Essays in Macroeconomic Policy and Behavioural Bias

Patrick Schneider

28 March 2025

A thesis presented for the degree of Doctor of Philosophy

Department of Economics
London School of Economics and Political Sciences
United Kingdom

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science (LSE) is solely my own work other than where I have clearly indicated that it is the work of others.

Chapters 1 and 2 are joint work with Patrick Moran (Federal Reserve Board of Governors). In these chapters, I contributed three quarters and a half of the work, respectively.

This thesis consists of 27,186 words.

Abstract

This thesis comprises three papers that examine the tension between macroeconomic policy and retirement policy in the context of behavioural bias. In the first paper, joint work with Patrick Moran, we develop a theoretical model to assess the distributional and welfare implications of granting early access to retirement accounts as a means of stimulating consumption, contrasting this approach with traditional fiscal stimulus measures. Using a heterogeneous agent model, we demonstrate that while household liquidity policy is an effective and popular stimulus tool, it burdens the poorer and more present-biased workers with the future costs of stimulus, making it more regressive than conventional fiscal policy. In the second paper, we explore how self-control issues influenced participation in an early withdrawal program during the COVID–19 pandemic, confirming that such behavioural factors played a significant role—a critical insight given that the typical illiquidity of retirement systems is often justified by concerns over individuals' self–control limitations. In the third paper, I compare two approaches to modelling present bias—quasi-hyperbolic discounting and temptation preferences. By recasting the latter in continuous time, I can directly compare the two frameworks, showing that quasi–hyperbolic discounting is a special case of temptation preferences under some common assumptions. Whilst being behaviourally equivalent, they are not welfare equivalent and so distinguishing between them is important. Differences in the behaviour of biased agents who are sophisticated provide opportunities for identification between the two approaches.

Contents

List of Figures

List of Tables

Dedication

For Nussi, راوي, and سمسم, the loves of my life.

Acknowledgements

I am extremely grateful to my supervisor Ricardo Reis for his guidance, mentorship, and support, and for the support of my other advisers Wouter den Haan and Ben Moll. I also gratefully acknowledge the financial support of the LSE Economics Department.

I thank my co–authors, Patrick Moran and Soroush Sabet, whose interventions in my research journey were perfectly timed and calibrated, and without whom this dissertation would simply not have been possible.

Chapter 1

Household Liquidity Policy

Patrick Schneider and Patrick Moran[1]

Abstract

We assess 'household liquidity policy', a novel approach to stimulating aggregate demand that relies on relaxed regulation instead of conventional fiscal tools. We analyse the effectiveness of these liquidity policies, focusing on a form that was widely used during the Covid–19 pandemic: early access to retirement savings accounts. In a heterogeneous agent model with retirement and present–biased households we find both liquidity and conventional fiscal policies can achieve similar boosts to aggregate consumption but have different distributional implications. Relative to fiscal policy, liquidity policy benefits wealthier workers, retirees, and future generations, due to its lower tax burden and added flexibility, but it is also regressive. Liquidity policy shifts the future financial burden of present–day stimulus onto poorer and more present–biased workers, who only feel the impact when it is too late to adjust.

[1]The views expressed in this chapter paper are solely those of the authors and do not represent the views of the Federal Reserve Board or the Federal Reserve System.

1.1 Introduction

Suppose the government wants to stimulate aggregate demand. The conventional way to do this is with government transfers or spending, funded by deficits. Such 'fiscal stimulus' boosts demand in part because people who are liquidity constrained tend to have high marginal propensities to consume (MPC) and so they spend more when they receive extra income (Aguiar et al., 2024; Carroll et al., 2021; Fagereng et al., 2021; Kaplan and Violante, 2014). The key to this mechanism is that policy relaxes household liquidity constraints, not necessarily that households receive more income. In principle, other policies that relax liquidity constraints can also stimulate aggregate demand, even if such policies leave total household resources unchanged. One example of this alternative approach is to give households early access to their otherwise illiquid retirement savings during periods of aggregate distress.[2] These policies were used sporadically for many years, but became common during Covid–19, with more than thirty countries granting some form of relief from retirement saving regulations (OECD, 2021). We call stimulus policy that uses regulation, rather than the government budget, 'household liquidity policy'.

In this paper, we compare household liquidity policy with conventional fiscal stimulus. Focusing on liquidity policies that relax retirement regulations, our paper makes two main sets of contributions. First, we develop a heterogeneous–agent model that captures the key tradeoffs to each of the different approaches, providing the first comparison of liquidity policy with fiscal stimulus. These tradeoffs stem from the difference in funding—fiscal policy is funded with future taxes that cause distortions and redistribution, whereas liquidity policy is self–funded, thereby causing no distortions nor redistribution, but also reducing the ability to commit resources to the future for those who need it.

Our second contribution is to show that, whilst both approaches are capable of the same short–term stimulus, liquidity policy is regressive. Comparing the two approaches, liquidity policy is more beneficial to wealthy workers, retirees and future generations, but for reasons unrelated to the stimulus itself—these individuals value the option to rebalance their portfolios, and wish to avoid the future taxes (and their distortions) implied by fiscal policy. Pursuing stimulus via liquidity policy effectively privatises the cost of aggregate demand management, and this cost is paid by the people who actually spend the money (the less wealthy, and more present–biased, workers) when they retire and find they have less to live on. Choosing liquidity policy therefore places the burden of stimulus on the shoulders of precisely the people retirement policy seeks to protect.

To begin, we develop a quantitative model that allows us to compare the aggregate

[2]By no means the only example, others of which include student debt, mortgage repayment, or rent deferrals; changes in regulations limiting collateralised loans; or changes to mandatory pre–payment of income taxes.

and distributional effects of the two different policies. Our starting point is a two–asset heterogeneous–household model featuring idiosyncratic risk and incomplete markets, where households save in a liquid account, with a borrowing constraint, or illiquid account. The two-asset framework helps to ensure that there is demand for liquidity, from households that are close to their constraints.[3] We combine this household with a fiscal authority constrained by a fiscal rule, yielding the standard heterogeneous–agent environment for analysing fiscal stimulus (as in Auclert et al., 2024b; Bayer et al., 2023; Kaplan and Violante, 2014).

To also analyse liquidity policy, we extend this model with three important features to capture the key tradeoffs. First, we include an overlapping–generations life–cycle with work and retirement phases (Blanchard, 1985; Yaari, 1965). This creates a need to provide for retirement, in addition to the standard precautionary saving motive. Second, a portion of the population is subject to naive present–bias, leading to over–consumption (and under–saving) in the present (Laibson, 1997; Maxted, 2024; Maxted et al., 2024). As a result, our model captures one of the principal rationales for government intervention in retirement saving—the view that many individuals are myopic and lack the ability to save for retirement if left entirely to their own devices (Feldstein, 1985). In light of the above, we add realistic retirement policy, which consists of two pillars: first, a tax–funded state pension that goes to all retirees, and second, defined–contribution individual retirement accounts with realistic tax subsidies, restrictions on early withdrawal, and the possibility of mandatory contributions, depending on the calibration. This retirement account replaces the illiquid asset that is standard in two–asset heterogeneous–agent models; here, illiquidity is due to regulations to address myopia, and differences in return are due to tax concessions.

We parameterise the model with standard values, and select the degree of present bias to match empirical estimates of the aggregate MPC.[4] The retirement regulations are set in an optimal policy exercise that sets mandatory contribution rates and tax concessions to maximise the well–being of prospective newborns into the society, subject to a constraint that everyone opts in at the start of their careers. This exercise is able to rationalise the contribution rates and tax concessions we observe in mandatory DC systems around the world. An implication of this calibration is that, in the stationary

[3]Throughout this paper, we will use the term 'liquidity' or 'liquid resources' to refer to resources that can be used close–to immediately for consumption. This would include cash, bank and saving deposits, available consumer credit, and investments in securities that can be sold at will. It excludes real property and other durable investments, holdings of private companies, and contingent assets like insurance policies and most retirement accounts.

[4]We solve the model using a new algorithm developed in Sabet and Schneider (2024), which is monotone and consistent, and so robust to parameter choices. This is in contrast to existing methods for solving continuous–time HA models with multiple endogenous state variables e.g. the drift–splitting approach in Kaplan et al. (2018), which is often unstable.

solution to our model, all but the poorest workers' portfolios are over–invested in the retirement account relative to what they would choose for themselves. This is the price of helping the present–biased save for retirement: they will only save if the money is illiquid, but because the government cannot distinguish type, nor career–stage, everyone is required to save at the same rate into the same illiquid environment.

To perform the stimulus experiment, we shock the model with unanticipated transfers (fiscal policy) and early access to retirement savings (liquidity policy) designed to boost household consumption by the same amount over one quarter. Matching the short–term stimulus from the two approaches allows us to compare their long–run implications. There are three main differences (i) the tax changes driven by the fiscal rule distort inter–temporal consumption smoothing, and this distortion is much greater under fiscal than liquidity policy; (ii) these taxes also cause redistribution, with lower future consumption by retirees and future generations subsidising the transfers received by workers; (iii) liquidity policy undermines retirement adequacy for present–biased workers, reducing their consumption upon retirement, and more so for the individuals who were initially less wealthy.

After describing the mechanisms through which the two policies operate, we then use our model to quantify their relative importance for household well-being. We find the compensating variation that would make each household indifferent between liquidity policy and its fiscal alternative. We show that household liquidity policy is better for wealthy workers, retirees, and future generations. These individuals do not benefit much from traditional fiscal stimulus, but are still liable for higher taxes under fiscal policy. Further, wealthy workers value the ability to rebalance their portfolio more than their less wealthy counterparts. By contrast, fiscal policy is preferred by working households with low wealth, a group disproportionately comprised of present–biased households. Aggregating across this heterogeneity, we find that roughly 70 percent of households prefer liquidity policy over traditional fiscal stimulus in our baseline calibration. In short, liquidity policy may be politically popular despite its regressivity, as it concentrates the costs of stimulus on a relatively small subset of society, namely the present–biased and low–wealth workers.

The degree to which one approach dominates the other depends on how the government sets retirement policy, and its fiscal rule. Aggregating the compensating variations, liquidity policy is marginally better for society the stricter is retirement policy, and the more aggressively the fiscal rule retires government debt. For example, liquidity policy is welfare *improving* for all but the poor present–biased types in our baseline calibration. This is because it allows some portfolio re–balancing, particularly valuable to wealthier workers. In an alternative calibration with lower mandatory contribution rates, wealthy workers are less over–invested in their retirement accounts, and there is much lower benefit from liquidity policy as a result. Similarly, the baseline calibration is based on an exponential fiscal rule that sets a half–life for debt gaps of 14 years (as in Galí, 2020). This rule front–loads the cost of stimulus under fiscal policy more than liquidity policy because debt is 'retired' at a faster rate than workers retire. If the fiscal rule is instead relaxed to match the rate of retirement, or be even looser, then the implied taxes are less onerous and the relative benefit of liquidity policy is reduced.

Related literature. This paper brings together two large strands of literature. On one hand, the influential heterogeneous agent macro literature explains fiscal and monetary policy transmission based on liquidity constraints and the distinction between liquid and illiquid assets (see e.g. Auclert et al., 2024b; Bayer et al., 2019; Kaplan et al., 2018; Kaplan and Violante, 2014). While this literature generally assumes that retirement accounts play an important role in household illiquidity, retirement policy is assumed exogenous, and none of these papers consider policies that alter the illiquidity of these accounts. On the other hand, there is a growing public economics literature that evaluates retirement policy and the optimal degree of illiquidity in retirement saving systems (see e.g. Amador et al., 2006; Andersen et al., 2024; Beshears et al., 2025, 2020; Moser and Olea de Souza e Silva, 2019). While these papers assume that the level of illiquidity in retirement systems is a societal choice, they are largely silent on macroeconomic considerations related to fiscal stimulus. To the best of our knowledge, our paper is the first to offer a positive and normative evaluation of household liquidity policy relative to traditional fiscal stimulus.

A growing empirical literature analyses past episodes of household liquidity policies (Andersen, 2020; Argento et al., 2015; Hamilton et al., 2024; Kreiner et al., 2019a; Preston, 2022; Schneider and Moran, 2024b; Shapiro and Slemrod, 1995). We bring these stimulus packages together under the banner of household liquidity policy and analyse them theoretically in a modeling environment that allows for direct comparison with conventional fiscal policy. This allows for positive and normative comparisons of the two approaches, which may help to design future stimulus packages.

Our work complements Hamilton et al. (2024), which analyses Australia's early withdrawal program during Covid–19, using detailed micro data to identify who withdraws and what they do with the money. The authors show empirically that around one in six working age people withdrew, the modal withdrawal was all of the $20,000 allowed, and these households on average spent 40% of the funds within eight weeks. They argue that this is evidence of present–bias, which they estimate in a structural model. Our paper makes a different but complementary contribution. While the above authors identify the MPC and the strength of present-bias, we take present-bias as given, and instead develop a model that captures the key trade-offs between household liquidity policy and traditional fiscal policy. This allows us to perform the first positive, normative, and distributional comparison of these two different approaches to stimulus.

Our model is also informed by empirical work by Schneider and Moran (2024b), who use a survey-elicited measure of psychological self-control, combined with the early release of retirement wealth in Australia, to estimate the relative importance of behavioral biases versus situational factors in accounting for early withdrawal from retirement accounts. The authors find that self-control heterogeneity plays an important role in predicting early withdrawal, and is a stronger predictor than other behavioural factors such as financial literacy, planning horizons, or personality traits. Overall, individuals in the top quintile of self-control issues are 60% more likely to withdraw than those in the bottom quintile. Informed by these empirical results, we also incorporate heterogeneity in present-bias into our quantitative model, evaluate how the two policies affect long-term retirement adequacy, and examine how the welfare implications of liquidity policy differ for individuals with versus without present-bias.

Some papers use quantitative models to explore the role of retirement accounts in stimulating the economy, but none compare liquidity policy to traditional fiscal policy. Love (2017) proposes a policy to stimulate the economy using counter–cyclical matching to retirement contributions, which he evaluates in a life-cycle model. Graves (2023) develops a HANK model to analyze the flight–to–liquidity that occurs following an increase in unemployment. He conducts one counterfactual exercise showing the effect of lower withdrawal penalties on aggregate consumption during Covid–19. Finally, Kaplan et al. (2020b) explore the tradeoff between health outcomes and economic impacts of policy choices during Covid–19 in the USA. They combine a HANK model with an SIR module of disease transmission, and use it to assess the impact of the various economic and health policies used in the USA. A part of the CARES Act that they model is the USA's removal of the withdrawal penalties from individual retirement accounts, as in Graves (2023), but this is not the focus of their analysis. Our paper (1) characterises the different channels through which household liquidity policy differs from traditional fiscal policy, (2) evaluates the distributional implications of the two policies, and (3) conducts a welfare analysis of the two policies, something that no previous paper has attempted.

We contribute to the broader heterogeneous–agent macro literature by providing a new micro foundation for the illiquid accounts commonly featured in two-asset macro models (Auclert et al., 2024b; Bayer et al., 2019; Kaplan et al., 2018; Kaplan and Violante, 2014). This illiquidity is generally modelled as an exogenous feature of the world, when in reality it is usually a result of government policy. Empirically, household budgets are made up of only two types of genuinely illiquid asset: housing and retirement savings (Fagereng et al., 2019). In both cases, much of the illiquidity is due to regulation, e.g. restrictions or penalties on withdrawals from retirement accounts, and limits to home equity withdrawal. Modelling it as such opens the option for liquidity policy in our environment. Our modelling approach also builds upon Beshears et al. (2025) and Maxted et al. (2024) who show the importance of present–bias for hand-to-mouth behaviour and fiscal policy, but do not consider household liquidity policy.

Finally, we contribute to the literature about retirement system design. A common thread in this literature is that imposed illiquidity is justified to help households overcome biases in their decision–making. The government has a role in mandating some form of retirement saving, and faces a problem of how to optimally balance the long–run commitment that households need against the short–run need for flexibility to insure working–life idiosyncratic risk, and also to balance the welfare of the behaviourally biased against those who are not (Amador et al., 2006; Beshears et al., 2025; Moser and Olea de Souza e Silva, 2019). Our government faces the same type of problem, but is constrained to consider mandatory DC systems, as implemented in many countries (OECD, 2023). We show that the optimal DC system involves mandatory contributions close to those actually observed, but does not involve any tax concessions, in stark contrast to the systems in place. Instead, we rationalise these tax concessions as necessary to encourage people to opt into the system, at least at the beginning of their careers. As such, these concessions serve a political, rather than policy, purpose.

One major feature of the public literature on retirement system design and reform is the impact of retirement systems on the decision to retire. These papers (e.g. Blundell et al., 2016; Kolsrud et al., 2024) emphasise the distorting effects retirement policy can have on labour supply and estimate the optimal design subject to these distortions and the fiscal externalities they impose. We abstract entirely from the retirement decision in this paper. Rather we take retirement to be a fact, which creates a need for extra savings, but we let the event itself arrive randomly. This simplifies the problem by removing a household decision without undermining the focus of our analysis, which is to explore the impact of different approaches to stimulus for a given retirement system.

Road map The paper proceeds as follows. In Section 1.2 we describe defined contribution retirement schemes, the relevant institutional setting for the paper, and identify two key parameters that we will map to our model. In Section 1.3 we detail the model, and calibrate its standard parameters in Section 1.4. In Section 1.5 we set up and solve the government's optimal retirement policy problem and show that this approach rationalises the key parameters identified in Section 1.2. We then turn to the stimulus policy experiments in Section 1.6, showing the two approaches are similar in aggregate, but have different distributional implications. We evaluate their differences with a welfare analysis in Section 1.7, finishing with robustness checks. Section 1.8 concludes.

1.2 Defined contribution retirement accounts

The need to provide for retirement is universal and government intervention to meet some of this need is also common. The government's involvement is justified for a variety of reasons. These include a redistributive motive, to help avoid poverty in retirement, a protective motive, to short–circuit the moral hazard created by the redistributive motive (i.e. households neglect to save for retirement, anticipating the government will bail them out), and a paternalistic one, correcting for biases that reduce working–life saving (Beshears et al., 2015; Feldstein, 1985). The means with which countries address these needs vary a lot, but usually involve some mix of state and private provisions, with the latter becoming increasingly important as many governments grapple with the strain of unfunded state pension provisions coupled with ageing populations (OECD, 2018).

Figure 1.1: Retirement system assets

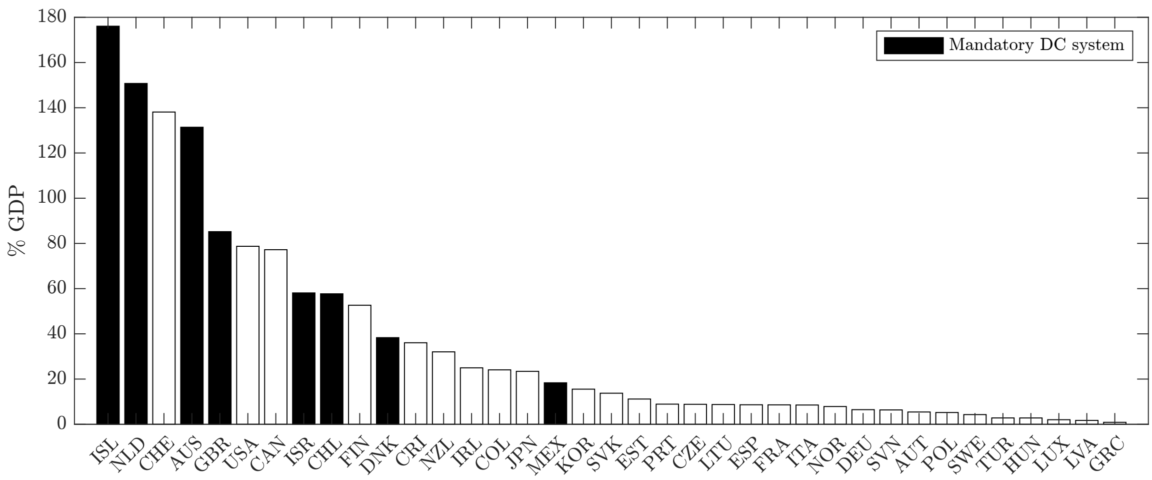

Source: values from OECD 2022 Total Pension Funds' Assets, % of GDP; classification as mandatory DC system from (OECD, 2023, Table 4.2)

Our focus in this paper is on countries with private, mandatory defined contribution (DC) schemes.[5] In DC pension systems, working–age people make contributions into regulated investment accounts, and they have a claim on the balance and accumulated returns upon retirement; these systems can build up substantial resources, as illustrated in Figure 1.1, which plots total retirement assets across OECD countries. The exact design features of DC accounts differ across countries, but they can be broadly understood in terms of rules defining (a) liquidity during working life, (b) contributions, (c) tax treatment, and (d) the state pension they are combined with.

Liquidity during working life Regulations affecting access to DC accounts differ across countries. In some settings, like the USA and UK, participants are allowed to

[5]Such schemes are common, and increasingly being adopted as countries attempt to reduce the fiscal burden of state–only systems facing ageing populations (OECD, 2018). Among OECD countries, for example, 20 have some form of regulated, private retirement savings vehicle—Australia, Belgium, Canada, Chile, Costa Rica, Denmark, Estonia, Germany, Iceland, Ireland, Israel, Lithuania, Mexico, Netherlands, New Zealand, Norway, Sweden, Switzerland, United Kingdom, United States (OECD, 2023). Of these, 8 mandate contributions into funded DC schemes—Australia, Chile, Denmark, Iceland, Israel, Mexico, Netherlands, United Kingdom (OECD, 2023), where the UK's contributions are a default, rather than mandatory.

withdraw prior to retirement, but they pay a penalty (10% in the USA, 55% in the UK). In others, like Australia, withdrawals are not allowed except in dire personal circumstances (e.g. terminal illness or extreme financial hardship) effectively making the accounts completely illiquid. In either case, the illiquidity is created by regulation, rather than because the assets are difficult to transact in.

Contributions Regulations affecting contributions generally govern (a) whether any contributions are mandatory and (b) limits on voluntary contributions. Mandatory contributions, when they exist, are usually set as a proportion of pre–tax employment income, commonly in the range of 10–20% (see Table 1.1 for some examples from OECD countries). Voluntary contributions to these schemes are often allowed as well, but are usually limited to maximum nominal amounts per year because they attract tax concessions.

Table 1.1: Mandatory contribution rates in OECD DC systems

	AUS	CHL	DNK	ICE	ISR	MEX	NLD	UK*
ξ (p.p.)	12	10	12	15.5	13	15	18.6	8

Source: OECD (2023) Table 3.4, p. 141. For all OECD countries with privately funded DC schemes, and no other mandatory private system. *The UK's is a default, rather than mandatory.

Tax treatment There are three potentially taxable flows in DC systems—contributions, investment returns, and withdrawals—and systems differ in whether each of these is taxed (potentially at concessional rates) or exempt, leading to a three-letter code describing them. In a system like USA 401(k)s, for example, contributions are made from pre–tax income, and returns are exempt as well, but withdrawals are taxed at the personal marginal tax rate, so it is coded EET. This is the most common approach. By contrast, Australia's Superannuation contributions from pre–tax income are taxed at a concessional 15% rate, as are returns, and withdrawals are tax free, so it is coded TTE.[6]

Figure 1.2 shows the combinations of mandatory contribution rates (ξ) and tax advantages across the collection of OECD countries with mandatory DC systems.[7] The tax advantage variable comes from OECD (2018), and represents the tax savings from a given flow of contributions over a typical working life, relative to if they were saved in a regular investment account. Tax advantages are ubiquitous in these accounts, though they range from quite small (worth a discount of around 10% in Chile) to substantial (around 50% in Israel and Mexico).

[6] In the model introduced in Section 1.3, we use a TTE system so that we can control the tax concession granted for contributions and returns inside the account. EET systems, whilst more common, typically apply taxes at full marginal rates, and so offer fewer degrees of freedom.

[7] Note this excludes countries that mix these with other mandatory private options.

Figure 1.2: DC system design in OECD countries

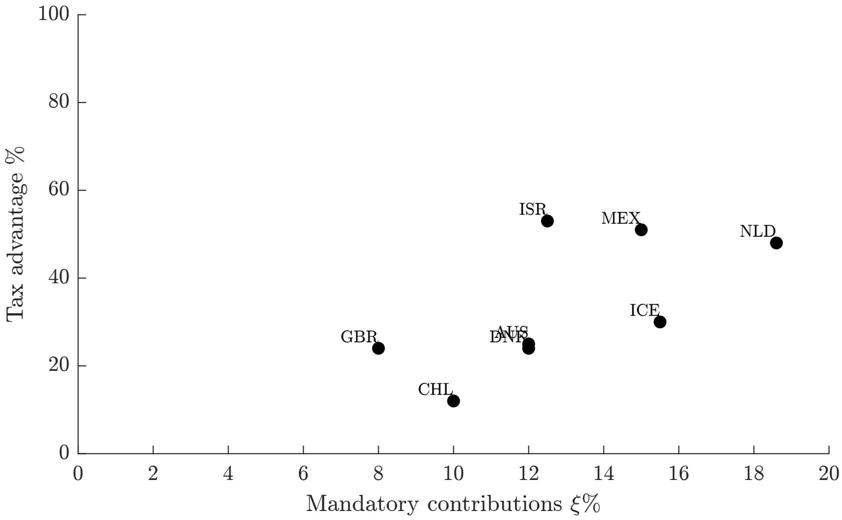

Source: mandatory contribution rates (OECD, 2023, Table 3.4) and tax advantage (OECD, 2018, Table 3.2)

In the model introduced in Section 1.3, the retirement system features parameters encoding mandatory contributions and tax concessions, and we rationalise the features seen here as the optimal policy when society features a portion of the population that is present–biased, and the government needs to ensure workers to opt–in to the system at the start of their careers. The tax concessions, then, serve the *political* purpose of building buy–in into the system.

State pension DC systems are usually introduced to reduce the burden of retirement provision on current taxation by shifting the responsibility onto households themselves. In mandatory DC systems, government retirement provision is usually still present, but less generous. We can see this by looking at the state pension replacement rate for an average income earner across countries with and without mandatory private systems, plotted across OECD countries in Figure 1.3. This replacement rate for an average earner is 31% on average for OECD countries with mandatory systems, compared with 56% on average for OECD countries with only mandatory public systems, a substantial difference.[8]

Early access as stimulus In response to economic crises, many countries have used access to these pools of resources to stimulate demand. Denmark was one of the first to introduce such measures during the Global Financial Crisis, allowing early withdrawals and temporary suspensions of contributions (Kreiner et al., 2019a). This approach be-

[8]These are population weighted averages of the mandatory public gross replacement rates for an average earner across OECD countries with and without a mandatory private system in place (OECD, 2023, Table 4.2)

Figure 1.3: Retirement system replacement rates for average earners, OECD countries

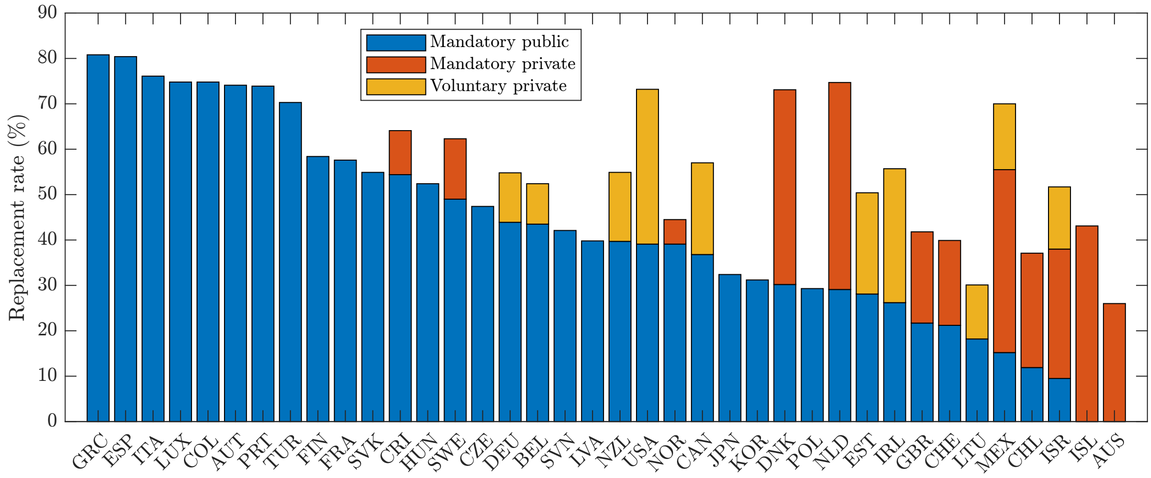

Source: OECD (2023) Table 4.2

came more widespread during the Covid-19 pandemic, in which three approaches were used across at least 31 countries: allowing limited withdrawals when they were previously banned, as seen in Australia, Chile, and Peru (Hamilton et al., 2024; Madeira, 2022; OECD, 2021); removing withdrawal penalties, as in the United States (Graves, 2023); and reducing or deferring mandatory contributions in countries like Singapore, Malaysia, and Vietnam.

These policies resulted in significant outflows from retirement systems, with the most dramatic cases observed in Latin America. In Peru, a staggering 18.3% of assets in retirement savings plans were withdrawn, followed closely by Chile at 14.6%. Iceland and Australia saw smaller, but still substantial, withdrawals during Covid–19 of 3% and 1.4% of assets, respectively (OECD, 2021, p. 27).

In the remainder of the paper we will analyse traditional fiscal and liquidity policy in an economy with government intervention into retirement policy motivated by paternalistic concerns. The model detailed in the next section features present–biased households who unwittingly save too little, and the government responds by forcing them to save in a defined contribution scheme with the features described above—limits on working–life withdrawal, mandatory contributions from labour income, and tax concessions. The liquidity policy we explore in Sections 1.6 to 1.7 is early withdrawal opportunities, like those used in Australia, Peru, and Chile. An alternative would be to reduce mandatory contribution rates. This works as stimulus as well, but it is more regressive because it grants relatively more liquidity to higher earners. We focus on withdrawal opportunities because it is the most directly comparable to fiscal transfers.

1.3 A model of households with illiquid retirement accounts

Our environment is a continuous–time, infinite horizon model featuring a measure of households, and a fiscal authority. Prices (the interest rate and wage) are fixed, and the stationary equilibrium between households and fiscal authority is reached with a tax rate that balances the government budget.

1.3.1 Households

Households are differentiated by their stage of life (working or retired) and four time–varying state variables: their idiosyncratic productivity z and employment state, the balance in their liquid account b, their illiquid account balance a, and their present–bias β. We collect these states into the vector $x = (b, a, z, \beta)$, where z captures life–stage, workforce status and employed productivity.

1.3.1.1 Life–cycle transitions

Households live through working–life and then retirement (Blanchard, 1985; Yaari, 1965), transitioning out of each phase with fixed Poisson intensities (δ_R, δ). When young, they work, and make consumption and asset allocation decisions. In retirement, they make the same decisions but can no longer receive the market wage. This creates a need for provision that is met personally by any assets they retire with, and collectively by the distribution of a fixed state–pension w_R. With these resources, retirees solve a cake–eating problem until they die, and are replaced by workers with no assets.

This lifecycle structure is a substantial simplification of how our careers typically progress. The most important departure from reality is that the transition into retirement is random, rather than a choice. This abstracts from an important part of the public economics literature that looks at the implications for retirement systems on incentives to retire, and the optimal pension reform to achieve affordability in the face of ageing populations (e.g. Kolsrud et al., 2024). Simplifying the transitions to be random gains us tractability, by reducing the households' choice set, without sacrificing structure that's important for our central question; transition rates are calibrated to match the average spans of working life and retirement.

1.3.1.2 Idiosyncratic risk

Working–age households are subject to two types of idiosyncratic risk to their income. First, they jump in and out of employment with Poisson finding and separation intensities (λ_f, λ_s). Second, whilst employed, their log–labour productivity is a diffusion that follows an Ornstein–Uhlenbeck process

$$d \ln z_t = -\theta_z(\ln z_t - \ln \bar{z})dt + \sigma_z dW_t$$

Where θ_z captures its persistence, $\bar{z}$ is the stationary mean, normalised to 1, W_t is a Wiener process, and σ_z is the weight on this noise. Newborn workers are employed, and draw their productivity from the stationary distribution of z. Retired workers are not subject to any idiosyncratic risk beyond the chance of dying.

1.3.1.3 Budget constraints

All households have access to two accounts for storing wealth—a liquid and an illiquid account. At any given point in time, households have two choice variables—consumption $c > 0$, funded from their liquid account, and voluntary transfers between the two accounts $d \in \mathbb{R}$.

