson, Brigitte Madrian, and Andrew Metrick.** 2003. "Optimal Defaults." *The American Economic Review*, 93(2): 180–5.
- Choi, James J., David Laibson, Brigitte C. Madrian, and Andrew Metrick.** 2004. "For Better or for Worse: Default Effects and 401(k) Savings Behavior." In *Perspectives on the Economics of Aging.*, ed. David A. Wise. University of Chicago Press.
- Congdon, William, Jeffrey Kling, and Sendhil Mullainathan.** 2011. *Policy and Choice: Public Finance Through the Lens of Behavioural Economics*. Brookings Institution Press: Washington, D.C.

- Conlisk, John.** 1996. "Why Bounded Rationality?" *Journal of Economic Literature*, 34(2): 669–700.
- Constantinides, George M.** 1990. "Habit Formation: A Resolution of the Equity Premium Puzzle." *The Journal of Political Economy*, 98(3): 519–543.
- Crawford, Ian, and Matthew Polisson.** 2014. "Testing for intertemporal nonseparability." *Journal of Mathematical Economics*, 52: 46–49.
- Cronqvist, Henry, and Richard Thaler.** 2004. "Design Choices in Privatized Social-Security Systems: Learning from the Swedish Experience." *The American Economic Review*, 94(2): 424–8.
- DECC, Department Energy & Climate Change.** 2013. "Removing the Hassle Factor associated with loft insulation: Results of a behavioural trial."
- Dhar, Ravi, and Klaus Wertenbroch.** 2012. "Self-Signaling and the Costs and Benefits of Temptation in Consumer Choice." *Journal of Marketing Research*, 49(1): 15–25.
- Diamond, Peter A.** 1982. "Aggregate Demand Management in Search Equilibrium." *Journal of Political Economy*, 90(5): 881–894.
- Duckworth, A. L., and M. L. A Kern.** 2011. "A meta-analysis of the convergent validity of self-control measures." *Journal of Research in Personality*, 45(3): 259–268.
- Duckworth, Angela L., and Martin E.P. Seligman.** 2005. "Self-Discipline Outdoes IQ in Predicting Academic Performance of Adolescents." *Psychological Science*, 16(12): 939–44.
- Ellsberg, Daniel.** 1961. "Risk, Ambiguity, and the Savage Axioms." *The Quarterly Journal of Economics*, 75(4): 643–69.
- Fama, Eugene.** 1965. "Random Walks In Stock Market Prices." *Financial Analysts Journal*, 21(5): 55–59.
- Festinger, Leon.** 1962. *A Theory of Cognitive Dissonance*. Stanford University Press.
- Fischbacher, Urs.** 2007. "z-Tree: Zurich Toolbox for Ready-made Economic Experiments." *Experimental Economics*, 10(2): 171–178.
- Frederick, Shane, George Loewenstein, and Ted O'Donoghue.** 2002. "Time Discounting and Time Preference: A Critical Review." *Journal of Economic Literature*, XL: 351–401.
- Fudenberg, Drew, and David K. Levine.** 2006. "A Dual-Self Model of Impulse Control." *American Economic Review*, 96(5): 1449–1476.
- Fuhrer, Jeffrey C.** 2000. "Habit formation in consumption and its implications for monetary-policy models." *American Economic Review*, 90(3): 367–390.
- Gailliot, Matthew T., Roy F. Baumeister, C. Nathan DeWall, Jon K. Maner, E. Ashby Plant, Dianne M. Tice, Lauren E. Brewer, and Brandon J. Schmeichel.** 2007. "Self-Control Relies on Glucose as a Limited Energy Source: Willpower Is More Than a Metaphor." *Journal of Personality and Social Psychology*, 92(2): 325–336.