Liquid account The law of motion for the liquid account is

$$\dot{b} = r_b(b)b + (1 - \xi(x))y(x) - d - \chi(d, a) - (1 + \tau_c)c + T(x) - \tau_b(x) \tag{1.1}$$

Drift in liquid assets comes from various sources. First, asset returns, where $r_b(b)$ is the balance–dependent rate of return on the liquid account. We assume positive balances attract a return of r_b, and that borrowing, whilst allowed, comes with an extortionate penalty $\omega >> 0$ so that $r_b(b < 0) = r_b + \omega$.[9] This assumption creates a soft–borrowing constraint, which will be important later. $y(x)$ is idiosyncratic income, assumed to be wz when working, w_U when unemployed, and w_R when retired. $T(x)$ captures fiscal transfers, c is consumption, which attracts a tax τ_c, and $\tau_b(x)$ is a state–contingent income tax function, discussed in full in Section 1.3.2.

Working households are required to contribute a proportion of their income $\xi(x)$ into their illiquid account (equal to zero in unemployed and retired states). As we will discuss in Section 1.3.2, this is one of the levers of regulation the government uses in retirement policy. As well as this, households may make voluntary transfers into (and out of) the illiquid account (d). These transfers are subject to a constraint $d \geq \gamma$ which is another lever of the government retirement policy, discussed in Section 1.3.2. Any voluntary transfers are subject to adjustment costs $\chi(d, a)$ which, following Kaplan et al. (2018),

[9] This setup reflects the empirical reality that very few households are actually borrowing constrained (Lee and Maxted, 2023), as in the wealthy hand–to–mouth literature (Kaplan and Violante, 2014). Instead, many hover close to zero liquid assets, rotate credit card balances (but not at their limit), and rarely exhaust *all* avenues for borrowing, which come with ever more onerous terms (pay–day loans, pawn shops, loan sharks etc).

have the structure

$$\chi(d,a) = \chi_0(d)|d| + \frac{\chi_1}{2}\frac{d^2}{a}$$

The convex nature of this function puts a handbrake on the voluntary transfer policy, necessary to avoid jumps in continuous time, and its structure leads to analytical solutions for d. The linear cost $\chi_0(d)$ may differ for withdrawals and deposits.[10]

Illiquid account The law of motion for the illiquid account is

$$\dot{a} = r_a a + \xi(x)wz + d - \tau_a(x) \tag{1.2}$$

Where r_a is the rate of return on the illiquid asset, and $\tau_a(x)$ is a state–contingent tax function, detailed in Section 1.3.2. Borrowing is not allowed in the illquid account, $a \geq 0$. As we will discuss in Section 1.3.2, and in contrast to the standard treatment in two–asset heterogeneous agent models (e.g. Kaplan et al., 2018), this illiquid account represents the households' regulated retirement accounts.

1.3.1.4 Present–biased Preferences

A portion of the population $\mu \in [0,1]$ is subject to present–bias, the rest are standard exponential discounters. The present–biased households have 'instantaneous gratification' (IG) preferences, the continuous–time analogue to quasi–hyperbolic discounting (Harris and Laibson, 2013; Laibson and Maxted, 2023; Maxted, 2024).

In working life, unbiased households' recursive preferences are as follows[11]

$$v(x_t) = \lim_{\Delta\to 0}\max_{c,d} u(c)\Delta + e^{-\rho\Delta}\,\mathbb{E}\left[v\left(x_{t+\Delta}(c,d)\right)\right] \tag{1.3}$$

Where the maximisation is constrained by the state transition functions that define $x_{t+\Delta}(c,d)$. $v(\cdot)$ is the value the household places on states x in time t, which comes from a combination of the utility they gain from optimal consumption for the present moment Δ, and the expected, discounted value placed on the state variables they're left with in the next moment ($\mathbb{E}[\cdot]$ captures all idiosyncratic risk transitions).

[10] This allows us to impose a withdrawal penalty on the illiquid account to replicate e.g. the USA's 10% tax penalty.

[11] This is derived by separating the present Δ from the future in the integral

$$v(x_t) = \max_{c_s,d_s}\mathbb{E}\int_t^\infty D(s-t)u(c_s)ds$$

where the discount function is $D(s-t) = e^{-\rho(s-t)}$

By contrast, the equivalent expression for present–biased households is[12]

$$v^{\beta}(x_t) = \lim_{\Delta \to 0} \max_{c,d} u(c)\Delta + \beta \cdot e^{-\rho\Delta}\, \mathbb{E}\left[v^E\left(x_{t+\Delta}(c,d)\right)\right] \tag{1.4}$$

The present–biased value in Equation (1.4) differs from (1.3) in two important ways. First, the continuation value is discounted by an extra $\beta \leq 1$ on top of the exponential discount $e^{-\rho\Delta}$. This is the source of present–biased behaviour—the IG agent values the future less than their exponential counterpart. Second, the continuation value $v^E(x_{t+\Delta})$ represents the value the household *believes* they will place on expected states in the future, which may not be how they actually value them. A 'sophisticated' agent holds correct beliefs: they will have the same value function in future as in the present (i.e. $v^E(x) = v^{\beta}(x)$) whereas 'naivete' wedges them apart.

We assume complete naivete i.e. $v^E(x) = v(x)$ where $v(x)$ is an exponential discounter's value function, defined in Equation (1.3).[13] This assumption simplifies the analysis because the solution can be reached in two steps: (1) solve the exponential discounter's problem to find $v(x)$ and use the solution as the IG consumer's continuation value to (2) solve the IG consumer's problem for each type.

1.3.1.5 Household problem and solution

During each stage of life, households choose consumption and voluntary transfers to maximise perceived value.

Exponential household Suppose we have substituted the drift in labour productivity with an N–state discrete process, and that this process, and jumps in and out of unemployment, are governed by Poisson intensities $\lambda^{z\to z'}$. During working life, the Hamilton–Jacobi–Bellman equation (HJB) is

$$\begin{aligned}\rho v(x) - \partial_t v(x) = \max_{c,d} &\left\{u(c) + \partial_b v(x)\cdot \dot{b}(x) + \partial_a v(x)\cdot \dot{a}(x)\right\} \\ &+ \sum_{z'} \lambda^{z\to z'}\left[v(x') - v(x)\right] + \delta_R\left[v_R(x) - v(x)\right]\end{aligned} \tag{1.5}$$

[12]As above, but where the discount is now the step function

$$D(s-t) = \begin{cases} 1 & \text{if } s-t=0 \\ \beta\cdot e^{-\rho(s-t)} & \text{if } s-t>0 \end{cases}$$

[13]This is an innocuous assumption. Maxted (2024) shows that under two assumptions—(1) CRRA utility and (2) soft–borrowing constraint—a problem with any degree of sophistication is isomorphic to a fully naive agent with a lesser degree of present–bias.

Where $\dot{b}(x)$ and $\dot{a}(x)$ are defined by Equations 1.1 and 1.2. This problem's FOC define the policy functions

$$\begin{aligned} u'(c(x)) &= (1+\tau)\partial_b v(x) \\ \partial_a v(x) &= \partial_b v(x)(1+\chi_d(d(x),a)) + \kappa(x) \end{aligned}$$

Where $\kappa(x)$ is the Lagrange multiplier on the withdrawal constraint.

During retirement, the equivalent HJB is

$$(\rho+\delta)v_R(x) - \partial_t v_R(x) = \max_{c,d} \left\{ u(c) + \partial_b v_R(x) \cdot \dot{b}(x) + \partial_a v(x) \cdot \dot{a}(x) \right\}$$

Where $\dot{b}(x)$ and $\dot{a}(x)$ are defined by Equations 1.1 and 1.2. This problem's FOC define the retired policy functions

$$\begin{aligned} u'(c_R(x)) &= (1+\tau)\partial_b v_R(x) \\ \partial_a v_R(x) &= \partial_b v_R(x)(1+\chi_d(d_R(x),a)) \end{aligned}$$

In either stage of life, the soft–borrowing constraint ensures the exponential household will never borrow and so these FOC always hold.

Present–biased household The biased households' perceived value and policies are recovered directly from the exponential household results.

Lemma 1.3.1 (Present–biased solution (Maxted, 2024))**.** Assuming (1) CRRA utility with risk–aversion σ, and (2) never–binding soft–borrowing constraint, the naive IG consumer's value and policies are scale transformations of the exponential discounter's equivalents

$$c^\beta(x) = \beta^{-\frac{1}{\sigma}} \cdot c(x) \quad \text{and} \quad v^\beta(x) = \beta \cdot v(x) \quad \text{and} \quad d^\beta(x) = d(x) \tag{1.6}$$

Proof. Derived in Appendix 1.A.1. □

The intuition behind this result, from Maxted (2024), is that an exponential discounter sets consumption so that marginal utility equals the marginal continuation value of liquid resources in future. The present–biased household does the same, but they perceive their marginal continuation value to be lower by β, and so consume more. Note that the present–bias only comes into play in decisions that trade between the present and future. The voluntary transfer choice is about balancing marginal values in the future, and as such is unaffected by present–bias.

These policies describe the optimal drift in the two accounts. Combined with the exogenous transitions in employment status, productivity, and life–cycle, they induce a stationary joint distribution $h(x)$ over all the households' state variables. We describe the Kolmogorov forward equation that characterises this distribution in Appendix 1.A.3.

Present–biased households over–consume relative to the exponential discounters they believe themselves to be by a factor of $\beta^{-\frac{1}{\sigma}} > 1$. As a result, left to their own devices, these households are left with less savings in retirement. Crucially, they regret ending up in this position: it is not the result of rational planning, but rather a series of mistakes that they would not have made if they could commit in advance to a state–contingent consumption plan. This regret leaves room for the government to intervene to help resolve this commitment problem with retirement policy.

State distribution The policy functions described above determine the drift in the endogenous state variables i.e. the balances in the liquid and retirement accounts. All other state transitions——between productivity and employment status within working life, and transitions between life stages——are exogenous. Together, these state–transition rules define how the distribution of households across the state-space moves around over time and therefore define the Kolmogorov Forward Equation (KFE). As detailed in Achdou et al. (2022), these movements are captured by the infinitesimal generator $\mathcal{A}$, the continuous–time equivalent of a discrete–time transition matrix, such that the KFE is

$$\partial_t h(x) = \mathcal{A}^* [h] (x) \tag{1.7}$$

Where $\mathcal{A}^*$ is the adjoint of the generator.

1.3.2 Retirement policy

The government sets retirement policy, which consists of the unconditional state pension w_R and regulations governing the households' illiquid accounts. In other macro papers, the illiquid account is usually taken to represent housing, or some other difficult-to-transact-but-attractive asset. In our model, it is an individual retirement account (IRA).[14] Our treatment of this account differs from the usual in some important ways.

First, the underlying asset is the same as in the liquid account, so the gross rates of return are equal.[15]

$$r_b = r_a = r$$

Second, the account's illiquidity stems from four regulatory parameters. The government

[14]Superannuation in Australia, IRAs and 401(k)s in the USA, SIPPs in the UK etc

[15]In developed countries at least, participation in DC plans doesn't change the span of assets available too much. Retail consumers may not be able to access alternatives like hedge funds themselves, but these make up a small portion of DC retirement funds, which tend to be mainly invested in market securities, or to hold real assets like infrastructure or commercial real estate that can also be accessed through market securities.

can penalise withdrawals during working life ($\chi_0(d < 0) = \chi_0$), or limit them directly by imposing the constraint $d \geq \gamma$. The source of the illiquidity in the retirement account is therefore regulatory; we set $\chi_0(d > 0) = 0$ and the convex adjustment cost χ_1 to be trivially low.

Third, in line with our discussion of DC schemes in Section 1.2, the government also mandates that a certain proportion $\xi \in [0, 1)$ of labour income be deposited into the retirement account. And , finally, the government offers a concession $\varphi \in [0, 1]$ so that contributions into and returns within the retirement account are subject to less tax. We discuss how the government government chooses these policy parameters in Section 1.5.

The mandatory contribution rate and tax concession affect the tax functions introduced in Equations 1.1 and 1.2 as follows:

$$\tau_b(x) = \tau\left[rb \cdot (b > 0) + (1 - \varphi\xi)wz\right]$$
$$\tau_a(x) = \tau(1 - \varphi)ra$$

The variable φ reduces the tax liability for mandatory contributions as well as asset returns within the account, implementing a concessional TTE system discussed in Section 1.2. If $\varphi = 1$, then contributions are made from pre–tax income, and returns are tax–free. We assume that retirees are free to withdraw from their retirement account, continue to receive tax concessions, but cannot deposit. At all stages in life, voluntary contributions are made from post–tax income, and withdrawals are not taxed.

Retirement policy gives life to the illiquid asset—it is illiquid due to regulations designed to discourage withdrawal during working life, and any difference in asset returns comes from preferential tax treatment. This setup closely mirrors the defined contribution systems discussed in Section 1.2. Retirement systems across the world are complex and vary from country to country (Beshears et al., 2015), but the setup in our model allows us to approximate many of their key features. In particular, we can span three of the first first four pillars in the World Bank's Conceptual Framework for classifying pension schemes (Holzmann et al., 2008): the state pension w_R protects households from poverty and affects redistribution (Pillar 0), mandatory contributions to the retirement account can be used to force personal saving in line with income, correcting for myopia and other errors (Pillar 2), and tax concessions in the retirement account can support voluntary savings as well (Pillar 3).[16]

Implementing a mandatory defined contribution system like the Australian Superannuation system, for example, is done by banning withdrawals and mandating contributions from labour income, as well as granting a tax discount of about 37.5% to an average

[16] What's missing is Pillar 1: a state–organised defined–benefit pillar like the USA's social security system, but this is beside the point for the purposes of this paper.

earner.[17] Under current legislation, the settings would be $(\chi_0, \gamma, \xi, \varphi) = (0, 0, 0.11, 0.375)$; and the average state pension per pensioner is around 20% average employed income.[18]

1.3.3 Fiscal authority

The fiscal authority pays benefits to the retired (w_R) and unemployed (w_U), services debt (rB), makes government purchases of the consumption good G, and makes discretionary transfers to households $T(x)$. It takes in taxes on consumption, capital and labour income. Deficits are funded by changes in government debt $\dot{B}$, defined below.

$$\dot{B} = \underbrace{G + w_U \pi_U + w_R \pi_R + rB + \int T(x) h(x) dx}_{\text{Spending}} \\ - \underbrace{\int \left(\tau_c \cdot c(x) + \tau_b(x) + \tau_a(x)\right) h(x) dx}_{\text{Tax revenues}} \tag{1.8}$$

Where π_U and π_R represent the measures of unemployed and retired people, and z_U and z_R represent those states (e.g. $\pi_U = \int h(b, a, z = z_U, \beta) dx$). Borrowing is restricted by an exponential fiscal rule (as in Angeletos et al., 2023; Auclert et al., 2023; Galí, 2020):

$$\dot{B} = -\mu \left(B - \bar{B}\right) \tag{1.9}$$

This rule sets an exponential decay of the government debt gap around some ideal level $\bar{B}$, where $\mu \geq 0$ (if zero, always hold debt constant; if ∞, return to $\bar{B}$ immediately).[19] We assume that government debt is fixed at the target in the stationary solution, and that the fiscal authority uses the consumption tax rate as the marginal tool to meet the fiscal rule in both the stationary and dynamic solutions. We also assume that this rule is suspended whilst any stimulus is active (whether fiscal or liquidity policy), and reactivated immediately after.

The fiscal authority has a balance sheet to manage and a rule to meet, but no objectives beyond this. Control of the discretionary policy levers and benefit levels are left to the government, whose choices the fiscal authority takes as given.

[17] Concessional contributions and returns within Superannuation are taxed at 15%, compared to a 23% income tax bill for a person on average income with no dependents.

[18] Australia's state pension is asset and income means tested so this average hides a lot of variation.

[19] Under this fiscal rule, debt must follow the path $B_{t+s} - \bar{B} = \left(B_t - \bar{B}\right) e^{-\mu s}$ yielding a half–life of $\ln 2/\mu$ in units of model frequency.

1.3.4 Stationary equilibrium

For fixed prices (w, r) and given retirement policy settings $(w_R, \chi_0, \gamma, \xi, \varphi)$, a stationary equilibrium is the set of working–life and retirement value functions $(v(x), v_R(x))$ and policy functions $(c(x), c_R(x), d(x), d_R(x))$, the measure over household states $h(x)$, and the consumption tax rate τ_c such that

- The values and policies solve the household problem (Equations 1.5 to 1.6)
- $h(x)$ is stationary (i.e. $\partial_t h(x) = 0$ in Equation 1.7)
- τ_c balances the budget (Equation 1.8)

In equilibrium, the household and fiscal authority constitute a combined model–block that takes prices and policy settings as given, and produces aggregate demand for a consumption good (the sum of household and government demand), net asset supply (total household wealth, net of government debt), and aggregate labour supply, which is an endowment process.

Stimulus policy This model allows us to analyse the two approaches to stimulus policy we're interested in. Stimulus is taken to mean policy interventions that increase aggregate consumption above its stationary level in the short–run. And the two methods available to the policy–maker are (a) fiscal transfers ΔT_t, funded by deficits and repaid via future tax changes that meet the fiscal rule, or (b) household liquidity policies i.e. regulatory changes to either restrictions on retirement account withdrawals $\Delta\gamma_t$, or mandatory retirement contribution rates $\Delta\xi_t$; we focus on the former in this paper.

1.3.5 Computational solution

We solve the model using a finite-–difference scheme (Achdou et al., 2022), with discrete, non–linear grids for the two endogenous assets, and three states for the idiosyncratic productivity process. The value–function updates are computed using the semi–implicit method (Achdou et al., 2022), and the policies are derived using the nested–-drift algorithm (detailed in Appendices 1.A.3 and 1.A.4), which is monotone and consistent, and is stable to parameter choices as a result (Sabet and Schneider, 2024). This is in contrast to existing finite-difference methods for solving continuous-time heterogeneous-agent models with two or more endogenous state variables, in particular the 'splitting the drift' variant of Achdou et al. (2022) used in Kaplan et al. (2018). The stationary solution makes further use of adaptive time–steps, detailed in Sabet and Schneider (2024), to ensure the starting guess can get appropriately close to the solution for local convergence to be guaranteed (Barles and Souganidis, 1991).

1.4 Calibration

We calibrate the model to match an advanced economy with a mandatory DC retirement system. To do this, most parameters are set to standard values in the literature. The remaining parameters govern retirement policy, which are set in the optimal policy exercise detailed in Section 1.5.

Our calibration achieves two goals that set the scene for our policy experiments. First, the retirement policy settings reflect and rationalise what we observe in countries with mandatory DC systems: mandatory contributions, no withdrawals allowed, and tax concessions. This is necessary to ensure a fair comparison between fiscal and liquidity policy. If retirement policy were too restrictive, liquidity policy will be both effective and unambiguously welfare improving; if too loose, then there isn't enough firepower for liquidity policy to be a viable substitute for fiscal. Second, the aggregate MPC among workers is empirically realistic. This is necessary to ensure the consumption boost from fiscal (the desired policy outcome) is matched by an appropriate increase in government debt (which sets the longer–run policy impact). This goal is achieved by estimating the present–biased share η that matches the model's aggregate MPC to empirical estimates, discussed below.

The calibrated parameters are outlined in Table 1.2, and detailed below[20].

Table 1.2: Calibrated parameters

Parameter	**Symbol**	**Baseline calibration**	**Source**
Retirement intensity	δ_R	1/160	40 year av. work–life
Death intensity	δ	1/80	OECD (2023)
Sep. & find. intensity	(λ_s, λ_f)	(0.0587,1.2)	Shimer (2005) & BLS
Log–income process	(ρ_z, σ_z^2)	(0.9136,0.0426)	Floden and Lindé (2001)
Risk aversion	σ	2	Standard
Discount rate	ρ	0.0025	Carroll et al. (2017)
Present–bias	(β_L, β_H)	(0.5,1)	Ganong and Noel (2020)
Present–bias share	η	0.5	Target MPC $\in [0.15, 0.25]$
Risk–free real rate	r	0.0051	2% p.a.
Borrowing penalty	ω	0.4024	500% p.a.
Wage	w	0.25	Numeraire
Unemployment benefit	w_U	0.1	Shimer (2005)
Income tax	τ	25%	OECD average
Consumption tax	τ_c	12%	Budget balance
Government spending	G	0.0238	G/GDP = 15%
Steady state debt	$\bar{B}$	0.1589	Debt/GDP = 25%
Fiscal rule	μ	0.0128	Galí (2020)
Adjustment costs	(χ_0, χ_1)	$(0, 0.001)$	Trivial

Life–cycle Working lives and retirement last an average of 40 and 20 years, respectively. This assumes a working life spanning 25–65, and matches the OECD average

[20]Time is continuous, with a base frequency of one quarter

expected life–years after retirement (OECD, 2023, p. 192). After death, retirees are replaced by employed workers with zero assets and productivity drawn from its stationary distribution.[21]

Working–life idiosyncratic risk Working households face idiosyncratic risk from jumps into, and out of, unemployment, and diffusion in their employed labour productivity. The jump transitions are governed by finding and separation rates of 0.0587 and 1.2 respectively; the former matches the quarterly separation rate in Shimer (2005), and the latter matches the mean 2.5 months spent in unemployment from the US Bureau of Labor Statistics.[22] The labour productivity process is calibrated to match the estimated AR(1) process in log–income residuals after individual characteristics effects are stripped out, from Floden and Lindé (2001).[23].

$$\ln z_t = 0.9136 \ln z_{t-1} + u_t,\ u_t \sim N(0, 0.0426)$$

We cast this in continuous time, following Achdou et al. (2022), so it becomes a monthly Ornstein–Uhlenbeck process, which we discretise over $k = 3$ points with reflecting barriers at on standard–deviation $\ln z \in [-\sqrt{0.0426}, \sqrt{0.0426}]$, and normalise so the stationary distribution of z, in levels, has a mean of one.[24]

[21]We also impose two extra rules on the transition from working life to retirement that households do not anticipate. The first is a 'forced–retirement' level in working households' illiquid account which does what it says on the tin. This limit is necessary to ensure the state–space is compact and that retirement balances don't get out of hand, but is set to $a_{max} = 15$, sufficiently high that it only affects a small measure of workers and keeps the retired population at realistic levels. And second, households that retire with negative total assets (liquid plus illiquid) are bankrupted. In the model, this means their gross positions in both accounts are returned to zero (the same as if they were born into retirement). Without bankruptcy, a small measure of households retire with the maximum debt, and they are stuck there because it is an absorbing state for present–biased households. This is a disastrous position to be in for these households—their consumption is close to zero and so their value is orders of magnitude lower than at other points in the state space. If the risk of destitution is not addressed, then the government's retirement policy is primarily focused on managing it, rather than the more prosaic concern of general retirement adequacy. The bankruptcy rule avoids the issue.

[22]Table A.12 'Unemployed persons by duration of unemployment: Monthly, Seasonally Adjusted'; recent average outside of recessions.

[23]This calibration is used in Maxted et al. (2024) and Guerrieri and Lorenzoni (2017). The estimates are from PSID data covering 1988–1991. More recent estimates of the same AR(1) process all get numbers around this i.e. with auto–regressive parameter at least 0.9, and standard–deviation at least 0.2 (e.g. Chang et al., 2013; Guvenen et al., 2023; Kaplan et al., 2020a).

[24]This yields a stationary distribution defined by

$$\{z_L, z_M, z_H\} = \{0.681, 0.9516, 1.4652\} \text{ and } \{\pi_L, \pi_M, \pi_H\} = \{0.2686, 0.4628, 0.2686\}$$

Where π_i represent the stationary probability of being in state i.

Preferences Households have CRRA preferences over consumption

$$u(c) = \frac{c^{1-\sigma} - 1}{1 - \sigma}$$

With a standard coefficient of relative risk aversion equal to $\sigma = 2$. Households discount the future at $\rho = 0.0025$, which corresponds to an annual discount rate of 0.99. Recall that households also face the risk of retirement and death. This boosts the effective annual discount rate to 0.97 and 0.94 for employed and retired households, within the range of standard estimates (Carroll et al., 2017).

Prices & fiscal The base rate of return is $r = 0.0051$ per quarter (2% p.a.) and the borrowing penalty is $\omega = 0.4024$ (500% p.a.), set to be extortionate to impose a soft–borrowing constraint. The annual wage for a unit of effective labour is the numeraire (quarterly wage $w = 0.25$) so all other monetary values are relative to this. The government pays an unemployed benefit $w_u = 0.1$ to match the standard replacement rate of 0.4 from Shimer (2005). The baseline income tax is set to $\tau = 25\%$, the OECD average personal income tax rate for a single person with no children on the average wage. In both the stationary solution and the dynamic exercises later, the consumption tax τ_c is set internally to meet the fiscal rule.

Steady state government spending is $G = 0.0238$, targeting 15% of GDP, and steady state debt levels are $\bar{B} = 0.1589$, or 25% annual GDP. In both cases, GDP here is taken to be the aggregate income of the steady state employed population multiplied by 3/2 to adjust for the capital share. The exponential parameter on the fiscal rule is set to $\mu = 0.0128$, such that debt gaps are closed with a half–life of 13.5 years, following Galí (2020).[25] This parameter has no influence on the stationary solution, in which the rule dictates budget balance.

Adjustment costs The adjustment cost function serves no purpose other than to deliver analytical solutions for the voluntary contribution policy d.[26] The baseline parameters that set common adjustment costs between working and retired households are set so the costs are trivial $(\chi_0, \chi_1) = (0, 0.001)$.

Present–bias The population is split into present–biased and exponential households. The present–biased have $\beta = 0.5$, similar to Ganong and Noel (2020) and Laibson et al. (2024), and the exponential have $\beta = 1$. The size of the present–biased population is set to $\eta = 0.5$. This leads to an average present–bias in the population of 0.75, which

[25] This matches the European Union's fiscal compact, which includes a provision that excess debt should be reduced by $1/20^{th}$ each year (Galí, 2020).

[26] It is simple to use this function to impose an early withdrawal penalty during working life, as in the USA.

is within the range of other estimates (Hamilton et al., 2024; Laibson et al., 2024), but a greater biased proportion than estimated by Ganong and Noel (2020), who find only 25% of their sample (unemployed workers) exhibit this degree of bias. This parameter has crucial importance for the model because it sets the aggregate marginal propensity to consume. As we show in Section 1.5.6, setting $\eta = 0.5$ leads to an equilibrium MPC in the centre of the range of empirical estimates, and we explore the sensitivity of our results to this parameter in Section 1.5.5.

Retirement policy framework We restrict attention to mandatory DC schemes backed by a fixed state pension because this is the relevant policy framework within which liquidity policy has been used. Withdrawals are not allowed $\gamma = 0$ and the withdrawal penalty is set to zero $\chi_0(d < 0) = 0$. The state pension is set to $w_R = 0.075$, a replacement rate of 30% average worker income, which matches the average replacement rate across OECD countries with a mix of state and private pension schemes.[27] The remaining retirement policy parameters are the mandatory contribution rate ξ and the tax concession φ. We saw in Section 1.2 that these parameters vary across countries but that contributions cluster in the range $8 - 20\%$ and tax concessions are universally granted to these illiquid savings environments. In the next section, we set these remaining parameters optimally.

1.5 Optimal retirement policy

In this section the retirement policy parameters are set to implement the optimal mandatory DC scheme. Within the constraints of the policy framework, the government chooses the optimal mandatory contribution rate (ξ) and tax concession (φ).

1.5.1 The government's problem

The government's problem is similar to that in the literature on paternalistic savings policies in that they must balance two tradeoffs. First, following Amador et al. (2006), there is both a need to provide households with commitment (to overcome present bias) and flexibility (to insure against idiosyncratic risk). This prompts government intervention to help the present–biased save, but puts a limit on how much mandatory saving is appropriate. Second, present–bias is heterogeneous but unobserved by the government. This introduces a need for the government to balance the interests of the biased against

[27]This is the replacement rate of average income provided by the state pension ('Mandatory Public'), weighted average across OECD states that combine state and private components in their pension schemes (OECD, 2023, Table 4.2, p. 153). Where the state is the sole pension provider, e.g. Austria, Spain, or Colombia, this weighted average is 56% in OECD countries.

the unbiased, potentially prompting screening or compensation (Beshears et al., 2025; Moser and Olea de Souza e Silva, 2019). Our setting differs from the extant literature by limiting the government's options to a DC system with uniform settings across the working–age population.

We evaluate welfare of these policy settings as follows.

Definition 1.5.1 (Social welfare criterion)**.** Social welfare in the stationary solution is defined as the expected long–run value for a prospective newborn.

$$W(\xi, \varphi) = \mathbb{E}_{(z,\beta)}[\hat{v}(0, 0, z, \beta; \xi, \varphi)] \tag{1.10}$$

$$\hat{v}(x; \xi, \varphi) = \mathbb{E} \int e^{-\rho s} u\left(c(x_s; \xi, \varphi)\right) ds \tag{1.11}$$

This is found in two steps. First we find the value of actual (rather than anticipated) behaviour[28] in Equation 1.11, referred to as the 'long–run value' (Bernheim and Taubinsky, 2018; Naik and Reck, 2024; O'Donoghue and Rabin, 2006). Using this value means the government anticipates but does not adopt the bias of its subjects when evaluating policy choices. Second we restrict attention to newborns (zero assets) and find their expected long–run value, based on the stationary distribution over (z, β) in Equation 1.10. Defining the criterion like this means comparisons of welfare under different policy regimes pose the question: 'which society would you prefer to be born into, anticipating potential present–bias?'.

Having defined a welfare criterion, we consider two approaches to the problem that differ in how powerful the government is. These lead to two distinct equilibria, defined below.[29]

Definition 1.5.2 (Social equilibrium)**.** The social equilibrium is a stationary equilibrium with retirement policy settings $(\tilde{\xi}, \tilde{\varphi})$ that maximise steady state social welfare.

$$(\tilde{\xi}, \tilde{\varphi}) = \arg\max_{\xi, \varphi} W(\xi, \varphi)$$

The social equilibrium is the outcome of an all–powerful government's problem, in the sense that they can guarantee compliance with the policy settings they choose. As we will show in the next section, the optimal tax concession is zero in the social equilibrium. To rationalise the ubiquity of these tax concessions, we suppose the government is subject to an extra participation constraint, optimisation under which leads to the 'buy–in' equilibrium.

[28]Both $w(x)$ and $v(x)$ are based on the incorrect assumption that the policy rules are $\check{c}(x)$.

[29]We don't mean that there are multiple equilibria, but that the government will select different equilibria under these different constraints.

Definition 1.5.3 (Buy–in equilibrium)**.** The buy–in equilibrium is a stationary equilibrium with retirement policy settings (ξ^*, φ^*) that maximise social welfare subject to a participation constraint: no newborn chooses to opt out of the system by setting their own personal retirement parameters to $(\xi_i, \varphi_i) = (0, 0)$.

$$(\xi^*, \varphi^*) = \arg\max_{\xi,\varphi} W(\xi, \varphi) \;\; \text{s.t.} \;\; \hat{v}(0, 0, z, \beta; \xi^*, \varphi^*, \tau_c) \geq \hat{v}(0, 0, z, \beta; 0, 0, \tau_c) \; \forall \, z, \beta$$

The buy–in equilibrium is the outcome when we imagine each household has the option at the start of their careers to opt out of the system[30]. In this case the government must cajole compliance for the system to be stable. The government wants to implement a pooling equilibrium where no–one opts out because (i) their policy tools don't allow for screening as in Moser and Olea de Souza e Silva (2019), and (ii) a separating equilibrium where only the present–biased opt in is not possible anyway, due to naivete. Note that this amounts to ensuring the rational agents opt in: assuming naivete, everyone thinks themselves rational.

1.5.2 The social equilibrium

Figure 1.4 plots the social welfare under different combinations of retirement policy settings, with warmer colours representing greater welfare. The social equilibrium is labelled, and picks the highest welfare point in the area plotted.

Result 1.5.4 (Social equilibrium)**.** The social equilibrium is $(\tilde{\xi}, \tilde{\varphi}) = (0.09, 0)$

This contribution rate is toward the lower end of the mandatory contribution rates in countries with compulsory DC schemes, with OECD examples outlined in Table 1.1.

Note that the social equilibrium involves no tax concession. The plotted results are for a population with $\eta = 0.5$, but the optimality of zero tax concession in social equilibrium holds no matter how small or large a share of the population are biased—it is never necessary to compensate people for a policy that helps them, and tensions between biased and unbiased workers are better resolved by adjusting the contribution rate than with tax concessions.

This result reflects a finding in the empirical public literature that a majority of people are not responsive to incentives to save for retirement, and those that are tend to save more in retirement accounts by reducing savings elsewhere (Chetty et al., 2014; Choukhmane and Palmer, 2024). The use of tax concessions to encourage retirement savings is therefore likely to incur a fiscal cost with little welfare benefit. In our setting, greater tax concessions also necessitate higher tax rates on consumption to make up for the eroded base. The result here tells us the marginal welfare benefit from the concessions is less than the welfare cost of their fiscal side–effects.

[30] Call this the 'Dubai option'.

Figure 1.4: Social welfare under policy settings

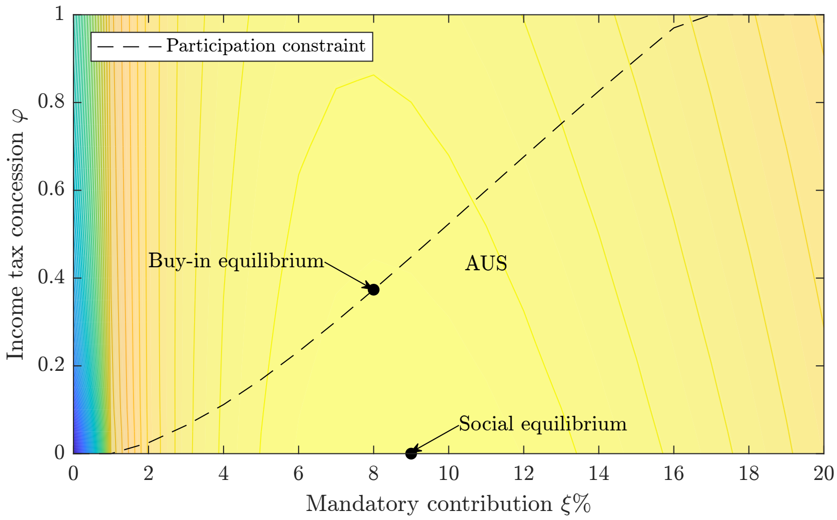

1.5.3 The buy–in equilibrium

Despite their apparent suboptimality, tax concessions are ubiquitous in actually–existing DC schemes, as discussed in Section 1.2. To rationalise them, we explore a political dimension to the government's problem—the need to implement a pooling equilibrium where agents opt–in to the retirement system at the beginning of their careers. The black dashed line in Figure 1.4 traces out the limited menu facing a government that has to deal with a participation constraint. In this setting, the tax concession ensures participation. Each contribution rate ξ has a minimum concession $\varphi(\xi)$ necessary to implement a pooling equilibrium, and this minimum is increasing in the contribution rate $\varphi'(\xi) > 0$. The government selects among these $(\xi, \varphi(\xi))$ to maximise welfare. The result is the buy–in equilibrium, with an optimal contribution rate lower than in the social equilibrium, and more tax concessions.

Result 1.5.5 (Buy–in equilibrium)**.** The Buy–in equilibrium is $(\xi^*, \varphi^*) = (0.08, 0.37)$

This result is in line with what's observed in the real world—substantial tax concessions coupled with mandatory contributions close to the OECD examples in Table 1.1—and this is our preferred calibration as a result. The tax treatment in the model is a TTE system, with concessions for the taxes on entry and returns. Direct comparison to other countries is difficult because most countries use an EET system. The one country that is comparable is Australia, identified on Figure 1.4, which taxes mandatory contributions from pre–tax income and Superannuation returns at 15%, a $\varphi = 0.4$ discount on the standard 25% rate for an average worker.[31] With a mandatory contribution rate currently at 11%, this is remarkably close to the buy–in equilibrium in our model.

[31] Specifically this is the average personal income tax rate on labour income for a single person with no children on the average wage, from the OECD's 'Labour taxation - average and marginal tax wedge decompositions' series in 2023.

1.5.4 The impact of retirement policy for working–age people

Retirement policy substantially raises working households' saving rates and retirement adequacy, and this effect is stronger for the present–biased households. Table 1.3 shows the impact of retirement policy for workers, by type. To see behaviour without mandatory savings first we set $\xi = 0$ and re–solve the model. In this scenario, working households' saving rate is 10% of earnings, with a large difference between present–biased and unbiased households (4% and 16% respectively). With mandatory savings, the average saving rate is higher (23%), with little difference between the two types of household.

Table 1.3: Effect of mandatory saving on workers, by type

	Average saving rate (% earnings)		**Median $\%\Delta C$ at retirement**	
	Present–biased	Unbiased	Present–biased	Unbiased
Without mandatory saving	4	16	−27	−13
With mandatory saving	22	24	16	−12

Similarly, retirement adequacy, which we measure as the median expected change in consumption from working to the first year of retirement, is much improved by the policy (see Table 1.3).[32] Without mandatory saving, unbiased people's consumption drops by a median of 13% upon retirement, and more than double this for present–biased households. With mandatory saving, the unbiased household's consumption drop barely changes, whereas biased households now see an increase in their consumption of 16%, as they move from being working–poor to having resources to spend.

1.5.5 Sensitivity to present–biased share

Table 1.4 shows how the government's solutions change with the present–biased population share.[33] The optimal mandatory contribution rate is increasing in the present–biased population share, reflecting the increasing level of need in society. The tax concession necessary in the buy–in equilibrium is increasing as well, to ensure early–career opt–in to the system.[34] The aggregate worker MPC in the buy–in equilibrium is also increasing in η, both due to the combination of a greater share of present–biased households, and the greater mandatory contribution rate.

[32] Specifically, we find the difference between the workers' consumption policy and the average expected consumption rate over the first year of retirement, using the employed worker's state distribution to identify moments.

[33] The grid–search is restricted to whole percentage–points in ξ.

[34] The optimal tax concession is zero in the social equilibrium across all levels of the present–biased population.

Table 1.4: government solutions with different population mixes

Present–biased share η	10%	20%	30%	40%	**50%**	60%	70%	80%	90%
Buy–in Equilibrium									
ξ^*	0.04	0.06	0.06	0.07	**0.08**	0.08	0.08	0.09	0.09
φ^*	0.11	0.23	0.23	0.30	**0.37**	0.37	0.37	0.45	0.45
Worker MPC	0.05	0.10	0.14	0.18	**0.22**	0.26	0.30	0.34	0.38
Social Equilibrium									
$\tilde{\xi}$	0.05	0.06	0.07	0.08	**0.09**	0.09	0.09	0.10	0.10

1.5.6 Model validation

Table 1.5 compares aggregate moments among workers in the stationary solution to the model to equivalents in the USA. The model produces an average quarterly MPC[35] of 0.22 among workers (0.24 for the whole population), within the range of recent empirical estimates for non–durable consumption (Fagereng et al., 2021; Ganong et al., 2023; Jappelli and Pistaferri, 2014; Kaplan and Violante, 2022; Kueng, 2018; Sahm et al., 2010, 2012). The MPC is this high because 40% of workers are classed as hand–to–mouth, in line with standard estimates from the literature (e.g. Aguiar et al., 2024; Kaplan et al., 2020a; Kaplan and Violante, 2022).[36]

Table 1.5: Aggregate moments for workers

Moment	**Model**	**Data**	**Source**
Quarterly MPC	0.22	$[0.15, 0.25]$	Kaplan and Violante (2022)
% HTM	40	41	”
Liquid wealth / labour income	1.0	0.6*	”
Fin. wealth / labour income	6.4	4.1*	”
Personal saving rate	7%	$[0, 10]$%	OECD range
Median Δc on retirement	-1%	-3.5%	Aguila et al. (2011)

*From the bottom 95% of the empirical wealth distribution.

The model comes close to aggregate wealth–to–income ratios: recording total financial wealth of 6.4 times labour income, compared to 4.1 in the USA data; and liquid wealth measuring one year's worth of labour income, compared to 0.6 in the data. Having greater total wealth is natural, given our assumption of a mandatory DC retirement system. The actual USA features a voluntary DC system, combined with government–backed Social Security that does not count in wealth statistics. The average household saving rate is 7%, higher than in the USA, but in the range of OECD rates.

The median drop in consumption upon retirement is -1%, which is well within estimates from the existing literature; for example Aguila et al. (2011) estimate a median change of -3.5% with substantial variation around this. The size and sign of this change has been the subject of a lot of empirical work. Initial estimates, focused on

[35]See Appendix 1.A.6 for the formula used to calculate this object in the model.

[36]Hand–to–mouth status is defined by having liquid assets less than half of monthly labour income (Kaplan and Violante, 2022).

food, found a substantial drop in consumption on retirement (Aguiar and Hurst, 2005; Banks et al., 1998; Bernheim et al., 2001). Later work explained this as substitution to home–production, meaning that such a measured drop doesn't reflect a loss in welfare (Aguiar and Hurst, 2005, 2013). Later work, expanding the definition to include other categories of spending, showed that there is no average drop among individuals who retire voluntarily, but that there is a lot of heterogeneity around this average across households. The change in spending at retirement is driven by both wealth (Aguila et al., 2011) and unobservable characteristics (Moran et al., 2021). Our median result matches the literature in that it is close to zero, and the average change is increasing in total wealth as well.

Across the distribution Figure 1.5a plots the distributions of liquid and retirement assets. Few households have any liquid assets at all, a result of their present–bias leading them to under–save. The distribution of retirement assets is approximately exponential, because it reflects the age distribution.[37]

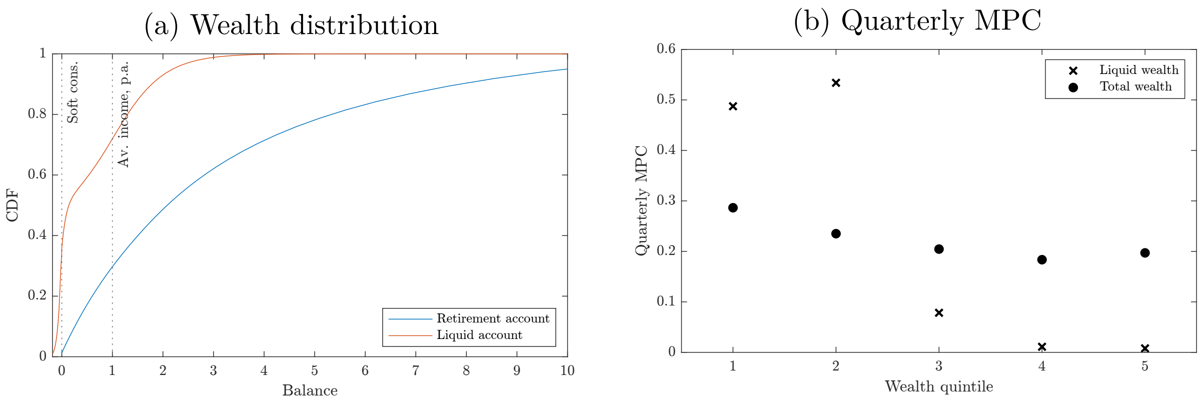

Figure 1.5b plots the household MPC across the quintiles of liquid and total wealth. The MPC is high in the first two quintiles of the liquid account distribution, which are closest to the soft–borrowing constraint, and declining rapidly after that. The MPC is roughly flat with total wealth. This mirrors reality well, as various studies have shown that the MPC is sharply declining in liquid wealth, but much less so with total wealth (e.g. Fagereng et al., 2021; Ganong et al., 2023).

This model is well situated for the stimulus policy experiments in the next section. The MPC matches empirical estimates, and so we can be confident that the sizing of fiscal policy will be appropriate. Similarly, because retirement policy is set optimally, the cost to society of liquidity policy is appropriate. One implication of the lack of voluntary contributions is that everyone will take advantage of opportunities to withdraw. Implemented like this liquidity policy will be identical to a fiscal transfer funded by a

[37] In the baseline calibration no–one makes voluntary contributions, and so balances in the retirement account are built of 8% contributions from labour income, and asset returns, accumulated over time.

simultaneous lump sum tax on retirement accounts. We will discuss this further in the next section.

1.6 Stimulus policy

In the remaining sections of the paper we explore the effects of household liquidity policy in comparison to conventional fiscal stimulus. We first establish that liquidity policy indeed works similarly to fiscal stimulus in aggregate: opening a window during which people can withdraw from their retirement accounts leads to such withdrawals, and the greater liquidity boosts consumption in much the same way as a fiscal transfer. From the perspective of the economy, the two approaches are similar—they increase liquidity in the present by incurring debt, but they differ in where this debt sits (on the government's books, representing an implicit liability for households in the form of future taxes, or against workers' retirement accounts), and the process through which it is repaid.

Having established that both approaches 'work' as stimulus in aggregate, we then explore the other implications that each stimulus approach has. These include distortions to inter–temporal choices from changing tax rates, inter–generational distributional effects, and reduced retirement adequacy for workers. To weigh these distinct implications appropriately, we use a welfare analysis to quantify the degree of households' preference for liquidity policy over fiscal stimulus. We show in a welfare exercise that liquidity policy is better for the retired, richer and unbiased workers, and future generations i.e. those who will (a) not gain much from the stimulus benefits of either policy, but will (b) pay more in taxes under fiscal, and (c) are more likely to be over–invested in their retirement accounts. We conclude by exploring the sensitivity of the welfare analysis to various dimensions of the problem—the strictness of retirement policy, the nature of the fiscal rule, and the size and targeting of the fiscal stimulus.

1.6.1 Defining the stimulus policy experiments

We assume that the planner wishes to induce a specific stimulus to aggregate household consumption over a set period of time, and it is deciding which intervention to use. In the baseline experiment we set the desired stimulus to be 5% of stationary average household consumption over a period of one quarter (i.e. $\Delta = 1$). To frame the results it is useful to define a policy–dependent aggregate that accumulates average household consumption over a period Δ

$$C(T, \gamma) = \int_0^{\Delta} \left(\int c_{t+s}(x; T, \gamma) \cdot h_{t+s}(x; T, \gamma) dx \right) ds$$

The stationary aggregate, for example, is $\bar{C} = C(0,0)$; and the target for a given policy mix (T, γ) is therefore

$$C(T, \gamma) = 1.05 \times \bar{C}$$

Both interventions are modeled as MIT shocks to the policy instrument that last for a duration of one–quarter.[38] The baseline fiscal intervention is a shock to T, paid to workers, that solves: $C(T^*, 0) = 1.05 \times \bar{C}$.[39] And the baseline liquidity policy intervention is similarly a shock to γ that lasts for one quarter such that $C(0, \gamma^*) = 1.05 \times \bar{C}$. In both cases, taxes adjust endogenously to meet the fiscal rule defined in Section 1.3.3. All other variables, like wages and the interest rate, are held fixed to isolate the direct effects of the policies.

1.6.2 Common aggregate stimulus

The first thing to establish is that liquidity policy works as stimulus in the model. The amounts needed to achieve a 5% consumption boost are below.

Result 1.6.1 (Stimulus equivalence)**.** The baseline interventions are close to the same magnitude at $T^* = 0.0706$ and $\gamma^* = -0.0723$.

The calibrated stimulus policies transfer the equivalent of 7.06% average annual income to working households under fiscal policy, or allow them to withdraw an amount equivalent to 7.23% of average annual income over a quarter under liquidity policy. In the US, this is equivalent to nearly \$5,500 in each case. The calibrated policy counterfactual is therefore large for a fiscal stimulus (e.g. the US CARES Act transferred \$1,200 per person and an extra \$500 per child), but small for a liquidity policy (e.g. in Australia people were allowed to withdraw up to \$AU20,000 from their superannuation accounts). The fact that the calibrated numbers are so close suggests policymakers should see them as equivalents, at least in terms of their aggregate impact on aggregate demand.[40]

[38] Unanticipated before time–0 but known thereafter.

[39] We choose to target workers in the baseline fiscal exercise so its magnitude can be compared to the alternative liquidity policy, which is only effective for the working–age. We explore an alternative where the whole population receives the transfer in Section 1.7.2.

[40] This aggregate equivalence would be ameliorated by a participation decision, which we do not model here. We explore the participation channel in Schneider and Moran (2024b), showing that 1/6 people withdrew when allowed in the Australian Covid–19 program, and we establishing that self–control issues were an important driver of this decision, alongside more standard drivers of demand for liquidity like assets and income shocks. Furthermore Hamilton et al. (2024) show that, conditional on participation, the MPC was 40%—higher than aggregate estimates. The high MPC can be explained by a combination of present–bias, and the self–selection of the most needy. A back of the envelope exercise then suggests an aggregate MPC out of liquidity policy of $0.4 \times 1/6 = 7\%$, relative to the 20% or so out of fiscal policy, meaning liquidity policy is only half as potent as fiscal stimulus, at least in the Australian context.

Figure 1.6 shows the impulse responses of aggregate consumption under the two different stimulus policies. Liquidity policy clearly has a stimulative effect on consumption, just as fiscal does, prompting a 5% increase in the aggregate consumption rate over the stationary equivalent ($\bar{C}$), as they were calibrated to do. After the immediate stimulus, liquidity policy produce less drag on consumption in the medium term (left panel). The two paths for consumption diverge within a few years, with the one under fiscal policy going negative as the stimulative effects wear off and the repayment plan kicks in. Under both policies this stimulus is mainly driven by workers who started off with less liquidity (right panel), as expected because these people tend to have higher MPCs (see Figure 1.5b).

Figure 1.6: Consumption response to different policies

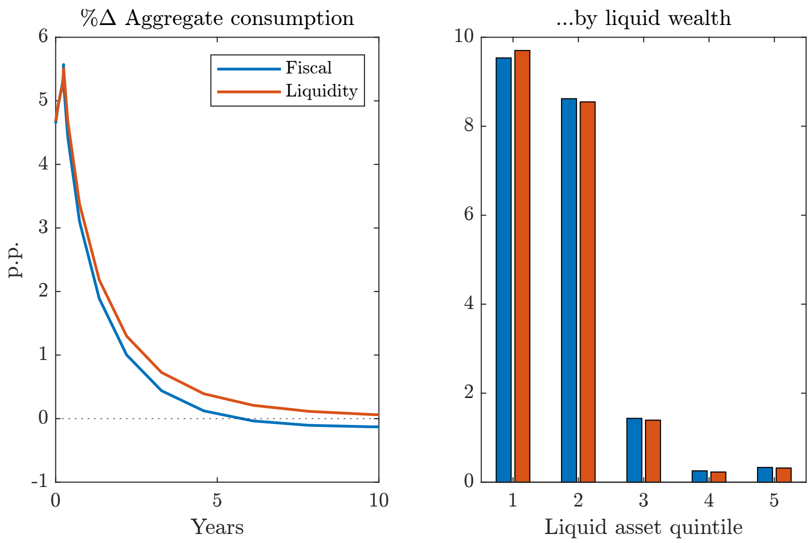

1.6.3 Difference in funding

The difference in funding mechanism is apparent by comparing the responses of government debt and household wealth under the two policies. In Figure 1.7, government debt jumps under fiscal policy as the transfers drive deficits, and then glides down as the debt is repaid with higher taxes. This increased government debt funds a concurrent increase in household wealth (visible in the middle panel) under fiscal policy, which reduces as households spend the extra liquidity on consumption and higher taxes. By contrast, there is no immediate impact on household wealth from liquidity policy (which only alters asset allocation), but a similar gradual decline as households (a) consume the extra liquid resources, and (b) retire with fewer illiquid resources.

The government debt accumulated under fiscal policy represents an implicit liability for households: the present value of their increased future tax bill, not recognised in their balance sheets. To show the consolidated household financial position, the right panel of

Figure 1.7: Funding channels

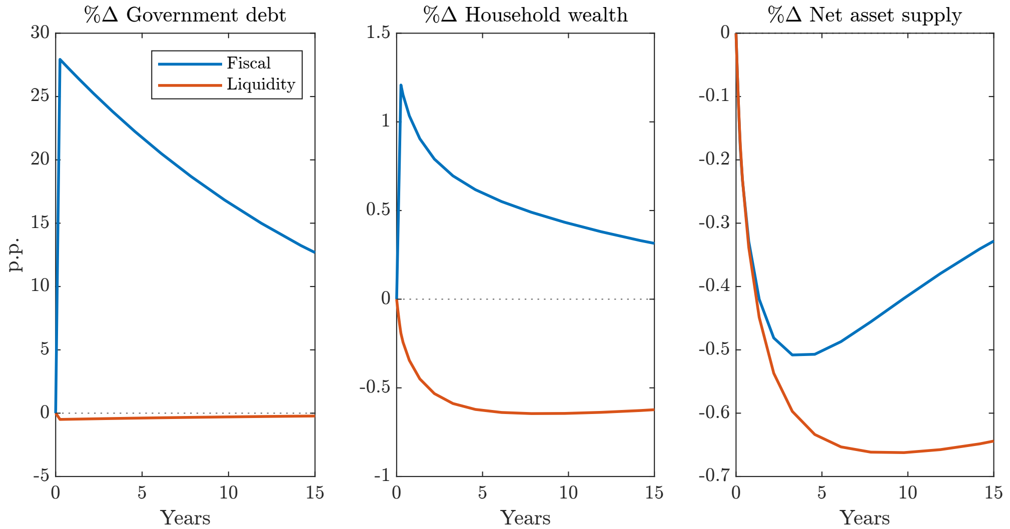

Figure 1.7 plots their net asset supply—the difference between households' wealth (the sum of liquid and illiquid assets) and government debt. This plot shows two things. First, the path of this aggregate is quite similar under the two policies for the first few years, and only starts to deviate in later years. The effects of both policies last a long time. Second, the duration of fiscal policy's impact clearly depends on the fiscal rule, but liquidity policy's effects necessarily span generations. They only wash out of the system after all affected workers retire, and die, which will be many decades after the stimulus occurred.

From the perspective of the aggregate economy, both policies do the same thing in qualitative terms. They increase liquidity for workers, driving short–term consumption, and they fund this with decreased illiquid wealth, whether in retirement accounts or greater future taxes. The difference between the two is where decreased illiquid wealth is situated, who repays it, and when. Under fiscal policy it sits as a debt on the government's balance sheet, and it is repaid with the (broad–based) consumption tax in a process determined by the fiscal rule. Under liquidity policy it sits like a debt on households' balance sheets (as a negative entry in their retirement accounts), and it is 'repaid' in lump sums by workers as they retire, in a process determined by the retirement rate. The primary reason the aggregate asset supply curves diverge in the right panel of Figure 1.7 is that the fiscal rule 'retires' debt at a faster rate than workers retire.[41]

Although the two approaches are qualitatively similar at the aggregate level, they have potentially different distributional and normative implications because of the difference in the size of the repayment tax base, and the redistribution and distortions caused by the repayment mechanism. The taxes used to repay increased government debt have a broad

[41] We explore an alternative fiscal rule that equalises these rates in Section 1.7.2.

base, applying to workers and retirees regardless of their status when the fiscal policy was implemented. By contrast liquidity policy is repaid from a narrower base—only the workers that extracted liquidity are forced to 'repay' by having lower future–value retirement savings—and it is not at all redistributive. One other difference between these repayment mechanisms at the aggregate level is whether they cause inter–temporal distortions.

1.6.4 Inter–temporal distortion

The two approaches have different implications for the path of tax rates, and the inter–temporal distortions these cause. Knowing the tax rate is growing (declining), households will shift consumption into the present (future)—a distortion away from the optimal smoothing path.

Lemma 1.6.2 (Euler distortions)**.** Assuming the borrowing constraint is non–binding, the changing consumption tax rate distorts the Euler equation

$$\mathbb{E}\left[\frac{\dot{c}(x)}{c(x)}\right] = \frac{1}{\sigma}[r_b(b)b + r(b)) - \rho] - \underbrace{(1+\tau_t)(1-\beta^{1/\sigma})\partial_b \hat{c}_t(x)}_{\text{Bias distortion}} - \underbrace{\frac{1}{\sigma}\left[\frac{\dot{\tau}}{1+\tau_t}\right]}_{\text{Tax distortion}}$$

Proof. Derived in Appendix 1.A.2.2 □

Figure 1.8 plots the consumption tax over time under the two policies (left panel), as well as the distortions to the optimal expected consumption growth path they imply (right panel). Under fiscal policy, consumption taxes climb steadily, peaking 0.4 p.p. higher than the stationary level about five years after the stimulus, and they glide down slowly over following decades. By contrast under liquidity policy the consumption tax drops with the stimulus. This is because the stimulus increases the tax base and the fiscal authority needs to keep the budget balanced.[42] It climbs back again slowly over following decades as aggregate consumption returns to normal.

Both policies imply short–term distortions to consumption growth because they both imply climbing tax rates over the first few years. These distortions effectively disappear under liquidity policy after the first few years, as the consumption tax rate stabilises close to the stationary level. But the distortions persist under fiscal policy for decades, with a maximal impact of causing expected annual consumption growth to be more than 4 percentage points higher than in the stationary solution.

We have established that both policies stimulate consumption in much the same way, but they differ in their funding mechanisms, and the inter–temporal distortions these imply. These results are all at the aggregate level, but the different approaches to stimulus

[42] We explore an alternative scenario with an asymmetric fiscal rule in Section 1.7.2

Figure 1.8: Euler distortions

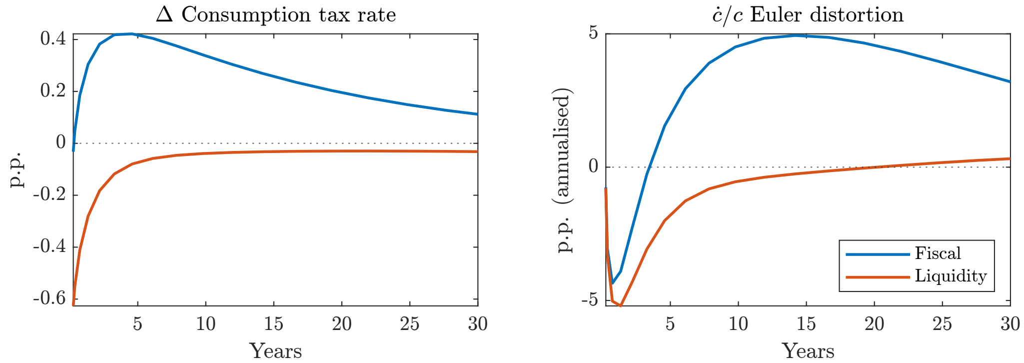

distribute their impacts quite differently across households states, and we explore this in the next section.

1.6.5 Distributional implications

1.6.5.1 Conflicting generational interests

Questions about the long–run implications of government policies are often framed as generational conflicts. It's useful to define three distinct groups within the population because they are affected by the policies differently. These are workers and retirees, those who are working whilst the stimulus occurs, and retired at time–0, respectively, and the future generations, those whose working lives begin after the stimulus ends.

Figure 1.9: Generational consumption

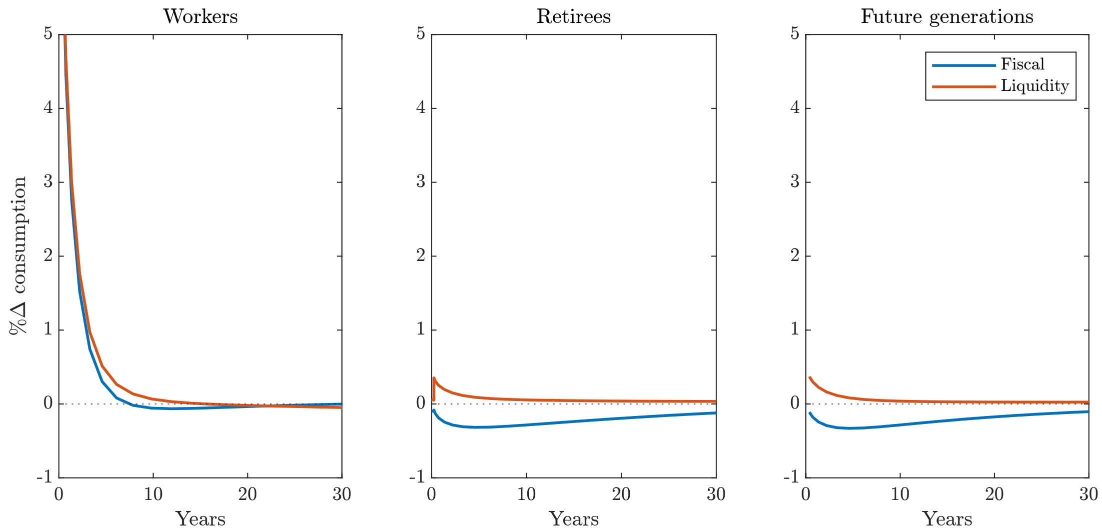

Figure 1.9 shows the average consumption path, relative to the stationary solution, for these distinct groups under each stimulus policy. Workers' consumption is boosted the most, by design, and there is little drag in the later years as the repayment plans for each policy kicks in. By contrast, retirees and future generations both experience a

drag on their consumption under fiscal policy—they must pay higher taxes and didn't receive any stimulus—and a boost under liquidity policy—caused by the tax distortions discussed in Section 1.6.4.

1.6.5.2 Retirement (in)adequacy

Reduced retirement adequacy—fewer resources on retirement for affected households—is the main negative implication of liquidity policy, and it only affects the workers who withdrew from their accounts and spent the money. To measure retirement adequacy, we compare the consumption rate that the cohort of workers at time–0 can expect the moment they retire under the different stimulus policies, relative to the equivalent in the stationary solution.

Figure 1.10 plots this change for the section of this cohort that retires 20 years after the stimulus, and breaks down contributions by wealth quintile and present–bias types.[43] The figure shows that liquidity policy has a much more severe impact on retirement adequacy than fiscal, which has almost no effect on adequacy. This impact is stronger the less wealthy workers are, and mainly driven by the behaviour of biased households at all points in the wealth distribution.

Figure 1.10: Relative outcomes at retirement, by wealth and present–bias type

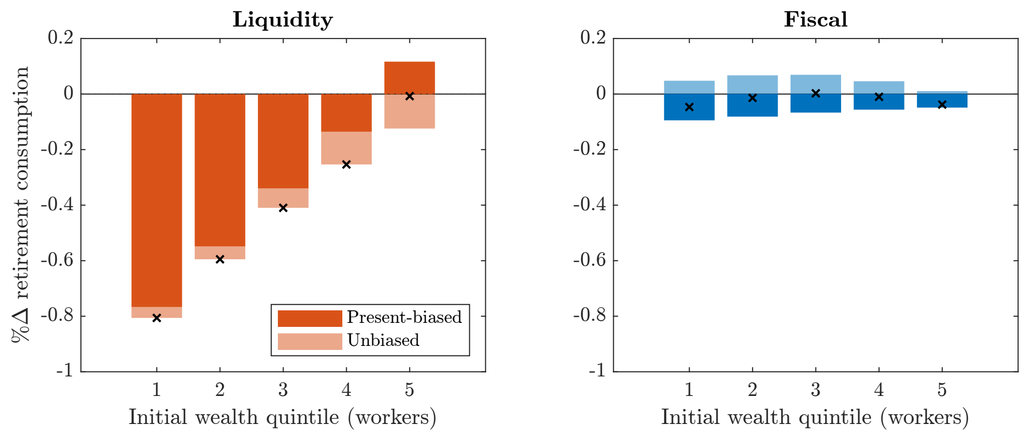

There are three forms of non–Ricardian behaviour at play in our model. First is the standard inability to smooth due to incomplete markets, second is the excess discounting caused by the OLG structure, and third is the over–consumption brought on by present–bias (Attanasio et al., 2024; Maxted et al., 2024). All stimulus relies on there being some non–Ricardian households. These results for retirement adequacy show that liquidity policy concentrates the duty of payment for stimulus mainly on the shoulders of the biased households. That is, exactly the group for whom the illiquidity in the retirement system is designed.

[43]Results are qualitatively the same for any retirement date.

1.7 Welfare

We have established that both approaches work to stimulate household consumption, but that they come with various contrasting implications in different dimensions. To bring this together and compare how their respective implications weigh against each other, and for whom, we compare welfare under each policy using a compensating variation (CV). At time–0, for each point in the state–space, we find the change in liquid assets that would be required to make the household indifferent to liquidity policy instead of the baseline fiscal intervention

$$\hat{v}(b + CV(x), a, z, \beta; 0, \gamma) = \hat{v}(x; T, 0)$$

We find this compensating variation from two perspectives: (a) the near–sighted evaluation finds the CV based on households' perceived value functions i.e. accepting any naivete they have, and (b) the long–run evaluation finds the CV based on a correct anticipation of how they will behave (Bernheim and Taubinsky, 2018; Naik and Reck, 2024; O'Donoghue and Rabin, 2006), defined in Equation 1.11. This compensating variation shows the intensity of a preference for fiscal policy—positive numbers indicate the need to compensate for using liquidity policy instead, and negative numbers indicate a willingness to *pay* for liquidity policy to be used.