- Gale, D.** 1967. "A Geometric Duality Theorem with Economic Applications." *Review of Economic Studies*, 34: 19–24.
- Gigerenzer, Gerd, and Reinhard Selten.** 2001. "Rethinking Rationality." In *Bounded Rationality: The Adaptive Toolbox.*, ed. Gerd Gigerenzer and Reinhard Selten. MIT Press.
- Goldman, Steven.** 1980. "Consistent Plans." *The Review of Economic Studies*, 47(3): 533–537.
- Goldstein, Daniel, Gerd Gigerenzer, Robin Hogarth, Alex Kacelnik, Yaakov Kareev, Gary Klein, Laura Martignon, John Payne, and Karl Schlag.** 2001. "Group Report: Why and When do Simple Heuristics Work?" In *Bounded Rationality: The Adaptive Toolbox.*, ed. Gerd Gigerenzer and Reinhard Selten. MIT Press.
- Grinstead, Charles M., and J. Laurie Snell.** 2003. *Introduction to Probability.* American Mathematical Society.
- Gul, Faruk, and Wolfgang Pesendorfer.** 2001. "Temptation and Self-Control." *Econometrica*, 69(6): 1403–35.
- Gul, Faruk, and Wolfgang Pesendorfer.** 2007. "Harmful Addiction." *Review of Economic Studies*, 74(1): 147–172.
- Hardisty, David J., Katherine Thompson, David Krantz, and Elke U. Weber.** 2013. "How to Measure Discount Rates? An Experimental Comparison of Three Methods." *Judgment and Decision Making*, 8(3): 236–249.
- Harris, Christopher, and David Laibson.** 2001. "Dynamic Choices of Hyperbolic Consumers." *Econometrica*, 69(4): 935–57.
- Havranek, Tomas, Roman Horvath, Zuzana Irsova, and Marek Rusnak.** 2013. "Cross-Country Heterogeneity in Intertemporal Substitution." *IES Working Paper.*
- Hilario, Monica R. F., and Rui M. Costa.** 2008. "High on Habits." *Frontiers in Neuroscience*, 2(2): 208–17.
- Houde, Sébastien.** 2014. "How Consumers Respond to Environmental Certification and the Value of Energy Information." *NBER Working Paper*, , (20019).
- Huberman, Gur, Sheena Iyengar, and Wei Jang.** 2004. "How Much Choice is Too Much?: Contributions to 401(k) Retirement Plans." In *Pension Design and Structure: new lessons from behavioural finance.*, ed. Olivia Mitchell and Stephen Utkus. Oxford University Press: Oxford.
- Iyengar, Sheena S., and Mark Lepper.** 2000. "When Choice is Demotivating: Can One Desire Too Much of a Good Thing?" *Journal of Personality and Social Psychology*, 79(6): 995–1006.
- Job, Veronika, Carol S. Dweck, and Gregory M. Walton.** 2010. "Ego Depletion-Is It All in Your Head? Implicit Theories About Willpower Affect Self-Regulation." *Psychological Science*, 21(11): 1686–1693.
- Kahneman, Daniel.** 2011. *Thinking Fast and Slow.* Penguin Group: London.
- Kahneman, Daniel, and Amos Tversky.** 1979. "Prospect Theory: An analysis of decision under risk." *Econometrica*, 47(2): 263–292.

- Kahneman, Daniel, and Amos Tversky.** 1984. "Choices, Values, and Frames." *American Psychologist*, 34(5): 341–50.
- KIPP.** 2011. "The Promise of College Completion: KIPP's Early Successes and Challenges."
- Knowlton, Barbara, Jennifer Mangels, and Larry Squire.** 1996. "A Neostriatal Habit Learning System in Humans." *Science*, 273(5280): 1399–402.
- Laibson, David.** 1997. "Golden Eggs and Hyperbolic Discounting." *The Quarterly Journal of Economics*, 112(2): 443–78.
- Laibson, David.** 2001. "A Cue-Theory of Consumption." *Quarterly Journal of Economics*, 116(1): 81–119.
- Lally, Philippa, Cornelia H. M. van Jaarsveld, Henry W. W. Potts, and Jane Wardle.** 2009. "How are habits formed: Modelling habit formation in the real world." *European Journal of Social Psychology*, 40(6): 998–1009.
- Lashley, K.S.** 1923. "The behavioristic interpretation of consciousness." *The Psychological Review*, 30(4): 237–406.
- Lipman, Barton.** 1991. "How to Decide How to Decide How To...: Modeling Limited Rationality." *Econometrica*, 59(4): 1105–1125.
- Loewenstein, George.** 2005. "Hot-Cold Empathy Gaps and Medical Decision Making." *Health Psychology*, 24: S49–S56.
- Loewenstein, George, and Ted O'Donoghue.** 2004. "Animal Spirits: Affective and Deliberative Processes in Economic Behavior." *Unpublished*.
- Mangels, Jennifer A., Brady Butterfield, Justin Lamb, Catherine Good, and Carol S. Dweck.** 2006. "Why do beliefs about intelligence influence learning success? A social cognitive neuroscience model." *Social Cognitive and Affective Neuroscience*, 1(2): 75–86.
- Metcalf, Janet, and Walter Mischel.** 1999. "A Hot/Cool-System Analysis of Delay of Gratification: Dynamics of Willpower." *Psychological Review*, 106(1): 3–19.
- Mirrlees, James.** 1987. "Economic Policy and Nonrational Behaviour." Working Paper 8728.
- Mischel, Walter, Ebbe Ebbesen, and Antonette Raskoff Zeiss.** 1972. "Cognitive and attentional mechanisms in delay of gratification." *Journal of Personality and Social Psychology*, 21(2): 204–218.
- Modigliani, F., and R. Brumberg.** 1954. "Utility analysis and the consumption function: An interpretation of cross-section data." In *Post-Keynesian Economics*, ed. Kenneth K. Kurihara. Routledge.
- Muraven, Mark, and Roy Baumeister.** 2000. "Self-Regulation and Depletion of Limited Resources: Does Self-Control Resemble a Muscle?" *Psychological Bulletin*, 126(2): 247–59.
- Muraven, Mark, Roy F. Baumeister, and Dianne M. Tice.** 1999. "Longitudinal Improvement of Self-Regulation Through Practice: Building Self-Control Strength Through Repeated Exercise." *The Journal of Social Psychology*, 139(4): 446–57.