1.7.1 Winners and losers from liquidity vs fiscal policy

Who benefits from liquidity policy over fiscal? Wealthy workers and retirees. Retirees don't stand to gain anything from the stimulus, but they must subsidise the stimulus to workers, and their consumption is further distorted by the changing taxes, as we saw in Figure 1.9. To highlight the importance of wealth for workers' preferences, Figure 1.11 plots the average compensating variation across wealth quintiles showing that enthusiasm for fiscal policy is declining in wealth. Wealthy workers are more likely to be Ricardian, more likely over–invested in their retirement account, and (like retirees) also expect to pay greater taxes under fiscal. Hence their preference for liquidity policy, which allows them to escape taxes and grants more flexibility in portfolio decisions that they can ignore if they wish. Figure 1.11 also disaggregates by type, showing that the majority of the benefit of fiscal over liquidity policy is accruing to biased households in the bottom quintile of the wealth distribution. Biased households are over–respresented in this quintile because their bias leads them into debt, and because of this they stand to lose the most from having lower retirement assets. This result confirms the intuition that many opponents to liquidity policy voiced—that it places the cost of support on the shoulders of the least well off, and the most in need of help to save for retirement.

Figure 1.11: Worker preference for fiscal policy, by wealth & type

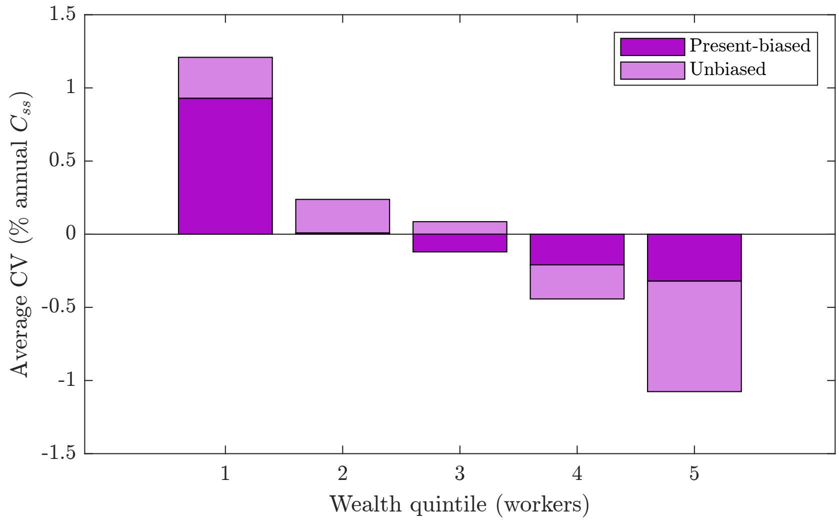

In aggregate, the best approach to stimulus depends on the welfare criterion. Liquidity policy is popular—it is preferred by an overwhelming majority, whether evaluated by households' near–sighted or long–run preferences (69% and 72%). And the average compensating variation across the whole population is negative—for the long-run evaluation the average CV is -1.6% of average annual stationary consumption—meaning there is more than enough money to fund compensation because the average person would be willing to pay to switch from fiscal to liquidity policy.

By contrast, a Utilitarian social planner would choose fiscal policy, under which average welfare is greater than under liquidity policy. Why the difference? Liquidity policy concentrates repayment for the stimulus on a small and relatively vulnerable group. It's popular because this small group doesn't constitute a majority. But it's a disaster for welfare because their marginal utility of consumption is *much* greater than for others in the population. Liquidity policy is preferred in average compensating variation terms for the same reason—their marginal utility is extraordinarily responsive to resources i.e. they can be bought off cheaply.

Bringing this together, the apparently optimal approach would be to use liquidity policy to do the stimulus, and to combine this with lump–sum transfers from wealthy workers, and the retired, to poorer and present–biased workers. Ironically, if these transfers were possible, then neither fiscal nor liquidity policy would be the optimal way to stimulate consumption. Instead, a program of lump–sum transfers moving liquidity from low– to high–MPC households would be better (Oh and Reis, 2012). Lacking such an instrument, we're left with the political question of how to weigh the needs of the many against those of the few.

1.7.2 Sensitivity

In this section we see how sensitive the baseline results are to different fiscal rules, targeting and magnitudes of the stimulus, and ex–ante retirement policy. In aggregate, liquidity policy is preferred to fiscal across all of our robustness checks around the baseline calibration for the retirement policy. It wins a majority vote in each case, and also costs a *negative* average compensating variation. This popularity is not due to naivete—the results in this section are all based on long–run evaluations, in which the evaluator correctly anticipates naive present–biased behaviour. In fact naive evaluations would be even more strongly in favour of liquidity policy. So the liquidity policy's widespread appeal is robust to these different assumptions. It is, however, quite regressive, with present–biased and less wealthy workers benefiting much more from fiscal stimulus. By contrast, if retirement policy were looser, then it is less binding on working–age people and this leads to reduced benefits from liquidity policy, tipping the balance back in favour of fiscal as the better approach to stimulus, both in aggregate and out of a concern for distributional impacts.

Fiscal rule symmetry The baseline results came from a symmetric fiscal rule—governments set tax rates to get debt back to target whether it is below or above this target. One of the implications of such a rule is that liquidity policy is accompanied by a persistent tax cut, to balance the government's books as the consumption tax's base is expanded. The populations we found to prefer liquidity policy may feel this way mainly because they like the tax cut. Although temporary tax cuts are often part of stimulus packages, we may also expect governments to apply fiscal rules asymmetrically, with the actual allowable debt level anywhere between zero and $\bar{B}$. In this case the fiscal rule would only activate when debt was above the upper bound, and otherwise taxes would stay fixed and deficits or surpluses allowed.

Table 1.6 shows that using an asymmetric fiscal rule reverses the average worker preference for liquidity policy. Workers on average now all prefer fiscal policy; the previous preference for liquidity policy among the unbiased workers was apparently driven by the tax cuts attached. Similarly, the retirees are still in favour of liquidity policy, but their enthusiasm is dampened (though not by enough to tip the scales to fiscal policy in aggregate).

Table 1.6: CV to prefer liquidity policy (% annual C_{ss})

	Baseline	Asymmetric	Consolidating	Accommodating
Total	-1.6	-0.1	-2.1	-0.8
Workers	-0.1	1.6	-0.5	0.8
Biased	*0.6*	*1.9*	*0.2*	*1.3*
Unbiased	*-0.8*	*1.3*	*-1.1*	*0.2*
Retired	-4.2	-3	-4.8	-3.5

Fiscal rule severity The timing of the fiscal rule may also matter. In the baseline results, the fiscal rule is set to mirror the EU's fiscal compact, following Galí (2020), which sets a target half–life for debt gaps of 13.5 years ($\mu = 0.0128$). Here we consider two alternative rules—the 'consolidating' one sets a stricter repayment schedule, with $\mu = 0.0256$ so the half–life is halved, and the 'accommodating' one sets a more lax schedule, with $\mu = 0.0063$. The accommodating rule is set so that government debt is 'retired' at the same rate as the working population. This is an interesting special case because it means the timing under liquidity and fiscal policy are essentially identical.

Figure 1.12: Aggregate impact of fiscal rule severity

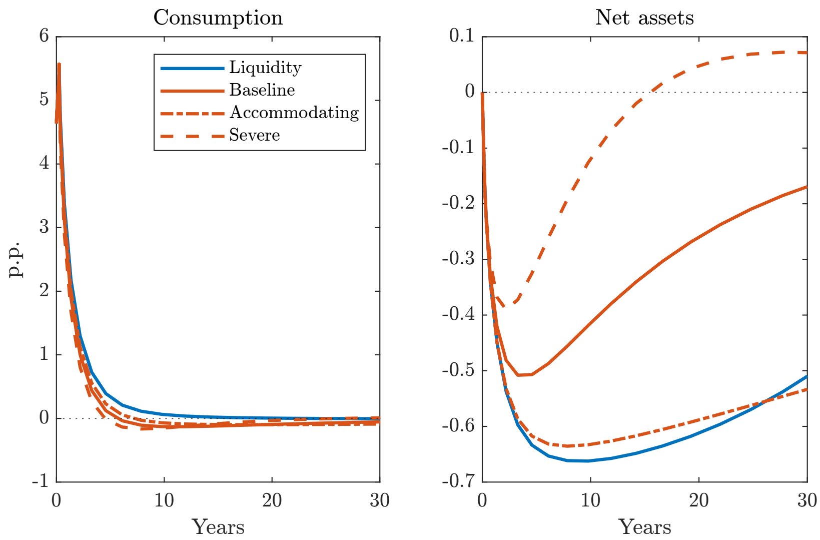

Figure 1.12 shows the aggregate consumption and net wealth plots under each alternative fiscal rule. They show what you'd expect—the drag on consumption is greater the stricter the rule, as net assets return to their stationary level quicker. Table 1.6 shows that the preference for liquidity policy diminishes with the intensity of the fiscal rule—it is greatest for the consolidating rule, followed by the baseline, and then the accommodating rule. As the fiscal rule relaxes, the taxes required to repay it are pushed further into the future and so they feel less onerous to the present generations, increasing the relative appeal of fiscal policy.

Targeting The baseline policy counterfactuals pitted liquidity policy against a fiscal transfer targeted at workers. This was designed to make the recipient group match, so they were more comparable. Here we explore how the main results differ when everyone receives the transfer.

If everyone receives the transfer then the total transfer per person needs to be smaller to achieve the same consumption stimulus. Specifically, to match the liquidity intervention we now need a transfer of $T^* = 0.0403$ for one quarter (versus 0.0706 when targeting workers), a total outlay for the government of around 10% less than when just targeting workers (retirees have a greater MPC than the average worker).

Table 1.7 shows that in this scenario, the overall long–run preference for liquidity policy is stronger. This is driven by a reversal of preferences for retirees, who now prefer to receive the fiscal transfer, and biased workers, who prefer to withdraw from their retirement accounts than receive a 40% smaller fiscal transfer. Unbiased workers are also much more in favour of the liquidity policy in this scenario. This result indicates how much the preference for, or against, redistribution was driving attitudes to fiscal policy.

Magnitude Table 1.7 shows long–run CVs under different magnitudes for the stimulus—50% and 150% the size of the baseline exercise—and their equivalent liquidity interventions. It shows a close–to linear relationship: liquidity is preferred overall in each case, with everyone's preference scaling up, or down, with the impulse magnitude.

Table 1.7: CV to prefer liquidity policy (% annual C_{ss})

	Baseline	Un–targeted	2.5% impulse	7.5% impulse	Lower ex–ante ξ
Total	-1.6	-2.1	-0.8	-2.6	1.4
Workers	-0.1	-4.9	-0.1	-0.1	4.1
Biased	*0.6*	*-3.9*	*0.2*	*0.8*	*5.9*
Unbiased	*-0.8*	*-5.9*	*-0.4*	*-1.0*	*2.3*
Retired	-4.2	2.8	-2.0	-7.0	-3.4

Looser retirement policy The baseline results all work in an environment with retirement policy set optimally in the exercise in Section 1.5. This policy is designed with the needs of newborns in mind, and one of its side–effects is to leave older workers generally over–invested in their retirement accounts, particularly if they are not present–biased. If we suppose instead that the government set retirement policy with the whole population in mind, but with the same participation constraint as the buy–in equilibrium, then the optimal policy sets contribution rates to about half what they are in the baseline, and with much smaller tax concessions as well

$$(\hat{\xi}, \hat{\varphi}) = (0.04, 0.12)$$

We redo the stimulus policy experiments with this calibration of the model's stationary solution. Fiscal and liquidity policy both now need to release more money to achieve the same stimulus $\hat{T} = 0.0816$ and $-\hat{\gamma} = 0.0875$ because the aggregate worker MPC is reduced to 0.2.

Table 1.7 shows that, with this starting retirement policy, fiscal is the better stimulus option for the average person. The average worker prefers it, with biased workers having the stronger (long–run) preference, whilst the average retiree still prefers liquidity policy because it means they avoid future taxes. Two forces drive the result for workers. With looser ex–ante retirement policy, there are both fewer wealthy workers who are over–invested in their retirement accounts, and more poorer workers who will suffer from the reduced retirement adequacy. Liquidity policy is thus an even more regressive stimulus option in settings where retirement contributions are low, or even not mandatory, as in the USA.

1.8 Conclusion

Household liquidity policy is becoming an increasingly popular tool used by governments seeking to stimulate the economy, growing from a few scattered uses during the global financial crisis, to a popular response used by over 30 countries during the COVID–19 recession. While a growing literature evaluates who withdrew and what they did with the money, our paper makes a different contribution. In short, we provide the first analysis of the distributional and welfare implications of household liquidity policy compared to traditional fiscal stimulus.

In our analysis, we capture many of the key trade-offs faced by policy makers thinking about the efficacy and equity concerns related to these two different approaches to stimulus. While this represents an important first step in understanding household liquidity policy, there are still a number of directions in which we could meaningfully extend our analysis, which we plan to pursue in future work.

At the most general level, our analysis could be expanded to capture the indirect, general equilibrium, effects of the stimulus policies in question. First, the stimulus approaches may have different implications for production. The model presented here works as a combined household–government block that can embed into a richer general equilibrium environment. Insofar as the aggregate effects coming out of this block are the same across the two approaches to stimulus, there will be no difference in their indirect effects on prices and production. However, if we extend the model to include labour supply, then inter–temporal distortions arising from taxes will drive a wedge between the two approaches that will bear on production—it will be more distorted under fiscal stimulus, than liquidity policy. Our future work will seek to quantify how important these production impacts are, and whether they have flow–on distributional implications.

Second, the source of the shock that precedes the stimulus may influence how we evaluate the two approaches. In our stimulus policy experiments we take it as given that the government wishes to achieve a 5% boost to household consumption. This allowed us to focus on the question of interest and to characterise the tradeoffs inherent in the two approaches. In practice, stimulus programs are implemented in response to shocks, and it may be that the nature of the original shock alters the relative desirability of liquidity versus fiscal policy.

Third, there may be important interactions with asset prices. Liquidity policy allows sales of assets at market lows, compounding the disadvantage to households that cash–out and spend the money because they realise capital losses (or at least reduced gains) by withdrawing from their retirement accounts and not reinvesting outside of them. Furthermore, in a similar vein, liquidity policy may exacerbate flight–to–safety dynamics we observe during times of aggregate stress, increasing funding costs for risky ventures.

Finally, we have modelled our stimulus policies as unanticipated shocks, so there is no role for anticipation. This was appropriate of the liquidity policies deployed in many countries during COVID–19 as this was the first time such measures were implemented in many places. Now that Pandora's box has been opened, however, forward–looking households may expect their retirement accounts to be somewhat more liquid than they were in the past. Such anticipation will alter households' engagement with retirement savings systems, and the repeated use of retirement resources for aggregate demand management will also alter the optimal design of retirement policy itself.

1.A Appendix

1.A.1 Present–biased household's problem

For a present–biased naif with bias parameter β, the HJB equation changes in two ways from the exponential equivalent—all terms except the flow utility are pre–pended by β, and the value in the HJB is the rational value, not the household's actual valuation.

Working life

$$\begin{aligned}\beta\rho v(x) = \max_{\hat{c},\hat{d}} &\left\{u(\hat{c}) + \beta\partial_b v(x)\cdot\dot{b}(x) + \beta\partial_a v(x)\cdot\dot{a}(x)\right\}\\ &+ \beta\sum_{z'}\lambda^{z\to z'}\left[v(x') - v(x)\right] + \beta\delta_R\left[v_R(x) - v(x)\right]\end{aligned} \tag{1.12}$$

This problem's FOC are

$$\begin{aligned}u'(\hat{c}(x)) &= \beta(1+\tau)\partial_b v(x) = \beta u'(c(x))\\ \partial_a v(x) &= \partial_b v(x)(1+\chi_d(\hat{d}(x),a)) + \kappa(x)\end{aligned}$$

Where the latter is the same as the exponential agent's, and so $\hat{d}(x) = d(x)$

Retired Similarly, the retired present–biased naif's problem is to solve the following.

$$\beta(\rho+\delta)v_R(x) = \max_{\hat{c}_R,\hat{d}_R}\left\{u(\hat{c}_R) + \beta\partial_b v_R(x)\cdot\dot{b}(x) + \beta\partial_a v(x)\cdot\dot{a}(x)\right\}$$

Where $\dot{b}(x)$ and $\dot{a}(x)$ are defined by Equations 1.1 and 1.2. This problem's FOC are

$$\begin{aligned}u'(\hat{c}_R(x)) &= \beta(1+\tau)\partial_b v_R(x) = \beta u'(c_R(x))\\ \partial_a v_R(x) &= \partial_b v_R(x)(1+\chi_d(\hat{d}_R(x),a))\end{aligned}$$

1.A.2 Euler equation

1.A.2.1 Present–biased naif, stationary

Following Maxted (2024), Appendix A.4, but with an added consumption tax. The following solves the Euler for the working household. First, find the derivative of the expected value Equation 1.5 with respect to the liquid asset b.

$$\begin{aligned}\rho\partial_b v(x) = {} & u'(c(x))\partial_b c(x) + \partial_{bb} v(x)\cdot\dot{b}(x) + \partial_b v(x)\partial_b\dot{b}(x)\\ &+ \sum_{z'}\lambda^{z\to z'}\left[\partial_b v(x') - \partial_b v(x)\right] + \delta_R\left[\partial_b v_R(x) - \partial_b v(x)\right]\end{aligned}$$

Apply the realised FOC i.e. that $u'(\hat{c}(x)) = \beta(1+\tau)\partial_b v(x)$ so $\partial_b v(x) = \frac{u'(\hat{c}(x))}{\beta(1+\tau)}$

$$\begin{aligned}(\rho - \partial_b r(b)b - r(b))\, u'(\hat{c}(x)) &= (1+\tau)\left(\beta u'(c(x)) - u'(\hat{c}(x))\right)\partial_b c(x) + \partial_b \hat{c}(x) u''(\hat{c}(x)) \cdot \dot{b}(x) \\ &+ \sum_{z'} \lambda^{z \to z'} \left[u'(\hat{c}(x')) - u'(\hat{c}(x))\right] + \delta_R \left[u'(\hat{c}_R(x)) - u'(\hat{c}(x))\right]\end{aligned}$$

Optimisation implies the relationship $u'(\hat{c}(x)) = \beta u'(c(x))$, which we can use to eliminate the first term on the RHS so the equation simplifies to

$$\begin{aligned}(\rho - \partial_b r(b)b - r(b))\, u'(\hat{c}(x)) &= \partial_b \hat{c}(x) u''(\hat{c}(x)) \cdot (r(b)b - (1+\tau_c)c(x) + other) \\ &+ \sum_{z'} \lambda^{z \to z'} \left[u'(\hat{c}(x')) - u'(\hat{c}(x))\right] + \delta_R \left[u'(\hat{c}_R(x)) - u'(\hat{c}(x))\right]\end{aligned}$$

Note that we can collect many of these terms into the time–derivative of expected marginal utility ($\mathbb{E}[du'(c(x))]/dt$), by Ito's Lemma, after we add and subtract $u''(\hat{c}(x))\partial_b \hat{c}(x)\hat{c}(x)(1+\tau)$

$$(\rho - \partial_b r(b)b - r(b))\, u'(\hat{c}(x)) = (1+\tau)u''(\hat{c}(x))\partial_b \hat{c}(x)(\hat{c}(x) - c(x)) + \mathbb{E}[du'(c(x))]/dt$$

And using CRRA utility we know $c(x) = \beta^{1/\sigma}\hat{c}(x)$

$$\begin{aligned}(\rho - \partial_b r(b)b - r(b)) &= (1+\tau)\frac{u''(\hat{c}(x))\hat{c}(x)}{u'(\hat{c}(x))}\partial_b \hat{c}(x)(1-\beta^{1/\sigma}) + \frac{\mathbb{E}[du'(c(x))]/dt}{u'(c(x))} \\ \mathbb{E}\left[\frac{\dot{c}}{c}\right] &= \frac{1}{\sigma}\left[\partial_b r(b)b - r(b) - \rho - \underbrace{\sigma(1+\tau)(1-\beta^{1/\sigma})\partial_b \hat{c}(x)}_{\text{Naive bias distortion}}\right]\end{aligned}$$

Hence the naif's Euler equation is a distorted version of the standard one, where the distortion scales with the bias, marginal propensity to consumer, and consumption tax.

1.A.2.2 Present–biased naif, dynamic

Using the same steps as above, but with an added term $\partial_t v_t(x)$ on the HJB equation, which eventually introduces an influence from drift in the tax rate.

$$\mathbb{E}\left[\frac{\dot{c}}{c}\right] = \frac{1}{\sigma}\left[(r_b(b)b + r(b)) - \rho - \sigma(1+\tau_t)(1-\beta^{1/\sigma})\partial_b \hat{c}_t(x) - \frac{\dot{\tau}}{1+\tau_t}\right]$$

The dynamic tax adds two distortions, relative to the stationary Euler. First, its level alters the bias distortion we found in the previous section, relative to its stationary level. Second, expected drift in the tax rate introduces inter–temporal smoothing distortions common to all agents. The analysis in the main text focuses on the latter.

1.A.3 Discretised value function updates and Kolmogorov forward equation

Suppose we know the household's optimal policy functions. The stationary solution to the household's problem can be expressed as a linear system, discretised over the state space, as follows

$$\rho\mathbf{V} = u(\mathbf{c}) + \mathbf{A}\mathbf{V}$$

Where $\mathbf{A}$ captures the finite–difference transition rates between all the states, from the perspective of the households (i.e. they anticipate retirement and death but not rebirth), taking into account their optimal policies, and $\mathbf{V} = \begin{bmatrix}\mathbf{v}' & \mathbf{v}_R'\end{bmatrix}'$ stacks the discretised working and retired values together. The solution to this linear system is

$$\mathbf{V} = [\rho\mathbf{I} - \mathbf{A}]^{-1} u(\mathbf{c})$$

Where updates are implemented using the semi–implicit scheme with update control step–size Δ

$$\mathbf{V}^{n+1} = [(1/\Delta + \rho)\mathbf{I} - \mathbf{A}^n]^{-1} [u(\mathbf{c}^n) + \mathbf{V}^n/\Delta]$$

We use the nested–drift algorithm in Sabet and Schneider (2024) to find the policy functions that define $\mathbf{A}^n$ and $\mathbf{c}^n$ for a given value guess $\mathbf{V}^n$, described in Appendix 1.A.4. In practice, we solve the retired value first, and then use this as an input into the working–life value solution; doing so reduces the computational burden of inverting the matrix in the semi–implicit update step. Following Achdou et al. (2022), the stationary measure[44] discretised over the same state grids ($\mathbf{g}$) is the solution to the linear system

$$0 = \tilde{\mathbf{A}}'\mathbf{g}$$

Where $\tilde{\mathbf{A}}$ adjusts $\mathbf{A}$ for the state transitions households do not anticipate, which are (1) rebirth as a zero–asset worker after dying in retirement, (2) forced retirement if illiquid assets reach the threshold, and (3) bankruptcy upon retirement if total assets are negative.

Present–biased agents For known policy functions, the solution process and formulae are the same, with the exception that their state transition matrix $\mathbf{A}^{\beta}$ solve for their *long–run* value function, not their perceived value.

1.A.4 The nested–drift algorithm

To find the optimal policy functions for a given value, this paper uses a version of the nested–drift algorithm introduced in Sabet and Schneider (2024). Combined with implicit updates of the value, described in Appendix 1.A.3, this is an approach to solving discretised systems of partial differential equations that is consistent and monotone, and so admits a solution that is stable to parameter choices. The below gives a sketch of the context for the algorithm, and its steps; see the original paper for more depth.

The problem Continuous time problems with multiple endogenous state variables, such as the one in this paper, typically have a solution defined by an HJB equation and first order conditions.

For the sake of communicating the algorithm, consider a simplified problem where households choose consumption and deposits between two accounts, where the latter are subject to adjustment costs. Their income is a random process y which switches between discrete states according to some transition intensities. The solution is thus defined by the HJB (and state transition equations) and first order conditions.

$$\begin{aligned}
\rho v(x) &= u(c) + \partial_b v(x) \cdot (r_b b + y - d - \chi(d,a) - c) + \partial_a v(x) \cdot (r_a a + d) \\
&\quad + \sum_{y' \in \mathcal{Y}} \lambda^{y \to y'} \left[v(x') - v(x) \right] \\
u'(c) &= \partial_b v(x) \\
\partial_a v(x) &= \partial_b v(x)(1 + \chi_d(d,a))
\end{aligned}$$

The above equations define the solution, but to implement it on a computer in a finite–difference scheme, we discretise the state–space x. Doing so introduces some imprecision that can be managed by using non–linear grids, placing a higher density of grid–points at areas where the value and policy functions are more likely to have curvature or kinks.

One form of imprecision is that the partial derivatives cannot be known exactly—we can find them with the forward and backward derivatives of the value at any point in the discretised state–space, or anywhere in between. So, which to choose?

Upwinding The standard approach to choosing between these options is to use 'upwinding'. That is, selecting the derivative that, when used, leads to drift in the state that goes in the direction assumed. In a single asset problem (e.g. if we fix $d = 0$ in the above), this consists of adding an additional condition to the solution. For a point in the state–space i the derivative is defined by

$$\partial_b v(x_i) = \begin{cases} \frac{v_{i+1}-v_i}{b_{i+1}-b_i} & \text{if } \dot{b}(x_i, c_i) > 0 \\ \frac{v_i - v_{i-1}}{b_i - b_{i-1}} & \text{if } \dot{b}(x_i, c_i) < 0 \\ u'(r_b b_i + y_i) & \text{if } \dot{b}(x_i, c_i) = 0 \end{cases}$$

Where the consumption policy is found from the FOC in the first two cases. In the third case, where there is no drift in the asset, the policy is identified by the drift equation and the derivative can be recovered from the first order condition.

In a two asset problem, things are more complex because one must consider the permutations of directions that each asset could be going in—both forward, both back, one still and the other drifting up, and so on. And also because multiple policies enter the one of the drift equations, and so the policies are jointly determined and can't necessarily be identified from the drift equations. In this case, they can be recovered using a root–finding algorithm based on the FOCs.

The nested–drift algorithm The algorithm we use here, originally from Sabet and Schneider (2024), solves this problem by exploring all possible cases (9 in a two–asset problem) of drift directions, using a root–finding algorithm to solve for policies in cases where necessary. The algorithm is monotone and consistent, and so satisfies the conditions necessary for local convergence (Barles and Souganidis, 1991). And it is also efficient because it places the most expensive root–finding steps last in the process, so they will only be reached if other alternatives have already failed.

The logic is as follows for each point in the state space x_{ijk}[45]:

1. Calculate key objects

   – The deposit policy that causes zero illiquid drift $\tilde{d} = -r_a a$

   – The directional derivatives of the value

$$V_b^F = \frac{v_{i+1,j,k} - v_{i,j,k}}{b_{i+1} - b_i} \quad \text{and} \quad V_b^B = \frac{v_{i,j,k} - v_{i-1,j,k}}{b_i - b_{i-1}}$$
$$V_a^F = \frac{v_{i,j+1,k} - v_{i,j,k}}{a_{j+1} - a_j} \quad \text{and} \quad V_a^B = \frac{v_{i,j,k} - v_{i,j-1,k}}{a_j - a_{j-1}}$$

2. Assume forward liquid drift, setting $\partial_b v = V_b^F$.

   – Calculate the consumption policy from the FOC

$$c^F = u^{-1}\left(V_b^F\right)$$

[45]Practically many of these operations are done simultaneously for all points in the state–space, but it's simpler to show one point at a time.

- Calculate the deposit policy from the FOC under each assumed illiquid direction, using the FOC for forward and backward drifts, and the state transition equation for zero drift. Check the resulting deposit policies against the drift in the illiquid asset they're based on

$$d^F = \begin{cases} d^{FF} & \text{if } d^{FF} > \tilde{d} \\ d^{FB} & \text{if } d^{FB} < \tilde{d} \\ \tilde{d} & \text{otherwise} \end{cases}$$

  This step nests the upwinding for the illiquid asset within the process for finding the liquid asset's drift.

- Find the liquid drift implied by (c^F, d^F)

$$\dot{b}^F = r_b b + y - c^F - d^F - \chi(d^F, a)$$

  and if it is positive move to the next point in the state–space. Otherwise continue.

3. Assume backward liquid drift, setting $\partial_b v = V_b^B$, and follow the equivalent steps to above to find the backward policies. This time, check whether the resulting liquid drift is backward

$$\dot{b}^B = r_b b + y - c^B - d^B - \chi(d^B, a) < 0$$

   If it is, move to the next point in the state–space. Otherwise continue.

4. Assume zero liquid drift, setting $\partial_b v = u'(c^0)$. Given this assumption, we know the consumption policy from the state transition equation

$$c^0 = r_b b + y - d^0 - \chi(d^0, a)$$

   and the deposit policy will be defined implicitly by the FOC

$$\partial_a v(x) = u'\left(r_b b + y - d^0 - \chi(d^0, a)\right)(1 + \chi_d(d^0, a))$$

   And we can solve for d^0 using this equation, and searching through different regions.

   (a) Check if $d^0 > \tilde{d}$ using $\partial_a v = V_a^F$. If not, continue

   (b) Check if $d^0 < \tilde{d}$ using $\partial_a v = V_a^B$. If not, continue

   (c) Set $d^0 = \tilde{d}$

This algorithm results in the household's optimal policy functions $(\mathbf{c}, \mathbf{d})$ for a given value $\mathbf{V}$. It nests the up–winding of the illiquid asset within that of the liquid one. This is the efficient order to use because the illiquid asset's state transitions depend on only one policy function, and so the costly route–finding in the final steps is only necessary if all alternatives have been exhausted.

1.A.5 Conditional expectations in continuous–time

The goal in this section is to find the expected path of a variable (a policy or state variable), given a starting point in the state space (x_0), from the perspective[46] of an agent at that point in the state–space.

$$y^e(x_0, t_n) = \mathbb{E}[y(x, t_n)|x_0 = x]$$

Suppose we have solved our problem, with time discretised over a sequence of N segments of a span from $t \in [0, T]$, where t_n denotes the end–point of the n–the segment such that $t_1 > 0$ and $t_N = T$. The solution is a sequence of policy vectors, transition matrices, and distributions (the measure, not the density) stacked over the state–space

$$\{\mathbf{y}(t_n), \mathcal{A}(t_n), \mathbf{g}(t_n)\}_{n=1}^{N}$$

Where $\mathbf{y}(t_n)$ is the policy that applies during the span $\Delta(n) = t_n - t_{n-1}$ and $\mathbf{g}(t_n)$ is the state–distribution at the point in time t_n.

Finding the distribution updates The sequence of distributions is found using the implicit, discretised Kolmogorov Forward equation, following Achdou et al. (2022)

$$\frac{\mathbf{g}(t_n) - \mathbf{g}(t_{n-1})}{\Delta(n)} = \mathcal{A}(t_n)'\mathbf{g}(t_n)$$

Where the starting distribution is the stationary solution $\mathbf{g}(t_0) = \mathbf{g}_{ss}$ The discretised KF equation rearranges into the implicit updating equation

$$\mathbf{g}(t_n) = \mathcal{L}(t_n)\mathbf{g}(t_{n-1}) = \prod_{j=1}^{n} \mathcal{L}(t_j)\mathbf{g}_{ss}$$

With the implicit transition matrix defined by $\mathcal{L}(t_n) = [\mathcal{I} - \Delta(n)\mathcal{A}(t_n)']^{-1}$.

Aggregation To build a sequence of aggregate policies, we use

$$Y(t_n) = \mathbf{y}(t_n)'\mathbf{g}(t_{n-1})$$

[46]Important if (a) expectations are not fully rational, or (b) agents don't internalise some transitions like reincarnation.

So the expected aggregate at the n–th point in the time grid is the policy from that point, integrated over the distribution at that point. If we substitute the definition of the distribution into this expression we have

$$Y(t_n) = \begin{cases} \mathbf{y}(t_1)'\mathbf{g}_{ss} & \text{if } n = 1 \\ \mathbf{y}(t_n)' \prod_{j=1}^{n} \mathcal{L}(t_j)\mathbf{g}_{ss} & \text{if } n > 1 \end{cases}$$

Focusing on the cases where $n > 1$, we can see that the aggregate is made up of three terms—the future policy: $\mathbf{y}(t_n)$, the starting distribution: $\mathbf{g}_{ss}$, and a matrix that affects the expected movement of measure around the state–space between the starting point and t_n: $\prod_{j=1}^{n} \mathcal{L}(t_n)$.