- Nassera, Jennifer A., Lauren E. Bradleya, Jessica B. Leitzscha, Omar Chohana, Kristy Fasuloa, Josie Hallera, Kristin Jaegera, Benjamin Szulanczyka, and Angelo Del Parigi.** 2011. "Psychoactive effects of tasting chocolate and desire for more chocolate." *Physiology & Behavior*, 104(1): 117–121.
- Nestler, Eric J.** 2001. "Molecular basis of long-term plasticity underlying addiction." *Nature Reviews Neuroscience*, 2: 119–128.
- Oaten, Megan, and Ken Cheng.** 2006. "Improved Self-Control: The Benefits of a Regular Program of Academic Study." *Basic and Applied Social Psychology*, 28(1): 1–16.
- O'Donoghue, Ted, and Matthew Rabin.** 1999. "Doing It Now or Later." *The American Economic Review*, 89(1): 103–124.
- O'Donoghue, Ted, and Matthew Rabin.** 2002. "Addiction and Present-Biased Preferences." *Working Papers*, 2-10: 1–56.
- O'Donoghue, Ted, and Matthew Rabin.** 2003. "Studying Optimal Paternalism, Illustrated by a Model of Sin Taxes." *American Economic Review*, 93(2): 186–191.
- Peleg, Bezalel, and Menahem Yaari.** 1973. "On the Existence of a Consistent Course of Action when Tastes are Changing." *The Review of Economic Studies*, 40(3): 391–401.
- Phelps, E.S., and R.A. Pollak.** 1968. "On Second-Best National Saving and Game-Equilibrium Growth." *The Review of Economic Studies*, 35(2): 185–199.
- Pollak, R.A.** 1968. "Consistent Planning." *The Review of Economic Studies*, 35(2): 201–208.
- Radner, Roy, and Michael Rothschild.** 1975. "On the Allocation of Effort." *Journal of Economic Theory*, 10(3): 358–76.
- Redelmeier, Donald, and Eldar Shafir.** 1995. "Medical Decision Making in Situations That Offer Multiple Alternatives." *The Journal of the American Medical Association*, 273(4): 302–5.
- Sallee, James M.** 2014. "Rational Inattention and Energy Efficiency." *Journal of Law and Economics*, 57(3): 781–820.
- Samuelson, Paul.** 1937. "A Note on Measurement of Utility." *Review of Economic Studies*, 4: 155–161.
- Schelling.** 1978. "Economics, or the Art of Self-Management." *AER*, 68(2): 290–294.
- Schwartz, Barry.** 2004. *The Paradox of Choice*. HarperCollins: New York.
- Simon, Herbert.** 1955. "A Behavioral Model of Rational Choice." *The Quarterly Journal of Economics*, 69(1): 99–118.
- Smith, Adam.** 1759. *The Theory of Moral Sentiments*. Printed for A. Millar, in the Strand; and A. Kincaid and J. Bell, in Edinburgh.
- Stigler, George J.** 1961. "The Economics of Information." *The Journal of Political Economy*, 69(3): 213–25.
- Stokey, Nancy, and Robert Lucas.** 1989. *Recursive Methods in Economic Dynamics*. Harvard University Press.