We can collect these terms in two different ways to get the same aggregate. In the standard aggregation, the matrix is applied to the starting distribution, and so the expected aggregate in t_n is the policy in that period integrated over the distribution at the beginning of that period. If we instead apply the matrix to the future policy, then it produces the expected policy in t_n from the point of view of a point in the state space at t_0: $\mathbf{y}^e(t_n)$. Aggregation is then achieved by

$$Y(t_n) = \mathbf{y}^e(t_n)\mathbf{g}_{ss} \text{ where } \mathbf{y}^e(t_n) = \mathbf{y}(t_n)' \prod_{j=1}^{n} \mathcal{L}(t_j)$$

The two approaches measure the same aggregate, but they create very different sub–aggregate objects—the first a future distribution, and the second an expected policy. The latter $\mathbf{y}^e(t_n)$ is exactly the conditional expectation we are seeking—it tells us the *expected* value of the policy y at some future point in time t for each *starting–point* in the state–space.[47]

With the conditionally expected policy in hand, we can use it for other purposes than aggregation. Say we wish to describe expected policies by some quantile of the distribution over states. This is very easy with above object, with only two steps required:

1. Calculate the relevant quantile for each point in the starting state space
2. Find the average of the conditional policy, conditional on the quantile x_0 is in

$$y^e(t_n, q) = \mathbb{E}[y^e(t_n, x_0) | x_0 \in q] = \frac{\int_{x \in q} y^e(t_n, x) dG(x)}{\int_{x \in q} dG(x)}$$

[47] This is a discretised implementation of the Feynman–Kac formula.

1.A.6 Marginal propensity to consume

Following Achdou et al. (2022) we calculate marginal propensity to consume over a discrete time–interval τ (always one–quarter in this paper) using the following formula

$$MPC_\tau(x_0) = \partial_b C_\tau(x_0) \;\; \text{where} \;\; C_\tau(x_0) = \int_0^\tau c^e(x_0, t) dt$$

Where the conditional expected consumption policy $c^e(t, x)$ is defined in Section 1.A.5.

Chapter 2

Self–Control and Early Withdrawal from Retirement Accounts

Patrick Schneider and Patrick Moran[1]

Abstract

Using a survey-elicited measure of psychological self–control and a policy change in Australia during COVID–19, we find that self–control issues significantly predict early withdrawals from retirement accounts. Individuals in the top quintile of self–control issues are 60% more likely to withdraw than those in the bottom quintile. Self–control is a stronger predictor of early withdrawal than other behavioral factors such as financial literacy, planning horizons, or personality traits. The effects are economically meaningful: eliminating self–control issues could have reduced early withdrawals by 24%—as large as the effect of adverse income shocks on withdrawals during COVID–19.

[1]An earlier version of this paper circulated with the title 'Situational and Behavioral Determinants of Early Withdrawal from Retirement Accounts'. This paper uses data from the Household, Income and Labour Dynamics in Australia (HILDA) Survey. The HILDA Project was initiated and is funded by the Australian Government Department of Social Services (DSS) and is managed by the Melbourne Institute of Applied Economic and Social Research (Melbourne Institute). The findings and views reported in this paper, however, are those of the authors and should not be attributed to either DSS or the Melbourne Institute. The views expressed in this paper are solely those of the authors and do not represent the views of the Federal Reserve Board or the Federal Reserve System.

2.1 Introduction

Many countries encourage households to save in individual retirement accounts that are illiquid during working life. One of the principal rationales for imposing this illiquidity is the view that many people are myopic and lack the ability to save for retirement on their own (Feldstein, 1985). Indeed, such systems play an important role in ensuring retirement adequacy, as many individuals reach retirement age with virtually no financial assets outside of the formal retirement system (Poterba, 2014). The defining feature of these accounts is that they are illiquid before retirement, in the sense that early withdrawals are either prohibited or attract substantial tax penalties (Beshears et al., 2015).

Whilst the rationale for these policies is that some people lack self–control, there is little domain–specific evidence on the role of behavioral biases in driving retirement leakage. This gap in our knowledge is largely driven by data limitations that impede simultaneous observation of self–control issues and real–world financial decisions. In this paper, we make two main contributions. First, we test one of the key rationales behind making retirement wealth illiquid, by evaluating what happens when retirement wealth is made more liquid. We find that self-control issues increase the likelihood of early withdrawal, thus supporting the view that illiquidity may be important for such individuals' later-life well-being. Second, after confirming the importance of self-control, we evaluate how it compares quantitatively to more situational factors that we already know are important. We find that self–control issues are as important, in aggregate, as situational factors like income or job loss.

Our study is possible because in recent years an increasing number of countries have turned to relaxing the illiquidity of these accounts during periods of aggregate distress. Such policies are a novel approach to supporting households and stimulating the economy, which we term 'household liquidity policy' in Schneider and Moran (2024a).[2] These policies attracted substantial controversy. The hope was that they could help liquidity constrained households who have been hit by adverse shocks, while also stimulating the economy in a downturn without burdening government budgets. But the fear is that they would undermine the goals of the retirement system in a way that disadvantages the people it's designed to help i.e. those with self–control issues.

We study an unexpected policy change in Australia that gave working-age individuals the ability to withdraw up to $20,000 AUD from their individual retirement accounts during the COVID–19 pandemic. To evaluate whether self-control issues were related

[2]Some examples include Denmark in 2009, and Australia, the United States, and many others in 2020 (OECD, 2021). And recent evidence shows that early access to retirement wealth can have a substantial impact on consumer spending. Hamilton et al. (2024) study the Australian policy that gave individuals the ability to tap into their individual retirement accounts during the COVID–19 pandemic. The authors find that individuals who withdrew spent more than 40% of the withdrawn money within the first two months.

to early withdrawals, we exploit a unique opportunity offered by the Household Income and Labour Dynamics in Australia (HILDA) survey. This survey is uniquely suited to our purposes as it includes measures of (i) early withdrawal decisions, (ii) self–control issues, and (iii) situational factors like income, wealth, and adverse shocks, as well as some measures of other psychological traits including financial literacy, planning horizons, and personality. The self–control measure comes from the 13 question Brief self–control Scale (BSCS), which is well-established in the psychology literature, and was included in the HILDA survey in the 2019 wave, giving us a measure of self-control that pre-dates COVID-19.[3] This survey gives us a unique opportunity to evaluate how much self-control issues actually color people's engagement with retirement savings.

Among working age Australians, we find that one in seven (roughly 14%) took advantage of the opportunity to withdraw from their retirement account during COVID–19. In line with the existing literature, we find that those who withdrew were on average younger and had lower income, fewer liquid resources, and more children (Hamilton et al., 2024; Bateman et al., 2023). Building on the existing literature, we provide the first evaluation of the role of self–control issues. In the raw data, we see that individuals in the top quintile of self–control issues are about 60 percent more likely to withdraw from their retirement account than those in the bottom quintile.[4]

Recognising that self-control issues are likely related to other factors like wealth that might drive demand for liquidity, we investigate the marginal importance of self-control by estimating a series of regressions where we include a growing set of situational and other behavioural behavioral characteristics. We show that self–control issues are significantly and meaningfully correlated with early withdrawal when controlling for demographics, income, adverse shocks, planning horizon, financial literacy, personality traits, and wealth. High self–control issues are associated with an 8.6 percentage point higher probability of early withdrawal, which is similar to the marginal effect of having more than three children (associated with an increase of 8.5 percentage points).

We also show that self-control issues are the most economically meaningful in comparison to other commonly used behavioural biases. The unconditional relationships between withdrawals and short planning horizons or financial illiteracy are as stark as for self-control. But planning horizons cease to be a significant predictor after we control for wealth, indicating that their effect is not direct. Financial literacy, by contrast, does appear to have a direct relationship to behaviour not mediated entirely by its effect on wealth, but the average marginal effect is weaker (increasing the probability of withdrawal by 2.8 percentage points) and illiteracy is also less common in the population.

We also document an important role for situational factors, in line with Coyne et al.

[3]Developed by Tangney et al. (2004), this measure has become popular in the psychological literature, and has been increasingly used in economic studies (e.g. Cobb-Clark et al., 2022). See Section 2.2.2.

[4]

(2022) and Andersen et al. (2024). To the best of our knowledge, we are the first to evaluate the relative importance of situational versus behavior determinants of demand for liquidity. We find that situational factors – in particular, adverse shocks – are ultimately more important than behavioral factors when it comes to predicting which individuals withdraw from their retirement accounts. Unemployed individuals are 5.8 percentage points more likely to withdraw, on average, while individuals who have suffered a pandemic-related negative income shock are 19.0 percentage points more likely. Thus, while both personality traits and adverse shocks are correlated with early withdrawal, we find that the average marginal effect of adverse shocks is larger, even though they are relatively low frequency events.

How do these individual–level factors contribute to the aggregate share of early withdrawals? Overall, we find that self–control issues account for a similar share of aggregate withdrawals as adverse income shocks. More specifically, we perform a back-of-the-envelope calculation to quantify the overall share of withdrawals that can be independently attributed to our main variables of interest.[5] While adverse income shocks are a stronger predictor of withdrawal at the individual level, they are also much less common, while self–control issues are relatively dispersed and widespread. As a result, eliminating either adverse income shocks or self–control issues would reduce the share of early withdrawals by about a quarter in both cases, all else equal.

We evaluate the sensitivity of our results to a number of different assumptions. First, we find that the importance of self–control is meaningful and precisely estimated when controlling for liquid and illiquid wealth, despite the fact that wealth is endogenous and may also be influenced by self–control issues. As a result, self–control issues matter above and beyond their documented effect on wealth (Attanasio et al., 2024). Second, while our baseline specification measures self–control as the first principal component of the BSCS, we find that our results are robust to an orthogonalized two–factor version of the BSCS, generally termed impulsivity and restraint by the existing literature (Maloney et al., 2012). In this case, we find that only the first factor (impulsivity) is significantly correlated with early withdrawal. The results presented here are not causal.[6] But this paper takes a first step to show that there is a positive relationship between self-control issues and retirement leakage, that it is economically significant, and as important in aggregate as more situational factors.

Related literature. Our analysis contributes to a growing empirical literature that evaluates demand for liquidity in retirement systems. In doing so, we bring together two separate strands of literature. On one hand, there's a large and growing literature

[5]We should note that this should be viewed as a lower bound estimate, as we also control for wealth and other factors that may be correlated with either self–control or adverse shocks.

[6]Rather causal inference requires a selection on observables assumption.

documenting situational demand for liquidity (Amromin and Smith, 2003; Andersen et al., 2024; Bateman et al., 2023; Coyne et al., 2022; Goda et al., 2022; Goodman et al., 2021; Hamilton et al., 2024). These papers document that individuals are more likely to withdraw from their retirement accounts (if allowed) following job loss, divorce, or other adverse shocks. This empirical literature roughly mirrors the "situational view" of illiquidity highlighted by Kaplan and Violante (2014, 2022). On the other hand, while there's a growing literature documenting the link between preferences and wealth accumulation (Ameriks et al., 2003; Banks et al., 2010; Epper et al., 2020; Goda et al., 2019; Stango and Zinman, 2023), we know of no studies evaluating the empirical link between preference heterogeneity and differences in the demand for liquidity. This latter mechanism roughly mirrors the "behavioral view" of illiquidity highlighted by Laibson (1997), Attanasio et al. (2024), and Maxted et al. (2024). We bring together these two different literatures, first by providing novel evidence on the role of behavioral bias in explaining demand for liquidity, and second by evaluating the relative importance of this bias compares to situational factors.

Our results speak to the complex trade-offs faced by policy makers that are interested in giving immediate financial relief to households, while also ensuring adequate resources for retirement. As such, our results are informative for the growing literature that uses quantitative models to evaluate the design of retirement account when agents suffer from present-bias (Beshears et al., 2025; Andersen et al., 2024; Choukhmane and Palmer, 2024; Schneider and Moran, 2024a). One challenge facing this literature is that we have little empirical evidence on how important present–bias is for driving interactions with the retirement system. Our analysis provides the first empirical evidence on how self–control affects demand for liquidity and, as such, may serve as important evidence to discipline such models as well as affirming that these self–control issues are important. Further, our survey-based approach is complementary to the studies that estimate present-bias in life-cycle consumption saving models (Kovacs et al., 2021; Laibson et al., 2024) which generally assume homogeneous preferences to identify the average level of present-bias. We take a very different approach, directly measuring self–control using a popular instrument from the psychological literature, then evaluating its relationship to observed financial decisions.

Understanding demand for liquidity and the determinants of early withdrawal is important for numerous reasons. First, given the growing prevalence of defined contribution retirement accounts, there's widespread concern about leakage from these accounts and the potential consequences for retirement adequacy. Goodman et al. (2021) find that for every dollar put into the US retirement system, 22 cents come out as early withdrawals. Choukhmane et al. (2024) show that early withdrawals are common, especially among low-income and minority savers, with almost one-quarter of Black savers making an early withdrawal each year. Goda et al. (2022) show that penalized withdrawals are more common among recent claimants of unemployment insurance. Second, there's growing interest in using retirement accounts to stimulate the economy. Indeed, at least 30 countries allowed early withdrawals or delayed contributions during COVID–19 as a way to support distressed households (Madeira, 2024; OECD, 2021). Third, recent research shows that such policies have a large impact on household spending, such as Kreiner et al., 2019b studying the release of retirement savings in Denmark in 2009 or Hamilton et al., 2024 studying the release of retirement savings in Australia in 2020. And while a growing literature mentions the potential role of behavioral biases (e.g. Bosch et al., 2020; Hamilton et al., 2024; Bateman et al., 2023), to the best of our knowledge, no previous study has evaluated how such biases may contribute to early withdrawal decisions.

Our analysis complements an influential recent paper by Hamilton et al. (2024) who also evaluate the early release of retirement wealth in Australia. The authors analyze the situational determinants of early withdrawal, then study how the policy affects consumer spending. Using high frequency spending data, the authors find that individuals who withdraw spend roughly 40% of the withdrawn funds within the first eight weeks. The authors argue that this indicates a sensitivity of consumption to income that is far greater than traditional models can predict, even with liquidity constraints, and show that the addition of present–bias is able to rationalize the behavior. We take a very different approach, exploiting survey–based measures of preferences to evaluate how early withdrawal varies with behavioral versus situational factors. Our results provide new, direct evidence that self–control issues played an important role in early withdrawals, supporting the interpretation by Hamilton et al. (2024).

Our analysis also complements recent papers that study the determinants of early withdrawal following the relaxation of withdrawal restrictions, a concept that we term 'household liquidity policy' in Schneider and Moran (2024a). In the Australian setting, Bateman et al. (2023) conduct a real-time survey and find that self–reported reasons for withdrawal were generally related to consumption smoothing. Preston (2022) documents the importance of numerous factors that contribute to early withdrawal, with a particular emphasis on income, job loss, financial literacy, and gender. They find that job loss and low financial literacy are important predictors of early withdrawal, a result that we confirm as well. In the US, Goda et al. (2022) evaluate the change in withdrawals at the age when the early withdrawal penalty is lifted, finding an important role for liquidity constraints and unemployment. We build upon the above studies by evaluating a wide range of behavioral factors that may influence demand for liquidity and then comparing such factors to the situational determinants that have already received substantial attention.

Finally, our paper builds upon a large literature using survey-based preference measures to evaluate the relationship between preferences and wealth. Ameriks et al. (2003) show that one's propensity to plan is correlated with wealth. Banks et al. (2010) show that measures of numeracy and cognitive ability are associated with greater wealth both before and after retirement. Goda et al. (2019) show that survey measures of present-bias and exponential-growth bias are both meaningful predictors of retirement wealth. Epper et al. (2020) document a strong correlation between time discounting and individuals' position in the wealth distribution. Stango and Zinman (2023) measure a wide range of behavioral biases and document that present-bias is negatively correlated with wealth and other financial conditions. Relative to the existing literature, we believe we are the first to evaluate the relationship between preferences and demand for liquidity. Our results indicate that self–control issues contribute meaningfully to retirement leakage.

2.2 Setting and Data

During the COVID–19 pandemic, many countries implemented policies allowing individuals to access their retirement savings to provide financial relief during the economic crisis. In the United States, the CARES Act permitted individuals to withdraw up to $100,000 from their retirement accounts without the usual penalties. Similarly, Canada allowed withdrawals from the Registered Retirement Savings Plan (RRSP), Australia allowed individuals to withdraw up to $20,000 AUD from superannuation funds, and Chile permitted withdrawals from their mandatory individual retirement accounts up to 10% of accumulated savings. Overall, at least 30 countries implemented policies that allowed for early withdrawal or delayed contributions to retirement accounts during the pandemic (Madeira, 2024; OECD, 2021).[7] Each of these events presents an opportunity to examine how people react to the new–found liquidity, and to identify how much the retirement system had been constraining people with self–control.

In this paper, we focus on the Australian experience because of the unique opportunity presented by high quality data measuring both self-control issues *and* early withdrawals. No such data exists, to the best of our knowledge, for any of the other countries that have used these policies recently. Furthermore, Australia's early withdrawal policy was one of the larger programs of this kind, and has already attracted considerable attention in the recent literature. As well as exploring how important self-control issues are for driving retirement saving behaviour, we can also compare their importance for engagement in this particular policy, relative to other better studied factors like job-loss.

[7]While such policies exploded in popularity during the COVID–19 pandemic, there were some pre-pandemic instances as well. For instance, Denmark in 2009 implemented a policy to stimulate the economy by allowing individuals to tap into their previously illiquid retirement accounts (Kreiner et al., 2019b).

2.2.1 Institutional Setting

Australia's system of mandatory retirement savings, known as superannuation, began in 1992 with the introduction of the Superannuation Guarantee (SG) scheme. Initially, the SG required employers to contribute 3% of employees' earnings into a superannuation fund. This rate increased incrementally over the years: to 6% by 1999, 9% by 2002, and 9.25% by 2013, with plans to eventually reach 12% by 2025. Superannuation accounts receive substantial tax benefits and are almost entirely illiquid before 'preservation age' (60 for most current workers), with only a few exceptions (e.g., financial hardship, compassionate grounds, and terminal illness). Australia's approach is similar to numerous other countries with mandatory defined-contribution (DC) systems.[8]

COVID–19 Early Release of Super program. In 2020, Australia introduced a policy allowing individuals to access up to $10,000 AUD from their individual retirement accounts by July 1, 2020, and an additional $10,000 AUD by December 31, 2020. The policy was widely publicized and saw significant uptake, with millions of Australians withdrawing funds. Most individuals who withdrew decided to withdraw the maximum of $20,000 AUD (Hamilton et al., 2024). Despite its popularity, the policy was controversial. Critics argued that it could undermine retirement security, while supporters saw it as a necessary measure for immediate financial relief.

Applications for early withdrawal from superannuation accounts were made online and required minimal supporting documentation (Bateman et al., 2023). While eligibility was supposed to be limited to individuals who had been financially affected by the pandemic, the conditions were relatively broad and covered more than 70% of the working age population (Hamilton et al., 2024).[9] Further, eligibility was entirely self–reported with no independent governmental verification.[10]

[8] Bateman et al. (2001) and Beshears et al. (2015) discuss mandatory saving in DC accounts. Countries with similar policies include Canada, Chile, Denmark, Peru, Vietnam, Singapore, and Sweden. While the U.S. also has mandatory retirement contributions in the form of Social Security, the Australian system differs in a few key aspects. First, Australia's SG scheme is a defined–contribution rather than defined–benefit pension system, as assets are directly earmarked to individuals, rather than pooled across society. Second, since contributions are mandatory and more uniform across the income distribution, Australia's superannuation system is designed to provide a close–to flat replacement rate of working–life income. In contrast, the U.S. system of mandatory social security contributions provides a higher replacement rate for those at the bottom of the income distribution.

[9] Residents needed to meet one of three criteria: (1) unemployment, (2) eligibility for a range of other government benefits, or (3) had been made redundant, working hours reduced by more than 20% or, if a sole trader, business suspended or revenue reduced by more than 20%.

[10] Despite not binding in practice, the presence of these rules may have deterred people who could have withdrawn, and wanted to, due to a perception that they would be punished for doing so.

2.2.2 Data

We use data from the Household, Income and Labour Dynamics in Australia (HILDA) Survey, a long-running longitudinal study that collects annual data on employment, income, and wealth from a large sample of Australian households. Initiated in 2001, HILDA follows a panel structure similar to the U.S. Panel Study of Income Dynamics (PSID), but with a substantially larger sample size: roughly 17,000 individuals across more than 8,000 households in the most recent wave. The survey collects detailed data on demographics, family structure, employment, income, and wealth. Further, HILDA is relatively unique among nationally representative longitudinal surveys in its occasional collection of detailed psychological traits, used by a variety of past studies (see e.g. Todd and Zhang, 2020). We use data from waves 18 to 21, collected between 2018 and 2021. The Brief Self–control Survey was conducted in HILDA's wave 19, between July 2019 and February 2020.

Sample selection. We restrict attention to individuals between the ages of 21 and 58 in 2020. The upper limit is motivated by the fact that 58 is the 'preservation age' at which superannuation accounts became partially liquid regardless of retirement status.[11] We further restrict the sample to individuals who were interviewed in all four survey waves between 2018 and 2021, given our desire to measure wealth (recorded every 4 years, last measured before the pandemic in 2018), personality traits (measured in 2019), and early withdrawals (measured in 2020 and 2021). Among this group, we further restrict our sample to individuals who responded to the 2019 self-completion questionnaire (SCQ), which measures personality traits and a host of other factors, and who did not miss 3 or more questions on the Brief Self–Control Survey.[12] Together, these restrictions leave us with a sample of 7,214 individuals, with observations spanning the 2018-21 waves of the survey.

Throughout our analysis, we focus on individual level data, as superannuation accounts are individually owned and controlled, and the decision to withdraw during COVID–19 was an individual decision. Almost all of our variables of interest are measured at the individual level, including self–control, and the other psychological variables with the exception of wealth, which is measured at the household level.

[11] At the time of the policy change in 2020, anyone aged 58 and above was allowed to implement a 'Transition To Retirement' strategy, moving any existing superannuation balance into a 'pension' account, exempting it from all taxes and imposing minimum and maximum withdrawal limits; people in this age group are still allowed to work and so have a tax arbitrage—they can withdraw the maximum from their pension account, and voluntarily contribute more into their superannuation account, reducing the tax liability to 15% on any earned income, up to the limit on concessional contributions.

[12] The SCQ is a 20 page survey consisting of questions that are difficult to administer in time-effectively in a personal interview. Conditional on meeting our other sample requirements, 94.4% of individuals complete the SCQ, and 97.4% of SCQ respondents answer all 13 questions of the BSCS.

Early Withdrawal. HILDA respondents were asked "Did you withdraw money from any of your superannuation [pension] funds because of the coronavirus crisis?" and, if yes, "What was the amount withdrawn?" In our data, 13.8 percent of working-age individuals withdrew early, matching estimates from other papers (Bateman et al., 2023; Hamilton et al., 2024) and official statistics. [13]

Self–control. In 2019, HILDA survey participants were asked to complete the Brief Self–Control Scale (BSCS), which is widely used in the psychological literature, and consists of 13 targeted questions on impulse control and goal adherence. Established by Tangney et al. (2004), the scale is designed to measure self–control—"the capacity to regulate attention, emotion, and behavior in the presence of temptation"—by asking respondents to score on a scale of 1 to 5 how much a series of 13 statements applies to them. The statements include "I am good at resisting temptation," "I often act without thinking," and "I am able to work effectively toward long-term goals." Previous work has found that this scale shows good internal consistency and retest reliability (Bertrams and Dickhäuser, 2009; Tangney et al., 2004), and that higher self–control is linked to better financial outcomes and disciplined behavior (Cobb-Clark et al., 2022).

We use Principal Components Analysis (PCA) to reduce variation in the 13 BSCS items to one dimension, a standard approach in the psychology literature (e.g. Manapat et al., 2021). The first principal component explains roughly one-third of the variation across the standardized 13 item scale, and the sign of each loading is as we would expect given the direction of phrasing.[14] We rescale the first principal component so that it ranges between zero and one, where zero represents no self–control issues on all 13 items, while one corresponds to full self–control issues. Figure 2.1 shows the distribution of self–control issues. Overall, we see that they are relatively widespread but vary substantially across individuals, with an average of 0.39, standard deviation of 0.16, and a long right tail to the distribution.

Other Psychological Traits. Whilst our focus is on self-control, there are various other behavioural biases that may justify illiquidity in retirement systems. To explore their relative importance, we augment our set of control variables with other psychological measures including financial literacy, the Big Five personality traits, and planning

[13] Previous research has shown that the vast majority of individuals who withdrew decided to withdraw the maximum amount permitted each round (Bateman et al., 2023; Hamilton et al., 2024). As a result, we focus on the discrete decision to withdraw, rather than the continuous decision of how much to withdraw. Indeed, analysis using the amount withdrawn (not reported) yields similar qualitative results to our baseline analysis.

[14] Appendix 2.A.1 shows the full list of items, the distribution of responses, and the factor loadings. We experimented with an additional factor in our empirical analysis, but found that it did not meaningfully change the results, and that only the first factor was significantly correlated with withdrawal.

Figure 2.1: Distribution of self–control Issues

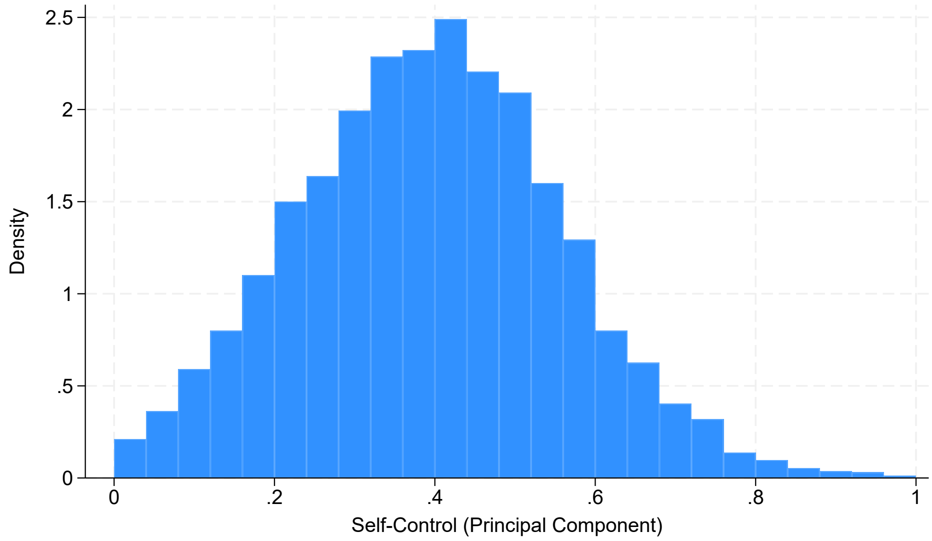

horizon.

Financial Literacy is measured using the well-established "Big-3" measure of Lusardi and Mitchell (2014), which is a binary measure equal to one if the respondent correctly answered all three questions related to interest rates, inflation, and diversification.

The Big Five personality traits – Openness, Conscientiousness, Extroversion, Agreeableness, and Neuroticism – are measured through a series of standardized questions. Past research has used HILDA data to document the importance of the Big Five personality traits for schooling and labor market outcomes (Flinn et al., 2018; Todd and Zhang, 2020).[15]

Planning horizon is measured based on individuals' response to the question "In planning your savings and spending which of the following time periods is most important to you?" The respondent can choose: next week; next few months; next year; next 2-4 years; next 5-10 years; or 10+ years ahead. While planning horizon is not a perfect measure of time preference, it is often used as a proxy when a direct measure does not exist in the data (e.g. Barsky et al., 1997; Samwick, 1998; Brown and Van der Pol, 2015). Past research has shown that planning horizon is correlated with time preference.[16]

[15] While the use of the Big Five traits in explaining economic outcomes is now well-established among economists (Almlund et al., 2011; Borghans et al., 2008; Heckman et al., 2021), there is much less evidence on self–control, perhaps because it has only recently been incorporated into large-scale surveys.

[16] Adams and Nettle (2009) show that planning horizon and discount rate are correlated, -0.19, with a p value < 0.001. While individuals with a higher time preference rate are likely to have a shorter planning horizon, socio-economic status and life expectancy are also likely associated with planning horizon. We thus control for income, wealth, and age in our empirical analysis.

Wealth. The HILDA Survey collects detailed data on household wealth through approximately 20 to 30 questions, covering a wide range of asset and liability categories. This includes information on real estate, financial assets, vehicles, business investments, and liabilities such as mortgages and personal loans. Given the reporting burden, the wealth module is only administered every four years. We use the most recent wave of wealth data prior to COVID–19, collected in 2018.

We divide wealth between three categories. The first is liquid wealth: the sum of cash holdings, equity investments, and bank accounts, net of credit card debt and overdue bills. Second is illiquid wealth: the sum of housing, other property, businesses, vehicles, and collectibles, net of mortgages and other debt. And third is superannuation wealth, which includes all superannuation accounts. Wealth is measured at the household level, and because of this we cluster standard errors by household.

Adverse shocks. We collect two important measures of adverse labor market shocks: unemployment and pandemic-induced negative income shocks. We record an individual as experiencing unemployment if they report unemployment in either 2020 or 2021.[17] Roughly 14 percent of our sample experienced unemployment during this period, much higher than usual. Second, we measure pandemic-related income shocks based on individuals' response to the question "Did the income you normally receive from paid employment increase or decrease because of the coronavirus? Or did it not change much?" which was asked to all individuals employed as of March 2020. In our sample, 17.6 percent of individuals reported a decrease in income due to the pandemic.

Demographics. We also collect a rich set of demographics for each individual. These include age, gender, education, marital status, number of children, and income (defined as financial year wages and salaries). All demographic variables are measured in 2020, the time when individuals were allowed early access to retirement wealth.

2.3 Analysis

In this section, we evaluate the relationship between self–control and early withdrawal. To set the stage, we first show how the probability of early withdrawal varies with self–control, as well as other observable characteristics or adverse shocks. We then explore how the marginal effect of self–control on early withdrawal changes when we control for these other factors in a regression specification in Section 2.3.3. We also explore the marginal effects of these other factors, and in Section 2.3.4 quantify their aggregate importance.

[17] We found that a more granular measure (time unemployed) did not substantially alter our results.

2.3.1 Descriptive Statistics

Overall, one in six (13.8%) working age individuals withdrew from their retirement account during the pandemic, in line with other papers that measure participation in alternative datasets (e.g. Bateman et al., 2023; Hamilton et al., 2024). This aggregate statistic masks meaningful heterogeneity in behaviour across various behavioural and situational dimensions, some of which are plotted in Figure 2.2. We see that early withdrawal is more common among people with greater self–control issues, supporting the idea that this is an important force driving people's interaction with the retirement system. But it is not the only force. Withdrawals are also more likely with shorter planning horizons, fewer assets (particularly liquid assets), and lower income.

The top–left panel of Figure 2.2a shows that the probability of early withdrawal increases markedly with self–control issues. In the bottom quintile of self–control issues, around 10 percent of individuals withdraw, while in the top quintile, roughly 17 percent withdraw, with a statistically significant difference. This is not the only behavioural trait that is related to withdrawal decisions. The top–right panel in Figure 2.2b shows that shorter planning horizons are also correlated with a higher probability of early withdrawal. And Appendix 2.A.2 shows a similar relationship between financial illiteracy and early withdrawal, as in Preston (2022).

Turning now to situational factors, we have measures of households' financial position and incomes, as well as adverse shocks suffered during the pandemic. Figures 2.2c and 2.2d show that the probability of early withdrawal is decreasing with wealth, whether liquid or illiquid, although the relationship with the former is slightly stronger, in line with the literature showing the primacy of liquid wealth in determining household spending (Kaplan and Violante, 2014). Figures 2.2e and 2.2f show how withdrawal varies with income and income shocks. The probability of early withdrawal is highest for individuals in the bottom two quintiles of the income distribution, adjusted for age, and declines gradually for higher income individuals. Further, the probability of withdrawal is much higher for those who have experienced adverse shocks, such as unemployment or pandemic related loss of income. Such shocks, while relatively rare, are strongly correlated with early withdrawal.[18] The importance of such shocks broadly lines up with the existing literature, such as Goda et al. (2022) who show in the U.S. setting that penalized withdrawals are more common among recent claimants of unemployment insurance and Choukhmane et al. (2024) who find that those who experience larger income declines are more likely to withdraw.

There are various other forms of heterogeneity which we report in Appendix 2.A.2.

[18] Such correlation is not a complete surprise, since eligibility was technically limited to individuals who had a reduction in work hours, loss of employment, or reduction in turnover. As noted in Section 2.2.1, however, eligibility was widespread, self-reported, and not verified. Past research shows that such rules did not have a binding effect, although it may still have deterred some withdrawals.