- Stroop, John Ridley.** 1935. "Studies of interference in serial verbal reactions." *Journal of Experimental Psychology*, 18(6): 643–662.
- Strotz, R.H.** 1955-1956. "Myopia and Inconsistency in Dynamic Utility Maximization." *The Review of Economic Studies*, 23(3): 165–180.
- Tangney, June P., Roy F. Baumeister, and Angie Luzio Boone.** 2004. "High Self-Control Predicts Good Adjustment, Less Pathology, Better Grades, and Interpersonal Success." *Journal of Personality*, 72(2): 271–323.
- Thaler, Richard, and Cass Sunstein.** 2008. *Nudge*. Yale University Press: New Haven.
- Thaler, Richard, and H.M. Shefrin.** 1981. "An Economic Theory of Self-Control." *Journal of Political Economy*, 89(2).
- Toffler, Alvin.** 1970. *Future Shock*. The Bodley Head Ltd.
- Tricomi, Elizabeth, Bernard W. Balleine, and John P. O'Doherty.** 2009. "A specific role for posterior dorsolateral striatum in human habit learning." *Eur J Neurosci*, 29(11): 2225–2232.
- Trust, Energy Saving.** 2013. "Roof and Loft Insulation."
- Tsvetanov, Tsvetan, and Kathleen Segerson.** 2014. "The Welfare Effects of Energy Efficiency Standards When Choice Sets Matter." *Journal of the Association of Environmental and Resource Economists*, 1(1/2): 233–271.
- Tversky, Amos, and Daniel Kahneman.** 1974. "Judgement under Uncertainty: Heuristics and Biases." *Science*, 185(4157): 1124–31.
- Tyson, Christopher J.** 2008. "Management of a capital stock by Strotz's naive planner." *Journal of Economic Dynamics & Control*, 32: 2214–2239.
- Van Boven, L, and G. Lowenstein.** 2003. "Social Projection of Transient Visceral Feelings." *Personality and Social Psychology Bulletin*, 29: 1159–1168.
- Verplanken, Bas, and Wendy Wood.** 2006. "Interventions to Break and Create Consumer Habits." *Journal of Public Policy & Marketing*, 25(1): 90–103.
- von Weizsacker, Carl Christian.** 1971. "Notes on Endogenous Change of Tastes." *Journal of Economic Theory*, 3: 345–372.
- von Weizsacker, Carl Christian.** 2013. "Freedom, Wealth and Adaptive Preferences." Max-Planck-Institute for Research on Collective Goods.
- Wang, XT, and RD Dvorak.** 2010. "Sweet future: fluctuating blood glucose levels affect future discounting." *Psychological Science*, 21(2): 183–188.
- Zauberman, Gal, and John G. Lynch.** 2005. "Resource Slack and Propensity to Discount Delayed Investments of Time Versus Money." *Journal of Experimental Psychology: General*, 134(1): 23–37.

## CHOCOLATES

---

Figure 4: Chocolates



## EXPERIMENTAL INSTRUCTIONS

<All Subjects assembled in Waiting Room>

Good afternoon everyone! Welcome to the CESS Lab. I'm Nick Chesterley, running the experiment: feel free to get in touch with me if you have any concerns or questions. You'll be given a handout at the end with my contact information on it. The experiment is intended to give us some insight into your behaviour in a number of different contexts. Please respond to the questions as honestly and thoroughly as possible. Today we're doing to do something a bit different; I'm going to take you to the lab one at a time, and you'll have a chance to randomly pick a workstation. My apologies for those of you who have to wait: we'll be as quick as possible. If you'd like to use your phones while you wait, that's fine, though I'll ask you to turn them off later on. If NAME could come with me, we'll start.

<Start taking people one by one into side room to pick a workstation>

Please choose a card to determine which workstation you'll be at. <Choose a card>

Thank you. Now, this experiment will also involve food. Do you have any food allergies? <Wait for answer>

Thank you.

If Treatment Group: Would you please try one of these chocolates, and rate it on a scale from 1-7? Thank you. Please take the rest of them to your workstation: feel free to have as many as you'd like. Apologies if you have to wait a few moments before we finish seating the others. <Subject is given the rest of the case to take to their workstation, and escorted to their seat>

If Control Group: Please take these chocolates to your workstation: feel free to have as many as you'd like. Apologies if you have to wait a few moments before we finish seating the others. <Subject is given the rest of the case to take to their workstation, and escorted to their seat>

<Once all subjects are at workstations> At this time, I'd ask everyone to switch off their phones, tablets, etc: the experiment only continues when everyone has completed the stage they're at, so please pay attention. Everyone should currently have a grey screen in front of them, once they move the mouse. If you don't, please raise your hand now. <Check>

<Start experiment> Everyone should have the first screen of the experiment in front of them, with an OK button in the bottom right of the screen. If you don't have that, please raise your hand now... Good. Now, please begin the experiment by clicking the okay button. <Participants work through experiment> <Participants reach end of Results Screen> Thank you. You are now looking at the results of your efforts, and the corresponding payments. If anyone has any questions, please feel free to raise your hand and I'll come over. <Collect remaining chocolate cups>

We now have a few questions for you to answer: please answer them honestly and thoroughly. For many of them, there are no right or wrong answers: we are interested in your ideas. Thank you all for your participation. We will call you by workstation number to receive your payment.