Figure 2.2: Probability of Early Withdrawal based on Situational and Behavioural Factors

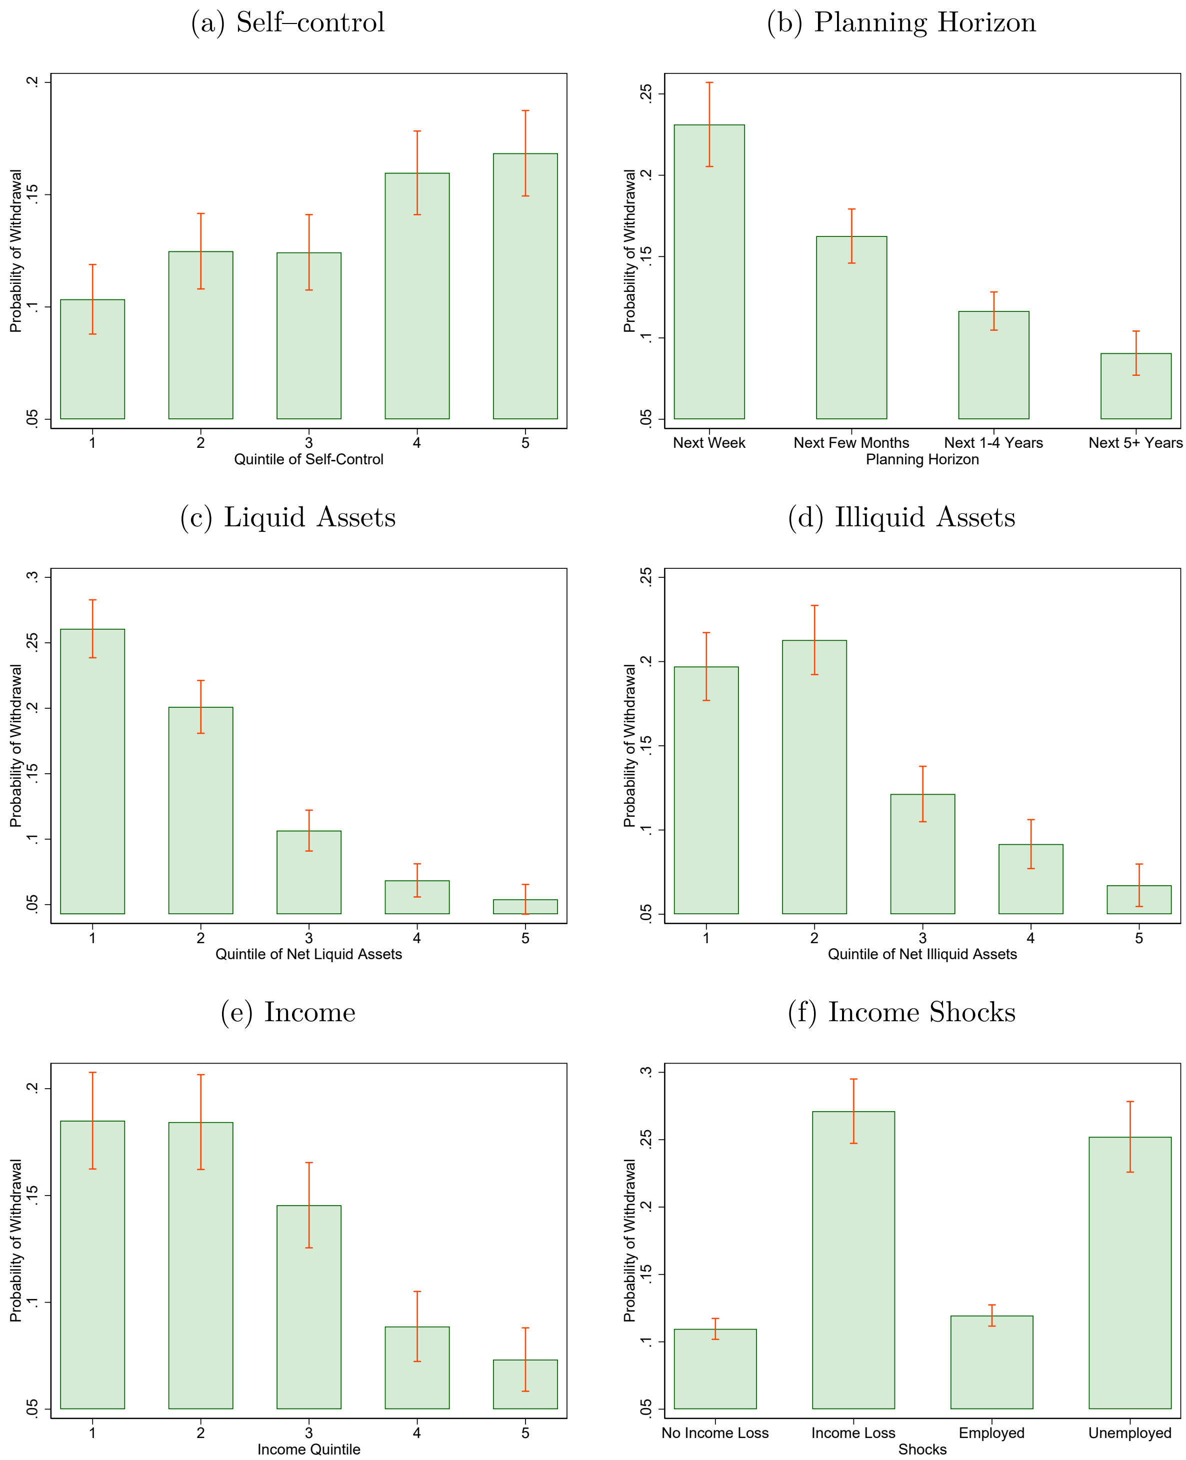

Note: Each figure shows the probability of early withdrawal based on a different observable characteristic. Income quintiles are computed within age group, since otherwise the income results are mostly driven by age effects.

Looking at age, individuals in their thirties are the most likely to withdraw, similar to the finding in Hamilton et al. (2024), which may reflect the fact that these individuals have already had a few years to accumulate wealth in their retirement account, but are still relatively constrained compared to their older counterparts.

That situational factors make people more likely to withdraw early is neither surprising nor necessarily concerning. In the design of retirement systems, one of the chief concerns is striking the right balance between commitment and flexibility. That people should use the flexibility they are afforded when they need it is a good outcome from the system *if these people are making good decisions.* Our results on self–control suggest that at least some people are not.

2.3.2 Empirical Specification

While the above section shows the unconditional probability of withdrawal for each of our main variables of interest, there are likely to be meaningful correlations between these variables. Issues with self–control, for instance, might be related to other psychological measures, such as planning horizon or financial literacy. Further, wealth may be greatly affected by psychological factors. Rather than assume ex-ante which of the potential determinants are most important, we include them all in a regression and test which have significant explanatory power.

To investigate the marginal relationship between self–control and early withdrawal we estimate the following logistic regression

$$\ln\left(\frac{p_i^{ew}}{1-p_i^{ew}}\right) = \beta_0 + \beta_1 \cdot \text{self–control}_i + \beta_2 \cdot \mathrm{X}_i + \epsilon_i \tag{2.1}$$

where p_i^{ew} is the probability of early withdrawal for individual i, self–control is the first principal component of the BSCS, and X_i is a vector containing a range of behavioural and situational control. These include measures of (i) demographics such as education, family size, age, sex, relationship status, log income, and a dummy for missing income, (ii) shocks like unemployment and loss of income income during COVID–19, (iii) the psychological factors financial literacy, planning horizon, and the big five personality traits, and (iv) wealth in liquid, illiquid, and Superannuation asset quartiles, as well as mortgage debt and mortgage payments. All estimates use responding person longitudinal weights, balanced between waves 18 to 21, and standard errors are clustered at the household level (at which wealth is measured).

2.3.3 Individual-Level Results

Table 2.4 reports the average marginal effects (AME) in a series of specifications, which build toward the full set of controls outlined in Equation (2.1). We find that while the relationship between self–control becomes weaker with the inclusion of other factors, it remains both statistically and economically significant. We also show that situational factors are significantly correlated with early withdrawal, and that self–control is the most important of the behavioural factors we consider.

Self–control and other psychological traits. Our main object of interest is the marginal effect of self–control issues, shown in the top row of Table 2.4. We find that self–control issues have an economically meaningful and significant relationship with early withdrawal. In specification (1), which controls just for demographics, we find that individuals with the highest level of self–control issues are 16 percentage points more likely to withdraw relative to those with no self–control issues, all else equal.[19] As we move through the specifications, adding controls for adverse shocks (2), behavioural factors (3), and wealth (4), we find that the AME of self–control diminishes but still remains economically meaningful. In specification (4), which includes all of our controls, we estimate an AME of 8.6 percentage points, which is similar to the effect of having 3 or more children. Based on this estimate, a one standard deviation increase in self–control issues (0.16) translates to a 1.4 percentage point increase in the probability of early withdrawal, while moving from the bottom to top quintile of self–control issues (i.e. from 0.17 to 0.63) translates to a 3.9 percentage point increase in the probability of early withdrawal, all else equal. This effect may be viewed as a lower bound if we believe that self–control issues also lead to lower wealth accumulation.[20]

A lack of self–control is not the only behavioural bias used to justify illiquidity in retirement accounts. We find that while these other psychological factors do also play a role in predicting early withdrawal, their estimated effects are weaker and less robust than that of self–control issues. Column 3 of Table 2.4 shows the marginal effects once we control for the full battery of psychological factors including financial literacy, planning horizon, and the big five personality traits. We find that financial literacy is correlated with a 4.2 percentage point reduction in the probability of withdrawal, although this relationship is nearly halved once we control for wealth in Column 4.[21] Further, we find that individuals with longer planning horizons have a lower probability of withdrawal. This effect disappears when we control for wealth, however, suggesting that the effect of shorter planning horizons on withdrawal is mediated mainly through wealth. Finally, we also evaluate the role of the big five personality traits (reported in Appendix 2.A.3). We find that greater emotional stability reduces the probability of withdrawal, but that

[19]Recall that our measure of self–control issues ranges between zero and one, so the AME tells us the implied impact, all else equal, of moving from no self–control issues to the maximum.

[20]If wealth is a mediator for self–control issues, then it is a bad control, absorbing variation that should rightly be attributed to self–control.

[21]Similarly, when predicting individual retirement wealth in the US, Goda et al. (2019) find that present bias and financial literacy are both important, with present-bias being the stronger predictor.

none of the other big five traits have an important effect. Overall, of all the psychological measures we consider, self–control is the most important determinant. This suggests that self–control issues may be the more important ones that retirement illiquidity guards against.

Adverse shocks. In line with the existing literature, we also find that adverse shocks play an important role in predicting early withdrawal. Unemployment and pandemic-related negative income shocks increase the probability of early withdrawal by 5.8 and 19.0 percentage points respectively. Our results indicate that negative income shocks are a stronger predictor of early withdrawal than self–control at the individual level. That said, it's important to note that the incidence of self–control issues is higher than either of these adverse shocks, a topic that we return to when evaluating the aggregate implications in Section 2.3.4.

Wealth. Finally, we also find that wealth is an important predictor of withdrawal. Individuals with low liquid assets are much more likely to tap into their retirement account, and liquid wealth plays a more important role than illiquid wealth in the spirit of Kaplan and Violante (2014). Of course, wealth is likely endogenous to personality traits such as self–control. Even when we control for wealth, however, we still see a significant and meaningful relationship between self–control and early withdrawal. This finding lends support to theories of present-bias contributing to high MPCs, above and beyond the effects of situationally low liquidity (Attanasio et al., 2024). In contrast, planning horizons cease to be important after controlling for wealth.

Our results complement recent analysis by Hamilton et al. (2024), who find that Australians who withdrew from their retirement accounts during COVID–19 spent around 40% of the money within the first two months, despite the modal withdrawal being the maximum $20,000 AUD. The authors state that this high MPC out of such a large amount is inconsistent with traditional models, where the MPC declines rapidly with shock size, and argue that early withdrawal is better rationalized by models with present-bias. We complement the above paper by evaluating the psychological determinants of early withdrawal using individual-level data on self–control issues, something the above authors can only infer. Our results provide clear evidence that self–control matters for early withdrawal. Further, our results show that heterogeneity in self–control is an important determinant of behaviour, lending support to recent models of retirement savings that explicitly model this form of heterogeneity (see e.g. Choukhmane and Palmer, 2024).

Our results also complement Goda et al. (2019), who predict retirement wealth in the US using a survey based measure of present-bias. The authors find that a one standard deviation increase in present-bias is associated with approximately $19,000 (10%) less retirement wealth at age 65. Two channels could cause this lower level of savings: fewer contributions or more withdrawals. While the setting in that paper differs from ours (namely, contributions are optional in the U.S. and withdrawals are generally permitted), our results support the idea that present–bias is likely to contribute to greater leakage from retirement accounts, absent regulations that preserve this wealth to later life.

2.3.4 Aggregate Implications

What do the individual level results in Table 2.4 imply in aggregate? To what extent is the aggregate propensity to withdraw from retirement accounts driven by psychological vs situational factors? To estimate the relative importance of these different factors, we need to think about how the composition of each varies across the population.

Figure 2.3 shows the overall rate of early withdrawals in the data (Bar 1) and under various counterfactuals (Bars 2 to 5) where we eliminate unemployment, pandemic related negative income shocks, self–control issues, and financial illiteracy. The counterfactuals are computed by setting each of these explanatory variables to zero in Equation (2.1), in effect turning off their direct influence on early withdrawal.[22] We then find and aggregate the fitted values to estimate the share of individuals that perform early withdrawal under each alternative assumption. This approach gives a lower bound of the effect of each factor because the traits are likely to have direct effects, as well as indirect ones mediated through income, wealth, or other controls; because we are only turning off the former, the aggregate importance we estimate does not include any indirect effects.

Overall, we find that self–control issues account for a similar share of withdrawals as negative income shocks. While negative income shocks have a larger AME, they are relatively concentrated; by contrast, self–control issues have a smaller AME but are much more widespread, and the net effect is about the same. More specifically, 17.6 percent of individuals in our sample were affected by pandemic-related negative income shocks. If we were to eliminate such shocks, the predicted early withdrawal rate would decline by 3.5 percentage points. In contrast, in our baseline sample, our measure of self–control has a mean value of 0.39. If we were to eliminate self–control issues by setting this value to zero for all individuals, while holding all other covariates fixed, we would predict the early withdrawal rate to decline by 3.3 percentage points.

Further, we see that self–control and negative income shocks both account for a larger share of early withdrawals than either unemployment or financial illiteracy. If we were to eliminate the direct effects of unemployment, we predict the early withdrawal rate to fall by 1.1 percentage points. This is despite the fact that the unemployment rate was unusually high during the pandemic, with 14 percent of our sample being unemployed during COVID–19. Similarly, if we were to eliminate financial illiteracy, the predicted

[22]We use the full specification, i.e. the estimates in column 4 of Table 2.4, for this exercise.

Figure 2.3: Implications for Aggregate Early Withdrawals

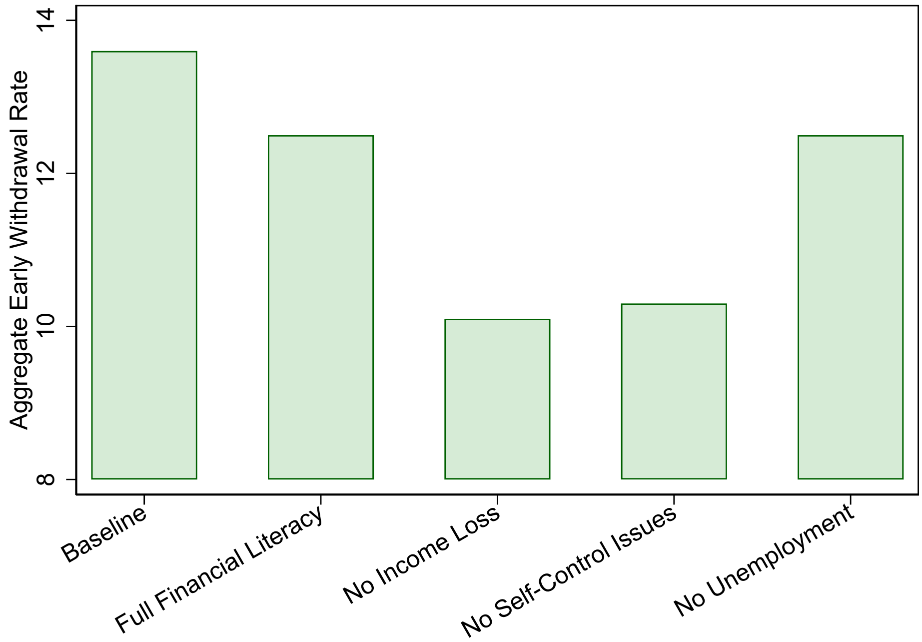

early withdrawal rate would only lower by 1.3 percentage points. The relative importance of self–control compared to financial literacy is consistent with Goda et al. (2019), who find that present bias is a more important predictor of retirement wealth than financial literacy. The above authors evaluate the behavioural determinants of wealth accumulation in an environment with few restrictions. Ours is the first study to measure how likely these restrictions are to be binding for people with self–control issues, driving their increased propensity to withdraw.

2.4 Conclusion

Our results highlight an important trade-off faced by policymakers: providing liquidity during economic distress while also ensuring that individuals with limited self–control can still build sufficient wealth for retirement. The recent trend of allowing households to withdraw from retirement accounts in times of aggregate economic distress amplifies the urgency of addressing this trade-off.[23]

In this paper, we examine the various factors influencing demand for liquidity, distinguishing between situational needs versus behavioral desires. Our results indicate that self–control issues do contribute substantially to early withdrawal. And while situational factors are generally a stronger predictor of early withdrawal at the individual level,

[23]While a full welfare analysis of this trade-off is outside the scope of the current paper, we return to this question in Schneider and Moran (2024a), where we develop a heterogeneous agent model to evaluate the distributional welfare implications of 'household liquidity policy' relative to traditional fiscal stimulus.

situational and behavioral factors are similarly important at the societal level.

Our result both supports the rationale for illiquidity in retirement systems to begin with—people do need help to put aside this wealth—and helps us evaluate the use of these funds for more short–term purposes. The Early Access scheme helped satisfy the need for short–term liquidity in the same way as traditional debt-financed fiscal stimulus payments. But it was attended by the fear that people who lack self–control will be more likely to tap into their retirement accounts and draw down their nest egg. We find that this was a well–founded fear.

2.A Appendix

2.A.1 Brief self–control Scale

Figure 2.4 shows the distribution of responses to the 13 items included in the Brief self–control Scale. Table 2.A.1 shows the estimated factor loadings for each of the 13 items. Overall, we see that the factor loadings go in the directions that we would expect based on the wording of each item. Further, we see that the estimated factor loadings, while relatively broad-based, are largest for items related to temptation and impulsive behavior.

Table 2.1: PCA Factor Loadings

	Question	Loading
a	I am good at resisting temptation	-0.2772
b	I have a hard time breaking bad habits	0.2916
c	I am lazy	0.2702
d	I say inappropriate things	0.2674
e	I do certain things that are bad for me, if they are fun	0.3185
f	I refuse things that are bad for me	-0.2331
g	I wish I had more self-discipline	0.3185
h	People would say I have iron self-discipline	-0.2100
i	Pleasure and fun sometimes keep me from getting work done	0.2656
j	I have trouble concentrating	0.2930
k	I can work effectively towards long-term goals	-0.2143
l	Sometimes I cannot stop myself from doing something, even if I know it is wrong	0.3247
m	I often act without thinking through all the alternatives	0.2907

Figure 2.4: Brief self–control scale questions and answers

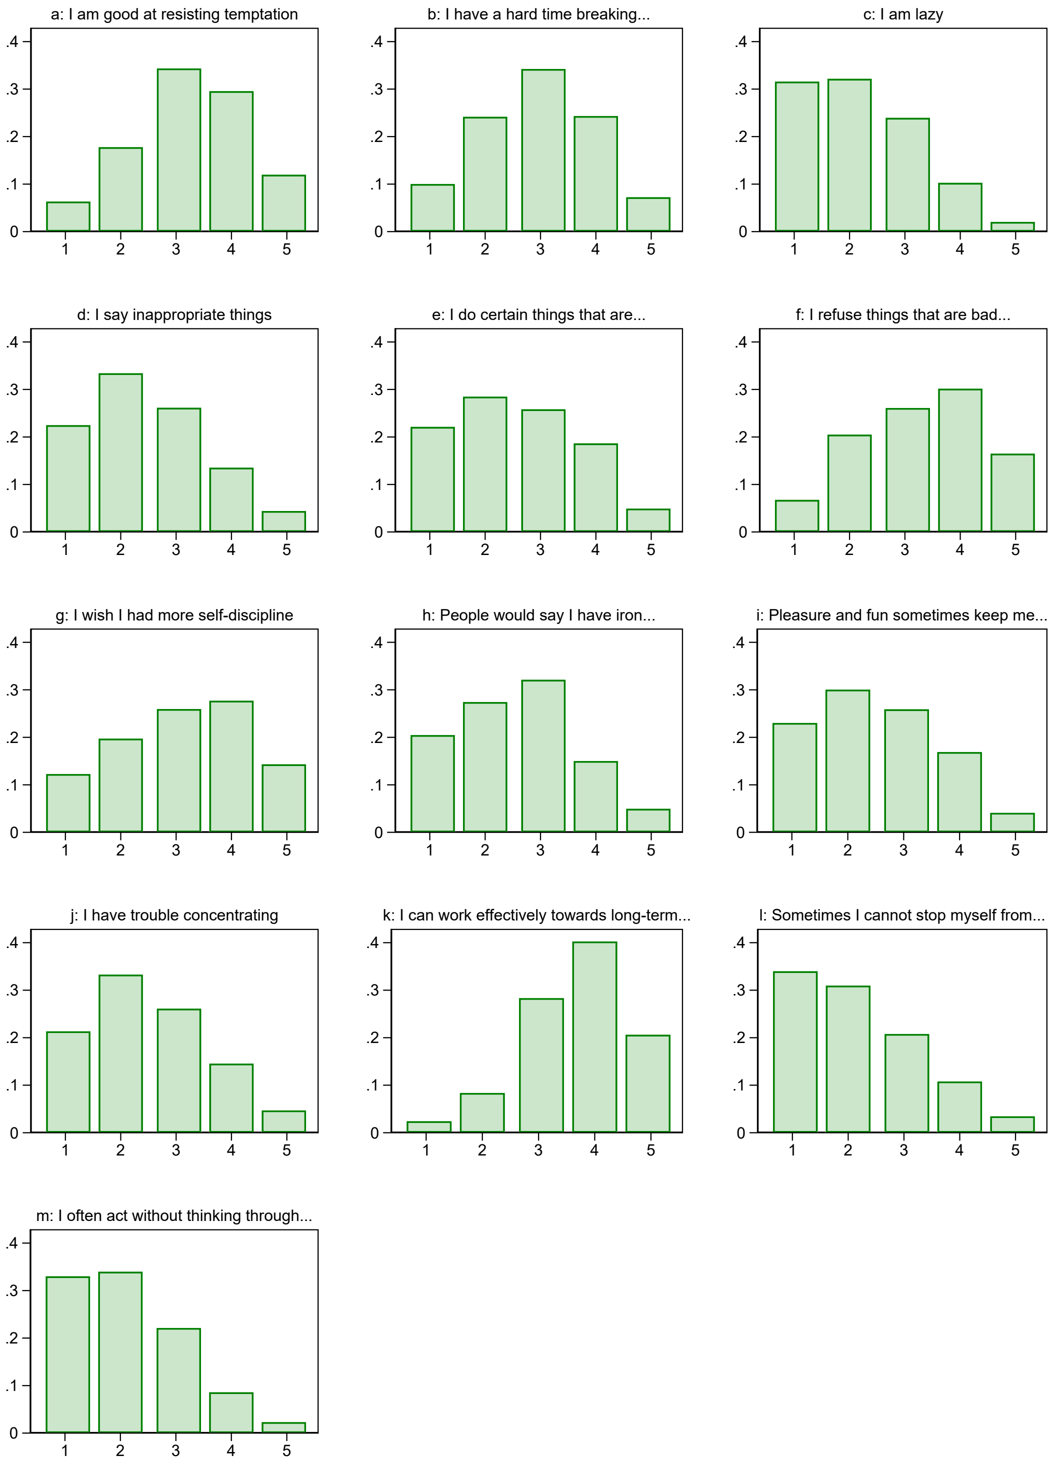

Note: Respondents are asked to rate how well each statement describes them, with responses ranging from 1 (“not at all”) to 5 (“very well”).

Of course, the Brief self–control Scale is not the only way to measure self–control issues. In general, there are two distinct approaches to measuring self–control, summarized by Cobb-Clark et al. (2022). The first relies upon responses to validated batteries of questions, following the canonical approach for measuring personality traits in the literature on personality psychology and economics (e.g. Almlund et al., 2011; Borghans et al., 2008; Heckman et al., 2021). The second approach is based on experimental economics, often measured on university students, which structurally estimates an individual's level of self–control based on their present-bias parameter β when estimating a $\beta - \delta$ model based on incentivized tasks (e.g. Andreoni et al., 2015; Andreoni and Sprenger, 2012; Augenblick et al., 2015; Augenblick and Rabin, 2019). In the present paper, we adopt the former approach using survey-based measurement. One benefit of this approach is that it can be embedded in large-scale household panel surveys that are nationally-representative and record a range of important economic outcomes. Both the Australian HILDA and German SOEP have recently incorporated such survey-based measurement of self–control into their large-scale panel surveys using the Brief self–control Scale.

2.A.2 Summary Statistics

Table 2.2 reports averages for all of our control variables in aggregate, and comparing withdrawing and non–withdrawing respondents.

Table 2.2: Summary Statistics

	No Withdrawal 6,538 (86.2%)	Early Withdrawal 1,047 (13.8%)	Total 7,585 (100.0%)
Unemployed			
0	87.9%	74.5%	86.0%
1	12.1%	25.5%	14.0%
incomeChangeCovidTrunc			
Not Decreased	85.1%	65.4%	82.4%
Income Loss from Covid	14.9%	34.6%	17.6%
Financial Literacy			
0	44.1%	56.5%	45.8%
1	55.9%	43.5%	54.2%
Self-Control Issues	0.390	0.423	0.394
bscs (standardized)	-0.003	0.213	0.027
Scores for component 1 (standardized)	-0.001	0.204	0.027
Age	39.210	38.305	39.085
educBins			
High School	67.5%	67.1%	67.4%
Postgraduate	7.3%	4.0%	6.8%
Undergraduate Bachelor	10.9%	7.1%	10.4%
Undergraduate Other	14.4%	21.8%	15.4%
male	0.450	0.489	0.455
Income	77,625.840	58,345.173	75,021.767
netLiquidWealth	79,730.663	29,160.334	72,750.158
netIlliquidWealth	608,484.840	309,511.091	567,215.820

Figure 2.5 shows the probability of early withdrawal conditional on various observable characteristics not shown in the main text. The probability of early withdrawal is highest for individuals in their thirties, which likely owes to the fact that these individuals have had time to accumulate wealth in their superannuation account, but still are early in their

life-cycle and therefore may be more exposed to other shocks. Turning towards education, we see that the probability of early withdrawal is lower for those who have completed a bachelors or postgraduate degree. The highest probability of early withdrawal is for those classified as "Undergrad Other," which reflects a number of undergraduate degrees including diplomas, certificates, and associate degrees, but not bachelor degrees.

Turning towards financial literacy, we see that the probability of early withdrawal is declining with the number of correct answers to the "big three" financial literacy questionnaire. Finally, turning towards wealth held in superannuation accounts, we see that the probability of early withdrawal is highest for those in the low-middle part of the distribution. Individuals in the bottom quintile have very little money to withdraw. Individuals in the top quintile are relatively wealthy and may have other forms of wealth that they can draw on before turning to retirement assets.

Motivated by these results, we include all of these variables as additional explanatory factors in our empirical specification discussed in Section 2.3.2.

2.A.3 Empirical Analysis

Table 2.3 reports the marginal effects for the full set of covariates included in our empirical specifications, including those omitted from Table 2.4 for the sake of expositional clarity.

Education initially appears to be an important predictor of withdrawal, although we find that most of this effect disappears once we control for wealth in specification (4). Further, although age appears strongly correlated with withdrawal in Figure 2.5, we find it is not an important predictor of withdrawal once we control for other factors.

We investigate the importance of the "Big Five" personality traits, which have been shown to be an important predictor of labor market outcomes (see e.g. Almlund et al., 2011; Borghans et al., 2008; Heckman et al., 2021; Todd and Zhang, 2020).[24] Overall, we find that most of these traits are unimportant when it comes to predicting early withdrawals. Of the big five traits, only emotional stability has a significant relationship, with greater emotional stability being correlated with reduced withdrawals. That said, none of the other traits have any significant relationship with withdrawal.

In specification (4), we also control for the presence of a mortgage and the size of mortgage payments, given the possibility that early withdrawal might be more likely for mortgagors. We find no evidence of such an effect conditional on our other controls.

[24] While the use of the Big Five personality traits in explaining economic outcomes is now well-established among economists, there is much less evidence on the role of self–control, perhaps because self–control has only recently been incorporated into large-scale household surveys.

Figure 2.5: Probability of Early Withdrawal

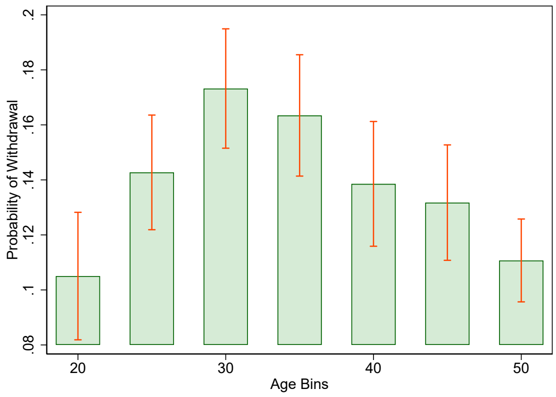

(a) Age

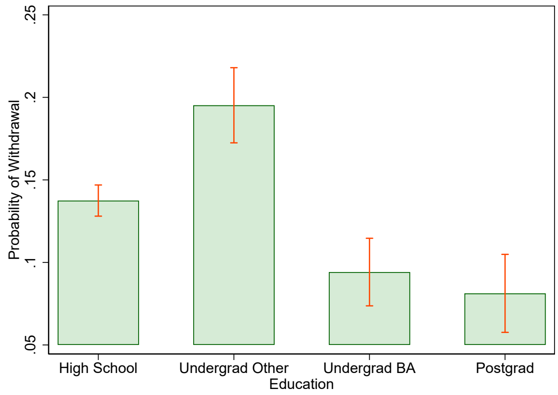

(b) Education

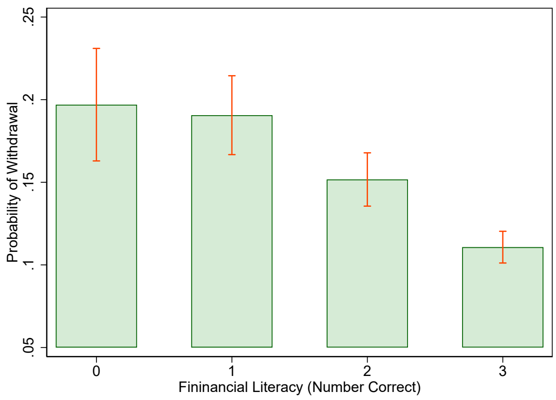

(c) Financial Literacy

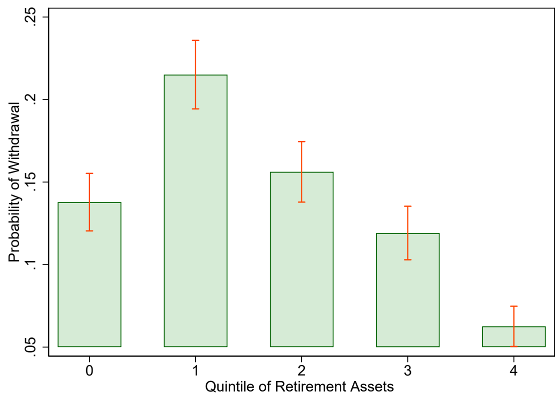

(d) Retirement Wealth

Note: Each figure shows the probability of early withdrawal based on a different observable characteristic. Retirement wealth is defined as the wealth held in one's superannuation account.

Table 2.3: Marginal Effects

	(1)	(2)	(3)	(4)
Self-Control Issues	0.16***	0.13***	0.11***	0.086**
	(0.036)	(0.034)	(0.036)	(0.034)
Log Income	−0.035***	−0.022***	−0.016***	−0.0097
	(0.007)	(0.006)	(0.006)	(0.006)
Postgraduate	−0.055***	−0.052***	−0.041**	−0.034
	(0.020)	(0.018)	(0.019)	(0.022)
Undergraduate Bachelor	−0.051***	−0.046***	−0.038**	−0.031*
	(0.017)	(0.017)	(0.017)	(0.017)
Undergraduate Other	0.035*	0.033*	0.031*	0.019
	(0.019)	(0.019)	(0.018)	(0.017)
Children: 1	0.070***	0.078***	0.069***	0.060**
	(0.024)	(0.025)	(0.024)	(0.023)
Children: 2	0.067***	0.067***	0.064***	0.055***
	(0.020)	(0.020)	(0.020)	(0.019)
Children: 3+	0.11***	0.11***	0.10***	0.085***
	(0.022)	(0.021)	(0.020)	(0.019)
agebins=30	0.016	0.015	0.019	−0.00085
	(0.027)	(0.027)	(0.025)	(0.025)
agebins=40	−0.021	−0.020	−0.0074	−0.012
	(0.024)	(0.025)	(0.023)	(0.023)
agebins=50	−0.043*	−0.042*	−0.025	−0.017
	(0.023)	(0.024)	(0.023)	(0.023)
male	0.033**	0.027**	0.033**	0.033***
	(0.013)	(0.013)	(0.013)	(0.013)
hasPartner	−0.028*	−0.021	−0.017	−0.0034
	(0.016)	(0.016)	(0.015)	(0.015)
incomeMissing	−0.41***	−0.27***	−0.21***	−0.15**
	(0.076)	(0.070)	(0.064)	(0.064)
Income Loss from Covid		0.19***	0.18***	0.19***
		(0.023)	(0.021)	(0.021)
Unemployed		0.068***	0.066***	0.058***
		(0.016)	(0.015)	(0.016)
Financial Literacy			−0.042***	−0.028**
			(0.013)	(0.012)
Planning Horizon: Few Months			−0.031*	−0.012
			(0.018)	(0.017)
Planning Horizon: 1-4 Years			−0.058***	−0.023
			(0.018)	(0.016)
Planning Horizon: 5+ Years			−0.065***	−0.023
			(0.020)	(0.019)
Big Five: Extroversion			0.016	0.016
			(0.014)	(0.014)
Big Five: Agreeableness			0.0082	0.014

			(0.018)	(0.018)
Big Five: Conscientiousness			0.014	0.021
			(0.017)	(0.017)
Big Five: Emotional stability			−0.033**	−0.033**
			(0.017)	(0.017)
Big Five: Openness			−0.0035	−0.015
			(0.015)	(0.015)
Liquid Assets: 2nd Quartile				−0.079***
				(0.017)
Liquid Assets: 3rd Quartile				−0.12***
				(0.017)
Liquid Assets: Top Quartile				−0.11***
				(0.022)
Illiquid Assets: 2nd Quartile				0.017
				(0.018)
Illiquid Assets: 3rd Quartile				−0.032*
				(0.020)
Illiquid Assets: Top Quartile				−0.049**
				(0.020)
Super Assets: 2nd Quartile				0.039**
				(0.018)
Super Assets: 3rd Quartile				0.023
				(0.019)
Super Assets: Top Quartile				−0.013
				(0.019)
mortgagePositive				0.12
				(0.090)
logMortgagePayment				−0.018
				(0.012)
Observations	7214	7214	7214	7214
Demographics	*Yes*	*Yes*	*Yes*	*Yes*
Adverse Shocks		*Yes*	*Yes*	*Yes*
Psych Controls			*Yes*	*Yes*
Wealth Controls				*Yes*

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.4: Marginal Effects

	(1)	(2)	(3)	(4)
Self-Control Issues	0.16^{***}	0.13^{***}	0.11^{***}	0.086^{**}
	(0.036)	(0.034)	(0.036)	(0.034)
Log Income	-0.035^{***}	-0.022^{***}	-0.016^{***}	-0.0097
	(0.007)	(0.006)	(0.006)	(0.006)
Children: 1	0.070^{***}	0.078^{***}	0.069^{***}	0.060^{**}
	(0.024)	(0.025)	(0.024)	(0.023)
Children: 2	0.067^{***}	0.067^{***}	0.064^{***}	0.055^{***}
	(0.020)	(0.020)	(0.020)	(0.019)
Children: 3+	0.11^{***}	0.11^{***}	0.10^{***}	0.085^{***}
	(0.022)	(0.021)	(0.020)	(0.019)
Income Loss from Covid		0.19^{***}	0.18^{***}	0.19^{***}
		(0.023)	(0.021)	(0.021)
Unemployed		0.068^{***}	0.066^{***}	0.058^{***}
		(0.016)	(0.015)	(0.016)
Financial Literacy			-0.042^{***}	-0.028^{**}
			(0.013)	(0.012)
Planning Horizon: Few Months			-0.031^{*}	-0.012
			(0.018)	(0.017)
Planning Horizon: 1-4 Years			-0.058^{***}	-0.023
			(0.018)	(0.016)
Planning Horizon: 5+ Years			-0.065^{***}	-0.023
			(0.020)	(0.019)
Liquid Assets: 2nd Quartile				-0.079^{***}
				(0.017)
Liquid Assets: 3rd Quartile				-0.12^{***}
				(0.017)
Liquid Assets: Top Quartile				-0.11^{***}
				(0.022)
Illiquid Assets: 2nd Quartile				0.017
				(0.018)
Illiquid Assets: 3rd Quartile				-0.032^{*}
				(0.020)
Illiquid Assets: Top Quartile				-0.049^{**}
				(0.020)
Observations	7214	7214	7214	7214
Demographics	*Yes*	*Yes*	*Yes*	*Yes*
Adverse Shocks		*Yes*	*Yes*	*Yes*
Psych Controls			*Yes*	*Yes*
Wealth Controls				*Yes*

Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Chapter 3

Yield To Temptation?

A comparison of present–biased preferences in continuous–time

Patrick Schneider

Abstract

I compare two ways of modelling present bias in dynamic consumption–-saving problems: quasi–hyperbolic discounting (QHD) and temptation preferences. I show how to implement temptation preferences in continuous time in a general way, where the temptation is to adopt an alternative discount function. Doing so allows for a direct comparison to results for the continuous–time limit of QHD ('instantaneous gratification' preferences (Harris and Laibson, 2013; Maxted, 2024)). I show that QHD consumers are behaviourally equivalent to naively tempted consumers, but that welfare is not equivalent. The models differ in various ways when consumers are sophisticated, and variation in the degree of sophistication offers opportunities for identification between the two models. Temptation preferences are more flexible than QHD, and so offer a promising path forward for incorporating present–bias into macro models.

3.1 Introduction

A growing body of evidence suggests behavioural biases are important for capturing consumption dynamics within macroeconomic models. Among these biases, a tendency to over–emphasise immediate gratification (present bias) is particularly prominent. Present bias is typically modelled with competing selves frameworks in one of two approaches.

The first and most prevalent approach is quasi-hyperbolic discounting, formalised by Laibson (1997) but stemming from earlier time–inconsistency models (Phelps and Pollak, 1968; Strotz, 1955). Here, consumers disproportionately value immediate rewards compared to future ones, causing short–term over consumption. While intuitive, quasi-hyperbolic discounting poses practical challenges, mainly due to the built in time inconsistency which complicates solutions. Perhaps due to these challenges, they have not been widely adopted in macro models, despite substantial evidence that present–bias is a relevant force (one notable exception is Maxted et al., 2024).

The second approach is temptation preferences, formalised by Gul and Pesendorfer (2001) and implemented in a consumption–saving setting in various papers (e.g. Attanasio et al., 2024; Gul and Pesendorfer, 2004; Kovacs et al., 2021). Here, consumers face costly temptations toward irresponsible consumption, and can mitigate these costs by giving in a little, leading again to short–term over consumption. This method is inherently more flexible and analytically convenient than the former, but has been used less in practice.

Despite their longstanding coexistence, there have been few comparisons of these two approaches, potentially because both are complex, but in different ways. Recent work has improved the tractability of quasi-hyperbolic discounting models by recasting them into continuous–time, termed 'Instantaneous Gratification' (IG) preferences (Harris and Laibson, 2013; Maxted, 2024). This work has demonstrated that, under some important assumptions, (1) IG consumers' policy functions are scale multiples of their rational equivalents, and (2) their value functions are positive affine transformations of their rational equivalents. These results make it much easier to work with and understand these preferences.

In this paper, I make the parallel extension of temptation preferences into continuous–time. The formulation I use is very general, and unifies various discrete–time temptation models by specifying that consumers are tempted to adopt an alternative (less responsible) discount functions. I use the new formulation to make a direct comparison of tempted and IG consumers' behaviour and welfare.

I show that IG preferences are behaviourally equivalent to *naive* temptation preferences, irrespective of the tempting discount function assumed.[1] Temptation preferences

[1] A similar equivalence is noted in O'Donoghue and Loewenstein (2004) and Fudenberg and Levine (2006), albeit in the much more limited setting of two–period models with restricted temptation preferences.

are therefore the more general option, because (a) they will deliver identical behaviour under naivete, and (b) relaxing the assumption of naivete leads to qualitatively different, and observable, behaviour that differs to consumers with IG preferences.

Temptation preferences are also more flexible than IG because the existence of solutions does not depend on functional form assumptions. All of the recent results in IG models are based on the assumption of constant relative risk aversion (CRRA) utility. While IG may be possible with other utility functions, these have not yet been developed, and it is not clear that they can be. As a result, IG preferences can only currently be used in a very narrow setting that excludes, for example, additive labour supply disutility. By contrast, temptation preferences can work with *any* utility function, and as a result are more easily implemented in a range of macroeconomic models.

While the IG model is nested within temptation in behavioural terms, they are not welfare equivalent. The welfare criterion in IG is the long–run preferences over biased behaviour. As Maxted (2024) shows, this will be a positive affine transformation of the rational agent's welfare function under his identifying assumptions. By contrast, temptation preferences work because the presence of temptation *harms* consumers—they are not biased in the sense that they're wrong about what they feel, just relative to what a rational agent would do. These preferences are tractable because they're time–consistent (sophisticated tempted consumers never regret their behaviour), but they achieve this consistency by distorting welfare instead. Because of this, whilst the two models can be behaviourally equivalent if consumers are naive, they will yield very different policy implications, and they will also yield different behaviour if consumers are sophisticated. This last fact presents opportunities for distinguishing between the two empirically.

The rest of the paper is structured as follows. In Section 3.2 I give an overview of the research on consumption's sensitivity to current income, in order to motivate the importance of behavioural preferences in macroeconomic models. In Section 3.3 I detail the basic consumption–saving model that will be the core of our analysis in the paper, and solve it for the rational benchmark case. In Section 3.4 I give an overview of IG preferences, the continuous time equivalent of the most widespread model of present–biased behaviour: quasi–hyperbolic discounting. And in Section 3.5 I introduce temptation preferences and show how to derive them in continuous–time in a general way. In Section 3.6 I use these results to compare behaviour and welfare under IG and temptation preferences, establishing naive behavioural equivalence and the different roles that sophistication plays. Section 3.7 concludes.

3.2 The case for present–bias in macro models

The aggregate marginal propensity to consume (MPC) is a central object in macro–policymaking, and an important target for models seeking to analyse stabilisation policy.

The MPC measures the sensitivity of consumption to current income and so governs how responsive consumption will be to fiscal transfers, one of the key stimulus tools available to governments.[2] The magnitude of this MPC has been a topic of contention, but recent evidence suggests that incorporating behavioural bias into theoretical models of consumption is necessary to match what we observe.

In the initial Keynesian analysis of government multipliers, the MPC was imposed by assuming consumption was a declining fraction of income, and the aggregate was an average of this across the population. In the 1950s through 1970s, the assumption was upset by theoretical contributions by Friedman (1957) and Modigliani and Brumberg (1954) which micro–founded consumption as the result of a forward looking problem by rational agents. This work led to the permanent income–life cycle hypothesis (PIH) that people will use savings to smooth through income fluctuations, and so current consumption should be a function of expected lifetime wealth. As a result, current consumption should be effectively independent from changes in current income, and more so if these changes were temporary, anticipated, or expected to be offset by future taxes with equivalent present value. The theoretical result implied that consumption should follow a random walk, supported empirically in Hall (1978). The PIH led to a strong policy prescription against Keynesian stimulus programs: efforts to juice the economy by redistributing from people saving resources to those who would spend them were futile, because *no–one* would spend them.

And yet they do. Starting in the 1980s, a vast literature has documented the 'excess sensitivity' of consumption to current income, in violation of the PIH. Time–series analyses showed that Hall (1978)'s result was upended once better econometric models were used (Campbell and Mankiw, 1989; Flavin, 1981). But the main evidence against the PIH has come from micro–data analysis of individuals' consumption sensitivity to changes in current income. The survey in Jappelli and Pistaferri (2010) gives a good overview of much of this literature, which tends to find current consumption *is* responsive to current income, whether it was predictable or not, but that there is substantial variation across the population in how responsive it is. This variation provides some clues as to what might be driving the apparently sub–optimal behaviour.

One of the most influential results is that individual MPCs are related to liquid resources, with less liquid people having higher MPCs (Baker et al., 2023; Fagereng et al., 2019; Gelman et al., 2022; Jappelli and Pistaferri, 2014; Karger and Rajan, 2020; Toczynski, 2023; Zeldes, 1989). This observation motivated the buffer–stock models of consumption and precautionary savings, starting in the 1990s (Aiyagari, 1994; Carroll et al., 2021; Carroll and Kimball, 1996; Deaton, 1989). In these models, households are subject to severe idiosyncratic risk that they can only partially insure due to borrowing

[2] MPCs can be measured in response to expected income at any horizon, termed inter–temporal MPCs in Auclert et al. (2024b). I will focus here on the MPC out of current income for brevity.

constraints. The fact of incomplete markets leads households that are close to their constraints to ration the resources they do have to avoid a sudden drop in consumption if they receive a bad shock. This rationing reflects short–term concern about hitting the constraint overwhelming long–term concerns about inter–temporal substitution or consumption smoothing i.e. their decisions are driven by a 'precautionary' motive. And in doing so it makes consumption much more sensitive to current income.

These models are the core of a recent wave of macro–models that seek to match aggregate MPCs (Auclert et al., 2024a; Bayer et al., 2019; Kaplan et al., 2018). They tend to feature households that are close to borrowing constraints for two reasons—some are poor, and some are wealthy but have most of their assets in illiquid form like housing or retirement accounts (Kaplan et al., 2014). These models rationalise investments in the illiquid asset with return differences: the illiquid asset pays such a good premium that rational households invest heavily in good times, only regretting it if they receive a bad shock or uncertainty increases. While these two–asset heterogeneous agent models have become the work–horse model for generating endogenous realistic MPCs, they are not a panacea. One issue, for example, is that the calibrated premium the illiquid account pays is implausibly high when compared to the return differences households actually face between genuinely liquid and illiquid options.[3]

The liquidity explanation of high MPCs has also been challenged by a range of recent papers identifying behaviour that is not consistent with precautionary saving. These include the observations that: measured MPCs can only be rationalised by implausibly low exponential discount factors (Gelman, 2022; Gerard and Naritomi, 2021; Hamilton et al., 2024; Laibson, 1997; Shapiro, 2005), people with substantial liquidity often have high MPCs anyway (Baugh et al., 2021; Boehm et al., 2025; Kueng, 2018; Olafsson and Pagel, 2018), consumption responds to current predictable income *losses* (which are insurable with saving) as well as gains (Baugh et al., 2021; Ganong and Noel, 2020; Gerard and Naritomi, 2021; Ni and Seol, 2014), people simultaneously carry high interest credit card debt and contribute to illiquid savings (Laibson et al., 2024), (lack of) liquidity is too predictable to be merely circumstantial (Parker, 2017), and measured MPCs are higher for credit that expires earlier, despite being fungible (Boehm et al., 2025).

This wealth of evidence suggests behavioural biases may be a necessary to addition to our toolbox if we want to explain household decisions. While a number of biases may be operating, much of the above behaviour can be explained with distortions to discounting captured by models of present–bias, the focus of the present paper.[4] In what follows

[3]Alternative theories for why people pile into illiquid but low return options appeal to other convenience–yields these assets might offer: owning housing allows one to escape the landlord risk, or over–withholding taxes might guard against the penalties that come with under–withholding when income is uncertain (Gelman et al., 2022). One such convenience yield I deal with in this paper is that illiquid assets don't tempt us to over–consume (Kaplan and Violante, 2022).

[4]Others in contention are distortions in expectation formation, and mental accounting or bounded

I detail the two main approaches to modelling present bias: quasi–hyperbolic discounting and temptation preferences. My contribution is to derive temptation preferences in continuous time in a general way for the first time, and use these results to do a direct comparison with Instantaneous Gratification preferences, the continuous time limit of quasi–hyperbolic discounting.

I am not the first to consider such a comparison. Krusell et al. (2010) build a model of temptation in which the tempting alternative is to adopt quasi–hyperbolic discounting such that the latter is the special case as the cost of temptation becomes infinite. My approach is much more general, and admits the Krusell et al. (2010) model as a special case. But by reframing the model in continuous time I am able to show that the temptation model nests quasi–hyperbolic discounting under *any* assumption about the tempting alternative if (a) the consumer is completely naive, and (b) the identifying assumptions in Maxted (2024) hold.

Toussaert (2018) tests for whether biased consumers suffer temptation or hyperbolic discounting in the lab. She identifies the difference by noting that temptation preferences lead people to seek commitment against strictly dominated options, whereas quasi–hyperbolic discounters will only seek commitment from options they would choose. She establishes that a substantial portion of the population is likely to experience temptation, providing evidence that this is the better model of present–bias to adopt. My work is theoretical, showing how these preference structures alter decisions in a consumption–saving framework. My behavioural equivalence result suggests that it will be difficult to tell the behaviour of these agents apart unless they are sophisticated. Identification schemes like the ones used in Toussaert (2018) may provide a useful way to approach identification in macro environments in future.

3.3 Consumption–saving model

The remainder of the paper compares different ways of modelling present bias in the same economic environment. This section details the environment—a continuous time consumption–saving problem where households have two assets, one a liquid transaction account in which borrowing is always possible but increasingly costly, and the other an illiquid saving account, and are exposed to idiosyncratic risk. This section starts by detailing the choices and constraints that households face in the general model environment, and I then solve the model for a household with rational preferences to serve as a benchmark.

rationality. Each of these, as well as temptation preferences, has in common that some part of decision marking is difficult, and the consumer must both make optimal decisions as well as exert effort over their decision making architecture—potentially limiting information gathering and processing, aiming for approximate optima, or reducing the pain from temptation.

3.3.1 The model

Households are differentiated by their liquid wealth b, illiquid wealth a, and labour productivity z.[5] Together, these three variables describe a household's 'state', and I sometimes group them together in the vector $x = (b, a, z)$. The model is in continuous time, and changes in these states over time are governed by the following processes:

$$\frac{\partial b}{\partial t} = r(b)b + (1-\xi)wz - c - d - \chi(d,a) \tag{3.1}$$

$$\frac{\partial a}{\partial t} = r_a a + \xi y + d \tag{3.2}$$

$$\frac{\partial \ln z}{\partial t} = -\theta_z \ln z + \sigma_z \frac{\partial W_t}{\partial t} \tag{3.3}$$

That is, additions to the liquid balance come from asset returns $r(b)b$ and labour income wz, net of the automatic illiquid contributions ξ, and reductions come from consumption c, voluntary contributions to the illiquid account (d), and the transaction costs these attract $\chi(d,a)$. The liquid asset return $r(b)$ is able to assume different values for different liquid balances. I assume it is fixed for positive balances at r_b and adopt the following assumption for negative balances.

Assumption 3.3.1 (*Soft borrowing constraint*)**.** The liquid return is continuous and weakly decreasing for negative balance $-r'(b) \geq 0\ \forall b < 0$, and becomes very large close to the natural borrowing limit.

Such a return function encodes a 'soft borrowing constraint' by deterring borrowing rather than banning it. Doing so ensures solutions are interior, which is helpful for some of the results below (Maxted, 2024). And the assumption is also realistic. Hard borrowing constraints are rare: extra credit is usually available as long as the borrower is willing to accept increasingly onerous terms e.g. credit card or payday loan rates are followed, in the limit, by risk of violence or criminal liability (Lee and Maxted, 2023).[6]

The illiquid balance evolves with asset returns $r_a a$, mandatory contributions ξwz, and voluntary contributions or withdrawals d. And log–labour productivity z follows an Ornstein–Uhlenbeck process around a stationary mean of $\bar{z} = 1$.

Within this environment, households choose a sequence of consumption and transfer plans between accounts $(\mathbf{c}, \mathbf{d})$ to optimise a general value function, which is composed of

[5]The model mirrors the setup in Kaplan et al. (2018) where households have access to a liquid and illiquid account, and face idiosyncratic risk. Relative to the original it is simplified in that the labour income process is an endowment.

[6]Despite being available, it's likely that these extra sources of credit come with greater costs, potentially both fixed and variable. Here I assume that these costs are all encoded in the interest rate, and this is important. If the costs were instead incurred prior to extracting the credit then they may reinstate hard constraints.

the expected discounted integral of the flow of utility

$$v_t(x) = \max_{\mathbf{c},\mathbf{d}} \mathbb{E}_t \int_0^{\infty} D(s) u(c_s(x_{t+s})) ds$$

Where the policy at a particular point in time $c_s(x)$ is a subset of the sequence of plans $\mathbf{c}$. In the stationary case, policies are the same at any point in time and $(\mathbf{c}, \mathbf{d}) = (c(x), d(x))$. $D(s)$ is the discount function applied to time s into the future. In the next section, I solve for the rational benchmark case.

3.3.2 Rational benchmark

In the rational benchmark, households have CRRA preferences and an exponential discount function $D(s) = e^{-\rho s}$. Hence stationary their value is

$$v(x_t) = \mathbb{E}_t \int_0^{\infty} e^{-\rho s} \left(\frac{c(x_{t+s})^{1-\sigma}}{1-\sigma} \right) ds$$

And the solution to their problem is defined by the Hamilton–Jacobi–Bellman (HJB) equation[7], state transition equations (3.1–3.3), and first order conditions

$$\rho v(x) = u(c(x)) + \partial_b v(x) \cdot \dot{b}(x) + \partial_a v(x) \cdot \dot{a}(x) + \mathcal{A}[v](x) \tag{3.4}$$

$$c(x)^{-\sigma} = \partial_b v(x) \tag{3.5}$$

$$\partial_a v(x) = \partial_b v(x)(1 + \chi_d(d(x), a)) \tag{3.6}$$

Where $\mathcal{A}[\cdot]$ is the infinitesimal generator of the income process. Note that under Assumption 3.3.1 the solution is interior, and so we do not need to account for the liquid state–constraint binding. In the following sections I detail how quasi–hyperbolic discounting and temptation models distort the value, and how naivete alters households' expectations of these distortions going forward.

3.4 Quasi–hyperbolic discounting preferences

The most widespread model of present–biased behaviour is the quasi–hyperbolic discounting model. This dates back to the time–inconsistency models of Strotz (1955) and Phelps and Pollak (1968), which grapple with problems where individuals or societies place disproportionate weight on the present when making dynamic choices. These were formalised for consumption–savings problems in Laibson (1997), in which the problem is boiled down to a distorted present–value function, which places an extra discount on all discrete entries except the current one. This so called $\beta - \delta$ model discounts a flow of utilities coming from a consumption plan $\{c_0, c_1, ..., c_T\}$ as follows

$$u(c_0) + \beta \cdot \sum_{t=1}^{T} \delta^t u(c_t)$$

Where $\delta < 1$ is a standard exponential discount factor between periods, and $\beta \leq 1$ captures the degree of present–bias. The starting point for these models was to formalise time–inconsistency, but it has been shown that this phenomenon leads to over–consumption and, potentially, commitment–seeking behaviour where these agents anticipate inconsistency and attempt to bind their future selves' hands.

3.4.1 Quasi–hyperbolic discounting in continuous–time

The model can be extended to its continuous–time limit (where the future is the next instant) and doing so yields to greater tractability, as well as clearer welfare results. In this limit, quasi–hyperbolic discounting is referred to as 'Instantaneous Gratification' (IG) preferences. They were first introduced and explored at length in Harris and Laibson (2013), with complementary analysis in Maxted (2024). What follows in this section is an overview of the main results from these papers.

Instantaneous Gratification (IG) preferences capture quasi–hyperbolic discounting in the continuous time limit, where there is a time–dependent discount function (Harris and Laibson, 2013)

$$D^{\beta}(s) = \begin{cases} 1 & \text{if } s = 0 \\ \beta e^{-\rho s)} & \text{if } s > 0 \end{cases}$$

That is, all time beyond the current instant is discounted at a rate of $\beta < 1$ on top of the standard exponential discounting. This discounting of the future makes the preferences time–inconsistent, as in the discrete–time formulation of the model. Assuming sophistication, i.e. the consumer correctly anticipates future bias, the IG consumer's current value $w^{\beta}(x)$ in the stationary solution is

$$\begin{aligned} w^{\beta}(x_t) &= \max_{c^{\beta}(x), d^{\beta}(x)} \mathbb{E}_t \int_0^{\infty} D^{\beta}(s) u(c^{\beta}(x_{t+s})) ds \\ &= \lim_{\Delta \to 0} \max_{c^{\beta}, d^{\beta}} u(c^{\beta})\Delta + \beta e^{-\rho\Delta}\, \mathbb{E}_t\, v^{\beta}(x_{t+\Delta}(c^{\beta}, d^{\beta})) \end{aligned} \tag{3.7}$$

Where the maximisation is subject to the state transition equations (3.1–3.3), and $v^\beta(x)$ is the long–run value the biased consumer places on their *future* expected stream of consumption, i.e. it anticipates but does not adopt future bias. Converting this into an HJB, by taking a second–order approximation around $\Delta = 0$, yields

$$\beta\rho v^\beta(x) = \max_{c,d} u(c) + \beta\left[\partial_b v^\beta(x)\cdot \dot{b} + \partial_a v^\beta(x)\cdot \dot{a} + \mathcal{A}[v^\beta](x)\right] - \frac{w(x_t) - \beta v^\beta(x)}{dt} \quad (3.8)$$

A solution requires that the current value satisfy $w(x) = \beta v^\beta(x)$, otherwise the final term in Equation 3.8 will explode. We can see why intuitively in Equation 3.7, where the first term goes to zero in the continuous–time limit i.e. if your value depends on current pleasure and discounted anticipated future pleasure, then the closer the future is, the more the latter term will dominate. The IG consumer's optimal choices are given by the FOC

$$u'(c^\beta(x)) = \beta\partial_b v^\beta(x) \quad (3.9)$$

$$\partial_a v^\beta(x) = \partial_b v^\beta(x)(1 + \chi_d(d^\beta, a)) \quad (3.10)$$

Based on the FOC, we can find the policy functions if we know $v^\beta(x)$. But this is where things get difficult for IG preferences. Due to the time–inconsistency, we cannot use Equation (3.8) to solve for $v^\beta(x)$ as we usually would.

The $\hat{u}$ construction Harris and Laibson (2013) show that while the value for this household does exist, and is unique, it cannot be estimated using any of the above equations. Instead, they show that one can recover $v^\beta(x)$ by solving a different problem for a rational agent (i.e. an exponential discounter) with a distorted felicity function $\hat{u}(\cdot)$. Assuming $u(\cdot)$ is CRRA over consumption, there is a particular $\hat{u}(\cdot)$ function that delivers a solution equal to the long–run value of the biased household i.e. $\hat{v}(x) = v^\beta(x)$. This value can be solved using standard methods because this agent's problem is time–consistent Achdou et al. (2022); Harris and Laibson (2013). And with this value in hand, we can use the IG agent's FOC (Equations 3.9 and 3.10) to find their policy functions.

Maxted (2024) builds on this result by showing that under the assumption of a soft borrowing constraint (Assumption 3.3.1), the $\hat{u}$ function simplifies to a positive affine transformation of the CRRA utility function.

$$\hat{u}(c) = \frac{\psi}{\beta}\frac{\left(\frac{1}{\psi}c\right)^{1-\sigma}}{1-\sigma} \quad \text{where} \quad \psi = \frac{\sigma - (1-\beta)}{\sigma}$$

This leads to two results (1) the $\hat{u}$ agent's value, and therefore $v^\beta(x)$, is a positive affine transformation of the rational value $v(x)$, and (2) the $\hat{u}$ household's behaviour is identical to a rational agent with a standard utility function's behaviour (because affine transformations don't upset ordinal rankings, choices will not differ).

Relation to rational benchmark These results allow us to express the IG agent's optimal behaviour as analytical functions of the rational benchmark derived earlier, obviating the need for the $\hat{u}(\cdot)$ gymnastics at all. To see this for the consumption function, for example, note from combining Equations 3.9 and 3.5 that the IG consumption policy satisfies

$$u'(c^\beta) = \beta \frac{\partial_b v^\beta}{\partial_b v} u'(c)$$

Given we know the IG agent's long–run value is a positive affine transformation of the rational one, the ratio $\frac{\partial_b v^\beta}{\partial_b v}$ is determined by the scaling factor in this transformation. In this case the scaling factor is $\frac{\partial_b v^\beta}{\partial_b v} = \frac{\psi^\sigma}{\beta} > 1$, and the relationship between the consumption policies is as follows

$$c^\beta(x) = \left(\frac{\sigma - (1 - \beta)}{\sigma} \right)^{-1} c(x)$$

Hence, as long as $\sigma > (1 - \beta)$ and $\sigma > 1$, there is a well–behaved relationship with the IG consumption increasing, relative to the rational benchmark, as the present bias becomes more severe $\downarrow \beta$, or as the elasticity of inter–temporal substitution $(1/\sigma)$ increases.

Similar logic leads to an even cleaner result for the transfer policy function $d^\beta(x)$, where the scaling factor appears on both sides of Equation 3.10, and cancels out. As a result, the IG agent makes the same asset allocation decision as the rational benchmark

$$d^\beta(x) = d(x)$$

As discussed in Maxted (2024), this result makes intuitive sense because asset allocation decisions compare different impacts on the future. Quasi–hyperbolic discounting distorts evaluations of decisions the span the future and the present, so it is natural that decisions not involving the present are not distorted.

This result implies that *IG agents do not pursue commitment devices*, a strong result relative to the existing literature, but one that has some empirical support (Laibson, 2015). The lack of commitment seeking rests on Assumption 3.3.1, which implies no amount of squirreling away funds in an illiquid account can actually bind the biased agent's consumption in future, because there are always other sources of credit to tap. As a result, they don't seek out such devices.

3.4.2 Extension to naivete

All of the discussion so far has presumed that households are sophisticated. That is they (correctly) anticipate that their bias will continue in the future. We can relax this assumption by introducing naivete, where households (incorrectly) expect their future selves to be subject to a degree of bias β^E that is less severe than their current level.

All of the above results hold with slight alterations to make a distinction between effects coming from current versus anticipated bias. I use the tilde notation to denote naivete. First, the current value is now

$$\tilde{w}^\beta(x_t) = \lim_{\Delta \to 0} \max_{\{\tilde{c}^\beta, \tilde{d}^\beta\}} u(\tilde{c}^\beta)\Delta + \beta e^{-\rho\Delta}\, \mathbb{E}_t\, v^{\beta^E}(x_{t+\Delta}(\tilde{c}^\beta, \tilde{d}^\beta)) \tag{3.11}$$

Where $v^{\beta^E}(x) = \int_0^\infty e^{-\rho s} u\left(c^{\beta^E}(x_s)\right) ds$ is the long–run value placed on the consumption plan of an agent with expected future bias β^E. The transfer policy $\tilde{d}^\beta(x)$ is unaffected by bias, real or imagined, and so is unchanged

$$\tilde{d}^\beta(x) = d(x) \tag{3.12}$$

and the consumption policy is

$$\tilde{c}^\beta(x) = \left(\frac{\beta}{\beta^E}\right)^{-\frac{1}{\sigma}} \left(\frac{\sigma - (1-\beta^E)}{\sigma}\right)^{-1} c(x) \tag{3.13}$$

This policy function embeds two interesting extremes. Under sophistication $\beta^E = \beta$ the function reduces to the sophisticated case derived earlier. Under complete naivete $\beta^E = 1$ the IG agent's consumption scales the rational benchmark by $\beta^{-\frac{1}{\sigma}}$.

We can use this formula to quantify the impact of naivete on consumption behaviour. Consider a stylised example where $(\sigma, \beta) = (2, 0.7)$, which are within the range of empirical estimates for the parameters. Using Equation 3.13 we can see that a sophisticated agent will consume 18% more than the rational benchmark, compared to 20% more for a complete naif. Naivete therefore increases consumption. The effect is quite small in this example, but it is larger the more intense the present bias.[8]

As all of these cases are scale multiples of the rational benchmark policy, they are

[8] The sign of this effect actually depends on parameters. Per Maxted (2024) "The intuition for this result is that naivete introduces two offsetting effects. On the one hand, the naif is more willing to save because the naif trusts their future selves. On the other hand, the naif is less willing to save because the naif believes that future selves will save enough on their own. The former effect dominates when the agent is relatively more willing to substitute intertemporally ($\sigma < 1$), and vice versa." (footnote 34, p. 28) Empirical estimates of σ place it well above 1, and so we can be confident that naivete should lead to lower consumption in the present.

also isomorphic to each other. Indeed if Assumption 3.3.1 holds, then any degree of sophistication is equivalent to a completely naive agent with a reduced discount.

Lemma 3.4.1 (*Inverse of Maxted (2024) Corollary 5*)**.** An IG agent with any degree of bias or sophistication parameters

$$(\beta, \beta^E) \in [0, 1] \times [\beta, 1]$$

consumes identically to a completely naive IG agent with a bias parameter β^* that satisfies

$$\beta^* = \frac{\beta}{\beta^E}\left(\frac{\sigma - (1 - \beta^E)}{\sigma}\right)^{\sigma}$$

Proof. This results from setting the scaling factor in Equation 3.13 equal to $(\beta^*)^{-\frac{1}{\sigma}}$ □

Lemma 3.4.1 implies that sophistication cannot be identified in consumption data alone under these preferences.

Comment on inflexibility These results mean the IG agent's policy functions can be found in a simple two–step process of (1) solving the rational agent's problem, and (2) backing out the IG agent's policies. This avoids the need to actually use the $\hat{u}$ function itself in the process of solving the model, but the results do rely on this construction.

As a result, dealing with IG preferences necessitates first finding the appropriate $\hat{u}$ function. Harris and Laibson (2013) backward–engineer it to deliver the value function equivalence for CRRA utility over consumption. But a different $\hat{u}$ is needed for any potential $u(\cdot)$. Other $\hat{u}$ constructions have been shown to exist. For example Maxted (2024) derives the new $\hat{u}$ function when the utility function is a CRRA wrapper around a Cobb–Douglas bundle of consumption and housing durable, and shows that the extension is relatively straightforward.

It's currently unclear, however, if these functions exist for more complicated problems—for additive disutility from labour, for example. For the moment, working with IG preferences can be cumbersome unless we either (a) stay within the bounds of CRRA utility wrappers, or (b) assume complete naivete (noting that the isomorphism between naive and sophisticated behaviour described above depends on the assumption of both CRRA utility and non–binding constraints). An alternative approach is to instead use *time–consistent* preferences that also induce present biased behaviour.

3.5 Temptation preferences

Temptation preferences are an alternative way of modelling present bias that are time–consistent. In this model, we imagine people derive utility from their choices, but this satisfaction is coloured by what their alternatives were. It is as though while making their choices, a demon stood on their shoulder whispering to them to be less responsible. While the demon doesn't succeed in driving behaviour, the act of resistance is costly. The cost depends on the gap in temptation utility between what you imagined and what you did, and so in the moment people can alleviate the cost by giving in a little.[9] This 'giving in' can mean over–consuming if the problem is a consumption–saving one. And if people recognise this pattern of behaviour in themselves (i.e. if they're sophisticated), they may take actions to limit what they can be tempted by in the future (i.e. seeking commitment devices), reducing the power these demons have to sway decisions in future.

Whilst temptation is at least as old as humanity, we owe the preference formulation to Gul and Pesendorfer (2001). In temptation preferences, people have rankings over choice sets as well as options within them and, if particular axioms hold, then temptation preferences in their most general form can be rationalised by the following value function

$$V(\mathcal{K}) = \max_{x \in \mathcal{K}} h(x) - \left[\max_{y \in \mathcal{K}} k(y) - k(x) \right] \tag{3.14}$$

Where the choices (x, y) are made from the consumer's choice set $\mathcal{K}$, x is the actual action taken and y is the tempting alternative the demon suggests. $h(x)$ represents the utility the path chosen generates, and the term in square brackets represents the cost of temptation: the difference between the temptation utility $k(\cdot)$ experienced by choosing the 'most tempting alternative' y compared to the actual choice x, evaluated with the tempting discount function $k(\cdot)$. With this structure, a consumer may prefer the limited choice set $\mathcal{S} \subset \mathcal{K}$ because the latter includes options that, while not chosen, will cause harm by giving the demon some ammunition.[10]

The key to using these preferences is being specific about what exactly the consumer is most tempted by i.e. the form of the temptation utility $k(\cdot)$. In this paper we are concerned with consumption–savings problems in which the choice is over streams of consumption flows and the choice–set is the budget–set. As I will show, it's natural for the temptation to be an alternative discount function.

Discrete–time example Before getting to the continuous–time derivation, we can illustrate how temptation works in a consumption–savings setting using the linear ver-

[9] Parents of small children may for example resist buying toys in the gift shop at museum exits by promising a smaller, cheaper toy from the supermarket on the way home.

[10] This feature of preferences is called 'set–betweeness', and is one of the new axioms introduced by Gul and Pesendorfer (2001) that permits a preference for commitment. The standard example of this demand for commitment (preferring limited choice–sets) in practice is the harm felt by recovering addicts of being around drug–users *even if* they manage not to relapse, because the willpower required to resist is itself costly. Standard preferences do not feature this because more options always weakly improve welfare.

sion of Dynamic Self–Control (DSC) preferences (Attanasio et al., 2024; Fudenberg and Levine, 2006; Gul and Pesendorfer, 2004; Kovacs et al., 2021; O'Donoghue and Loewenstein, 2004).

In this situation consumers are tempted to discount the future entirely, and their temptation utility based on a consumption stream $\{c_t\}_{t=0}^{\infty}$ therefore ignores all future consumption $k(c) = \lambda u(c)$. In a simple one–asset setup with interest rate r, stochastic income y and discount factor δ, the Bellman equation is

$$\begin{aligned} V(x) &= \max_c u(c) - \lambda \left[u((1+r)b + y) - u(c) \right] + \delta \, \mathbb{E} \left[V(x') \right] \\ b' &= (1+r)(b-c) + y' \end{aligned}$$

Here the most–tempting feasible plan is to consume all available cash on hand $c^* = (1+r)b+y$, presuming no borrowing is allowed. When $\lambda = 0$ these preferences embed the rational case. When $\lambda > 0$ the closer actual consumption is to exhausting resources the less the temptation cost is felt, and so a rational consumer will choose greater consumption than their un–tempted equivalent to relieve this pressure. We can rearrange the Bellman equation slightly to show two separate forces introduced by temptation

$$V(x) = \max_c \; (1+\lambda)u(c) - \lambda u((1+r)b + y) + \delta \, \mathbb{E} \left[V(x') \right] \tag{3.15}$$

Here we can see that temptation ($\lambda > 0$) introduces two effects—it boosts the value of present–day consumption by $1 + \lambda$, and it reduces felicity by an amount that depends on cash–in–hand. This latter force means the disutility from the most–tempting alternative introduces a force akin to negative money in the utility function models (Poterba and Rotemberg, 1986) i.e. it makes liquid assets relatively less attractive for a given financial return by imposing a sort of inconvenience yield. In cases where there are alternative storage options that are not tempting, this force will lead consumers to store less wealth in the transacting asset and more in those alternatives. That is, they will seek out commitment devices.

3.5.1 Temptation in continuous–time

In this section I extend the temptation model to continuous–time in a general way, based on the consumption–saving problem introduced in Section 3.3 and where I make two assumptions

Assumption 3.5.1. The costs of temptation comparisons is linear.

Assumption 3.5.2. Consumers are tempted by an alternative discount function

$$D(s) \neq e^{-\rho s}$$

The first is a common simplifying assumption (e.g. in Attanasio et al., 2024; Fudenberg and Levine, 2006; Kaplan and Violante, 2022; Kovacs et al., 2021; Krusell et al., 2010) and the second allows us to embed various forms of temptation preferences in a general framework.

To start, define the inputs into the temptation value (Equation 3.14). The choice set contains all the feasible consumption and deposit plans for each point in the state–space. The standard value placed on a consumption and deposit plan $(\mathbf{c}, \mathbf{d})$ that respects state–transitions (where bold text represents a stream of consumption plans into the future) is

$$h(\mathbf{c}, \mathbf{d}, x) = \mathbb{E} \int_0^\infty e^{-\rho s} u(c(x_{t+s}, s)) ds$$

and the temptation value place on the same plan under our two assumptions is

$$k(\mathbf{c}, \mathbf{d}, x) = \lambda \, \mathbb{E} \int_0^\infty D(s) u(c(x_{t+s}, s)) ds$$

The temptation value differs from the standard in two ways: consumption flows may be discounted differently to standard $D(s) \neq e^{-\rho s}$, and they receive a linear boost $\lambda \geq 0$. Combining these, we can define the value function for a general tempted consumer as

$$v^\lambda(x) = \max_{\mathbf{c}, \mathbf{d}, \hat{\mathbf{c}}, , \hat{\mathbf{d}}} \mathbb{E} \int_0^\infty e^{-\rho s} u(c(x_{t+s}, s)) ds - \lambda \, \mathbb{E} \int_0^\infty D(s) \left[u(\hat{c}(x_{t+s}, s)) - u(c(x_{t+s}, s)) \right] ds$$

This setup embeds various different approaches to temptation preferences that differ in the discount function the consumer finds tempting.

Rational benchmark First note that these preferences embed rationality as a special case. The linear boost λ determines how intensely temptation is felt, and can be used to switch it off entirely: in the case where $\lambda = 0$, the above resolves to the rational value.

Alternatively, if consumers are tempted by their options in the same way as they actually evaluate them, i.e. they set the tempting discount function to the exponential one $D(s) = e^{-\rho s}$, then the above resolves to the rational value as well because the optimal $c(s) = \hat{c}(s)$ and so there is no difference between the tempting plan and the actual one.

For temptation to have any effect, there must therefore be some alternative way of viewing options $k(c) \neq \lambda \cdot h(c)$ *and* the cost of resisting temptation must not be zero $\lambda > 0$.

Greater discounting In Kaplan and Violante (2022) people are tempted to adopt a much higher discount rate. We can encode that here with $\lambda > 0$ and the exponential discount function with $\hat{\rho} > \rho$

$$D(s) = e^{-\hat{\rho}s}$$

These preferences yield the following HJB equation, derived in Appendix 3.A.2

$$\begin{aligned}\rho v^{\lambda}(x) = \max_{c^{\lambda}, d^{\lambda}} \; &(1+\lambda)u(c^{\lambda}) - \lambda u(\hat{c}(x)) + (\hat{\rho} - \rho)\left(\hat{k}(x) - k(\mathbf{c}^{\lambda}, \mathbf{d}^{\lambda}, x)\right) \\ &+ \partial_b v^{\lambda}(x) \cdot \dot{b} + \partial_a v^{\lambda}(x) \cdot \dot{a} + \mathcal{A}\left[v^{\lambda}\right](x)\end{aligned}$$

Where $\hat{k}(x) = k(\hat{\mathbf{c}}, \hat{\mathbf{d}}, x)$ is the temptation value of the tempting consumption and deposit plan, and the latter are solved in a separate problem with the discount rate $\hat{\rho}$. Solving the general problem amounts to treating it as a time–consistent problem with the distorted felicity function

$$(1+\lambda)u(c(x)) - \lambda u(\hat{c}) + (\hat{\rho} - \rho)\left(\hat{k}(x) - k(\mathbf{c}^{\lambda}, \mathbf{d}^{\lambda}, x)\right)$$

There are two distortions here. The first is from currently felt temptation utility, and the second is from the anticipated future temptation and the clash of discount rates used to evaluate this future. In a stationary solution, $k(\mathbf{c}^{\lambda}, \mathbf{d}^{\lambda}, x)$ is simply a function of the chosen policies. In a dynamic solution, this object will depend on the sequence of consumption plans going forward. Similarly, $\hat{k}(x)$ depends on the solution to the tempting problem, which will depend on the anticipated sequence of prices and shocks in dynamic solutions. From the consumer problem's perspective, these are based on future behaviour and so they are not taken into account in solving the problem for the present.[11] As a result, they do not enter the FOC and so optimal choices will satisfy

$$\begin{aligned}(1+\lambda)u'(c^{\lambda}(x)) &= \partial_b v^{\lambda}(x) \\ \partial_a v^{\lambda}(x) &= \partial_b v^{\lambda}(x)(1 + \chi_d(d^{\lambda}(x), a))\end{aligned}$$

The distortions do impact the value itself, however, and so the optimal choices will be distorted by the anticipation of future temptation. With the effect of current temptation reducing the relative value of the liquid asset and so driving more consumption and more transfers out of the liquid account.

Hyperbolic discounting In Krusell et al. (2010) consumers are tempted to be quasi–hyperbolic discounters. We can encode continuous–time equivalent here with $\lambda > 0$ and

[11] As the derivation in Appendix 3.A.2 shows they are the result of taking the limit $\lim_{\Delta \to 0} k(\mathbf{c}_{\Delta}, \mathbf{d}_{\Delta}, x_{t+\Delta})$

the IG discount function discussed in Section 3.4

$$D(s) = \begin{cases} 1 & \text{if } s = 0 \\ \beta \cdot e^{-\rho s} & \text{else} \end{cases}$$

These preferences reduce to the following HJB equation (see Appendix 3.A.2)

$$\rho v^\lambda(x) = \max_{c^\lambda, d^\lambda} (1 + \lambda)u(c^\lambda) - \lambda u(\hat{c}(x)) + \partial_b v^\lambda(x) \cdot \dot{b} + \partial_a v^\lambda(x) \cdot \dot{a} + \mathcal{A}\left[v^\lambda\right](x)$$

And where $\hat{c}(x)$ is the solution to the sophisticated IG problem discussed in the previous section i.e. it is some multiple of the rational agent's consumption policy. With this tempting consumption policy in hand $(v^\lambda, c^\lambda, d^\lambda)$ can be solved the usual way because the above is a time–consistent problem for an agent with a distorted felicity function. These preferences embed *IG* in the limit as resistance to temptation becomes infinitely costly $\lambda \to \infty$ (Krusell et al., 2010).

Ignoring the future Most implementations of temptation preferences in the consumption–saving literature have used the DSC preferences introduced by Gul and Pesendorfer (2004) in discrete–time. This is where people are tempted to ignore the future entirely (Attanasio et al., 2024; Fudenberg and Levine, 2006; Gul and Pesendorfer, 2004; Kovacs et al., 2021; O'Donoghue and Loewenstein, 2004).

We can encode this in continuous–time with $\lambda > 0$ and the piecewise discount function

$$D(s) = \begin{cases} e^{-\rho s} & \text{if } s \leq \Delta^\lambda \\ 0 & \text{if } s > \Delta^\lambda \end{cases}$$

That is, future utility flows are discounted as normal, but only over a limited span Δ^λ. This is a little different to DSC preferences because continuous–time doesn't admit a clean distinction between now and the future. In the limit as Δ^λ goes to zero, the tempting action is to ignore all but the current instant, the tempting consumption rate explodes, and so actual consumption would explode to meet it if $\lambda > 0$. To keep the most tempting alternative finite it's therefore necessary to fix a positive span of time over which the tempted consumer's imagination wanders. For example, a natural assumption would be to mirror the discrete time literature and set Δ^λ equal to one unit of the model frequency e.g. one year.

As well as having a positive Δ^λ it's necessary to impose a terminal condition that rules out Ponzi schemes for the tempting consumption plan, particularly if there is no limit to potential credit, as in Assumption 3.3.1. To this end, I assume that consumers are not tempted to incur debt when they start with a positive liquid balance, nor are those already in debt tempted to increase it.[12]

$$\hat{b}(\Delta^{\lambda}) \geq \begin{cases} 0 & \text{if } b(0) \geq 0 \\ b(0) & \text{else} \end{cases}$$

Note that there is a relationship between Δ^{λ} and λ. If the tempting party is over a longer time–span ($\uparrow \Delta^{\lambda}$), the most–tempting flow is lower as the starting stock of assets is spread over a longer time, so some adjustment $\uparrow \lambda$ that restores the correct intensity is needed. Knowing the duration that a consumer is tempted by is an impossible identification challenge. Given this, we should expect to recover different values of λ from the same empirical exercise measuring behaviour at different frequencies, unless the most–tempting alternative is spread out over the same span.[13]

For given assumptions about the temptation parameters $(\lambda, \Delta^{\lambda})$ these preferences reduce to the following HJB equation (see Appendix 3.A.2)

$$\rho v^{\lambda}(x) = \max_{c,d} \; (1+\lambda)u(c) - \lambda u(\hat{c}) + \partial_b v^{\lambda}(x) \cdot \dot{b} + \partial_a v^{\lambda}(x) \cdot \dot{a} + \mathcal{A}\left[v^{\lambda}\right](x) \tag{3.16}$$

Where $\hat{c}(x)$ is the solution to the cake–eating problem over the limited span Δ^{λ}, funded in part by the most tempting voluntary transfer function $\hat{d}(x)$. As with the hyperbolic discounting case, once we know this most–tempting alternative consumption policy, solving for $(v^{\lambda}, c^{\lambda}, d^{\lambda})$ is simple.[14] The FOC yield the policy functions, and the HJB provides the solution for the value.

$$\begin{aligned} (1+\lambda)u'(c^{\lambda}(x)) &= \partial_b v^{\lambda}(x) \\ \partial_a v^{\lambda}(x) &= \partial_b v^{\lambda}(x)(1 + \chi_d(d^{\lambda}(x), a)) \end{aligned}$$

The HJB Equation 3.16 is the continuous–time analogue to Equation 3.15. The force of temptation boosts the utility experienced in the current moment, but also drags it down because of the tempting consumption alternative $\hat{c}$. This alternative will be closely related to cash on hand, and so this second force acts like negative money–in–utility function, making the liquid asset relatively less attractive for non–pecuniary reasons and so increasing (decreasing) the voluntary contributions to (withdrawals from) the illiquid

[12] This assumption can be justified with reference to mental accounting—people only daydream about spending money they can see, for example, or at least their demons don't tempt them with becoming criminals to fund their short party.

[13] The empirical exercises in Attanasio et al. (2024) and Kovacs et al. (2021) both use a frequency of one year, and find similar estimates of the temptation parameter.

[14] In practice, I have found a finite–difference scheme using the nested–drift algorithm (Sabet and Schneider, 2024) works well, but that it is important to start with a guessed value that assumes no temptation effect on utility, and to let the algorithm discover this force as it updates.

account.

The second and third of the above cases have the same equations describing the solution: the HJB is in the same form and policy functions determined by the same FOC. The only difference is the most–tempting consumption plan $\hat{c}(x)$. This will be the case for any tempting alternative discount function that discounts future increments at the rate ρ. With tempting discount rates that don't match this, the extra adjustment term in the first case is necessary, complicating the solution. In what follows, I will focus on the simpler case.

Relation to rational benchmark These results allow us to express the tempted agent's optimal behaviour relative to the rational benchmark policies defined earlier. Focusing on the consumption function, the FOC implies

$$u'(c^\lambda(x)) = \frac{1}{1+\lambda}\frac{\partial_b v^\lambda(x)}{\partial_b v(x)}u'(c(x))$$

Marginal utility of the tempted consumer is therefore distorted away from their rational equivalent by two forces—first, it is lower due to the pressure of current temptation $1/(1+\lambda) \leq 1$ and, second, it is lower due to the anticipation of future temptation reducing the value of carrying resources forward $\frac{\partial_b v^\lambda(x)}{\partial_b v(x)} \leq 1$. Both forces drive the consumer to over–consume relative to the rational benchmark.

The latter force has the opposite sign in IG preferences, where self–awareness reduces the impact of present bias, and the drive to over–consume. By contrast, with temptation preferences, the sophisticate over–consumes for two reasons. First, to relieve the cost of resisting his own temptation, and second to reduce the same torment for himself in the future. I will revisit this in Section 3.6 because it offers an opportunity for identifying between the two preference structures.

3.5.2 Extension to naivete

As with IG preferences, we can relax the assumption that tempted consumers correctly anticipate their future temptation by introducing naivete, where households expect their future selves to be subject to some lesser degree of temptation.

Naivete is difficult to define in the temptation model in continuous time. This is because of the potential for a clash between the horizon over which naivete is defined (separating the present from future) and over which the tempting alternative is defined (which depends on the tempting discount function). One approach could be to use a time–dependent temptation intensity.[15]

$$\lambda(s) = \begin{cases} \lambda & \text{if } s = 0 \\ \lambda^E & \text{else} \end{cases}$$

Where λ^E captures the degree of temptation experienced in the future. But this on its own leads to pathological results in cases approaching complete naivete i.e. $\lambda^E = 0$, because tempting plans will blow up as the temptation is to entirely ignore the future, leading to infinite c to avoid the cost of that temptation. To avoid this complexity, I make the following assumption

Assumption 3.5.3. A naif of any degree expects to be sophisticated in future with $\lambda^E \leq \lambda$. In the present, their most–tempting actions are those of a sophisticate with their current degree of temptation λ.

This assumption leads to the general value function

$$\tilde{v}^{\lambda}(x_t) = \lim_{\Delta \to 0} \max_{\tilde{c}^{\lambda}, \tilde{d}^{\lambda}} \left(u(\tilde{c}^{\lambda}) - \lambda \left[u(\hat{c}) - u(\tilde{c}^{\lambda}) \right] \right) \Delta + e^{-\rho\Delta} \, \mathbb{E} \, v^{\lambda^E}(x_{t+\Delta})$$

This assumption implies a contradiction in the consumer's thinking: they feel bias in the present and do not expect it to continue, but they *do* expect it to continue in the off–equilibrium path their demon is offering, from which $\hat{c}$ is derived. Under this assumption, the degree of naivete or sophistication does not affect the tempting plan felt in the moment of deciding ($\hat{c}$), just the perceived continuation value. This assumption is useful because it helps arrive at a solution, but it introduces time inconsistency into the calculation of the tempting plan $c(s)$, and so we cannot derive these naive temptation preferences from the general form.

The solution to this problem is the HJB and first order conditions

$$\begin{aligned} \rho v^{\lambda^E}(x) &= (1+\lambda) u(\tilde{c}^{\lambda}(x)) - \lambda u(\hat{c}(x)) + \partial_b v^{\lambda^E}(x) \cdot \dot{b}(x) + \partial_a v^{\lambda^E}(x) \cdot \dot{a}(x) + \mathcal{A}[v^{\lambda^E}](x) \\ u'(\tilde{c}^{\lambda}(x)) &= \frac{1}{1+\lambda} \partial_b v^{\lambda^E}(x) \\ \partial_a v^{\lambda^E}(x) &= \partial_b v^{\lambda^E}(x)(1 + \chi_d(\tilde{d}^{\lambda}(x), a)) \end{aligned}$$

Focusing on the consumption policy, the more naive an agent is, the more marginally valuable liquid resources will seem going forward, and so optimal consumption will reduce. Similarly for voluntary transfers to the illiquid account—naivete about future temptation

[15] O'Donoghue and Loewenstein (2004) introduce naivete in their model in the same way, noting that it could arise from genuine misunderstanding, or alternatively from time–inconsistency akin to quasi–hyperbolic discounting. Under that latter channel, the current consumer feels and responds to temptation, but doesn't care that their future selves will feel this as well when making decisions.

will cause consumers to under–estimate the inconvenience yield of future liquid resources, and they will keep relatively more resources in liquid form going forward.

Complete naivete Consider now the special case of complete naivete, where a consumer feels tempted in the moment, but anticipates that they'll be rational in future so $\lambda(s) = 0$ if $s > 0$ and as a result $v^{\lambda^E}(x_{t+\Delta}) = v(x_{t+\Delta})$. In this special case the consumer's optimal choices will be the same under *any* assumption about the tempting discount function.

Lemma 3.5.4 (*Tempting naive behavioural equivalence*)**.** If Assumption 3.5.3 holds, then different models of temptation, i.e. assumptions about $D(s)$, are behaviourally equivalent when the tempted consumer is completely naive.

Proof. The value function of completely naive consumers is

$$\tilde{v}^{\lambda}(x_t) = \lim_{\Delta \to 0} \max_{c,d} \left(u(c) - \lambda \left[u(\hat{c}) - u(c)\right]\right) \Delta + e^{-\rho\Delta}\, \mathbb{E}\, v(x_{t+\Delta})$$

Where $v(x_{t+\Delta})$ is the rational valuation of the future states. The cost of temptation in the present – $\lambda u(\hat{c})$ – is sunk. This is the only term in the naif's value that would be affected by the tempting discount function $D(s)$. Therefore the form of this tempting discount function cannot alter optimal choices. □

This result will be useful in comparing the different models of present–bias.

3.6 Identification between present–bias models

How could we tell the difference between these two models? Which is the more appropriate to use in macro–economic models? Now that I have defined both types of preferences in continuous time, I compare results and look for ways they differ. I start with the simplest case of complete naivete, and show the strong result that the two approaches are behaviourally equivalent, albeit not welfare equivalent. Identification between them, therefore, relies on the behaviour of sophisticates. I then extend to partial or full sophistication, showing that behaviour diverges, with tempted consumers seeking commitment devices more, and changing their consumption in different ways to IG consumers.

These results hold under the assumptions in Maxted (2024) that the borrowing constraint never binds, and CRRA utility. If we relax the first then IG consumers will seek commitment devices but only if they constrain behaviour, whereas tempted consumers will seek commitment devices that restrain their demons, and so will place more of a premium on commitment devices than IG consumers. These differences under sophistication offer opportunities for identifying between the two present–bias formulations if sophistication is measurable.

3.6.1 Naivete and identification

Assume that our agents are completely naive. That is, if they are IG consumers then $\beta^E = 1$, and if they're tempted then $\lambda^E = 0$. Under this assumption, there is a relationship between their bias parameters that delivers exactly the same behaviour.

Proposition 3.6.1 (*Naive Behavioural Equivalence*)**.** An IG consumer with future discount $\beta < 1$ and a tempted consumer with cost of temptation $\lambda > 0$ and tempting alternative consumption $\hat{c}(x)$ will make the same choices if (a) they are both completely naive, and (b) their parameters are related like this[16]

$$\beta = \frac{1}{1+\lambda}$$

Proof. The naive tempted value is a location–scale transformation of the naive IG value. To see this, note the IG consumer's value function for given choices (c, d) is

$$\tilde{w}^{\beta}(x_t) = \lim_{\Delta \to 0} u(c)\Delta + \beta e^{-\rho\Delta}\, \mathbb{E}_t\, v(x_{t+\Delta}(c, d))$$

And the tempted consumer's equivalent is

$$\tilde{v}^{\lambda}(x_t) = \lim_{\Delta \to 0} \left[(1+\lambda)u(c) - \lambda u(\hat{c}(x_t))\right] \Delta + e^{-\rho\Delta}\, \mathbb{E}_t\, v(x_{t+\Delta}(c, d))$$

If $\beta = 1/(1+\lambda)$ then these values are location–scale transformations of each other

$$\tilde{v}^{\lambda}(x) = (1+\lambda)\tilde{w}^{\beta}(x) - \lambda u(\hat{c}(x_t))dt$$

At the choices (c, d) all marginals are the same. This is true of any choices (c, d), and so it is true of the optimal choices as well. Therefore optimal choices will be the same. □

Intuitively, naive agents share a continuation value of resources, and only differ in the source of their present bias. In the moment, their decisions are guided by a boost to momentary utility (if tempted) or a reduced emphasis on the future (if IG). The relative evaluation of the present versus the future is therefore the same: emphasising the present more or the future less are two sides of the same coin.

Corollary 3.6.2 (*Nesting*)**.** Adopt the Maxted (2024) assumptions: (a) CRRA utility, and (b) non–binding borrowing constraints. There is always a naive tempted agent that behaves identically to an IG agent with *any* degree of sophistication.

[16] Fudenberg and Levine (2006) and O'Donoghue and Loewenstein (2004) both note a similar same relationship in a two–period setting, where the continuation value is the same by definition. This result is more general, applying to infinite–horizons with exponential discounting, and I also show in following text that it implies nesting of the IG model within the temptation model.

Proof. This stems from the combination of Proposition 3.4.1 that sophistication in IG preferences is isomorphic to naivete with a different discount, and Proposition 3.6.1. □

Under the Maxted (2024) assumptions, any behaviour (sophisticated or naive) under the IG model is nested by the temptation model as a special case. Note this holds regardless of the form of temptation adopted. This follows from Result 3.5.4 that temptation models are behaviourally equivalent to each other under the assumption of naivete. Krusell et al. (2010) showed that quasi–hyperbolic discounting can be nested within temptation models when consumers were actually tempted by this discount function. While this remains true, I have shown a stronger result that IG is nested under complete naivete with $\beta = 1/(1+\lambda)$, regardless of the tempting alternative discount function.

Welfare While the two approaches are behaviourally equivalent, they are not welfare equivalent. They seem to be from the perspective of the agents due to their naivete, but the appropriate welfare criterion to adopt with naive behavioural preferences erases this naivete. Taking the optimal policies as given (c^*, d^*), the appropriate welfare criteria are

$$w^\beta(x_t) = \beta\,\mathbb{E}\int_0^\infty e^{-\rho s}u(c^*(x_{t+s}))ds$$

$$v^\lambda(x_t) = (1+\lambda)\,\mathbb{E}\int_0^\infty e^{-\rho s}\left[u(c^*(x_{t+s})) - \frac{\lambda}{(1+\lambda)}u(\hat{c}(x_{t+s}))\right]ds$$

These are not at all the same. The stream of utility coming from the consumption plan are equivalent up to a multiple. But the tempted consumer experiences actual harm from the temptations their budget set offers them, distorting the consumer's welfare, and more so the greater the distance between the tempting and actual consumption $\hat{c}(x)$ and $c(x)$.

Identification The non–equivalence in welfare is important for policy–making as it will alter how we evaluate different options. But it is useless for identifying which of the two preferences is the better model to use. If people are genuinely naive, then observable variables (behaviour) will be identical. Identifying the difference between these preferences therefore hinges on the behaviour of sophisticates.

3.6.2 Sophistication and identification

Assume now that our agents are only partially naive. That is, if they are IG consumers then $\beta^E \in [\beta, 1)$ and if they're tempted then $\lambda^E \in (0, \lambda]$. The presence of some awareness of bias changes these consumers' behaviour. In the IG case, it does so in a way that is isomorphic to remaining naive but having a lower discount rate. In the tempted case the effect is more pronounced.

General Euler equation The difference is apparent in the general Euler equation describing consumption growth for the tempted consumer.

Lemma 3.6.3 (*Tempted Euler equation*)**.** Assume an agent is naive with $\lambda^E \leq \lambda$, and utility is CRRA over consumption with relative risk aversion σ. Whenever the consumption function is locally differentiable in the liquid asset, it satisfies the Euler

$$E\left[\frac{\dot{\tilde{c}}^\lambda(x)}{\tilde{c}^\lambda(x)}\right] = \frac{1}{\sigma}\left[r - \rho - \frac{\lambda^E}{1+\lambda}\left(\frac{\hat{c}(x)}{\tilde{c}^\lambda(x)}\right)^{-\sigma}\hat{c}_b(x) - \sigma\left(1 - \left(\frac{1+\lambda^E}{1+\lambda}\right)^{\frac{1}{\sigma}}\right)\tilde{c}_b^\lambda(x)\right]$$

Proof. Derived in Appendix 3.A.3 □

Recall that a rational agent's expected growth rate of consumption will equal

$$E\left[\frac{\dot{c}(x)}{c(x)}\right] = \frac{1}{\sigma}\left[r - \rho\right]$$

Expected consumption growth under temptation preferences is distorted from this rational benchmark by two forces, captured in the two extra terms above. First, expected future temptation and second, naivete. Both are distortions to the discount factor. The first scales with the *tempting* MPC, adjusted by the ratio of marginal utility from the actual to tempting consumption flow. The second scales with the actual MPC. Both forces lead to over–consumption in the present (lower expected consumption growth), and distort behaviour more for people on steeper parts of their consumption function.

This Euler embeds two important special cases. At the two ends of the naivete spectrum, one of these forces is switched off. With full sophistication $\lambda^E = \lambda$, and so only the first force is present. In this case, and assuming DSC preferences hold, the Euler has exactly the same properties as the discrete time sophisticated equivalent that has been used elsewhere (Attanasio et al., 2024; Kovacs et al., 2021). With full naivete $\lambda^E = 0$ the first force drops out and only the second remains. In this case, the Euler is identical to the naive IG agent's (with $\beta = \frac{1}{1+\lambda}$) derived in Maxted (2024).

In one respect this equivalence result reduces our knowledge. It takes a previous identification strategy off the table. That is the Euler equation method in Kovacs et al. (2021) and Huang et al. (2015). These papers estimate the degree of temptation in a life–cycle model, assuming the Euler has the sophisticated form in discrete time above. By showing expected consumption growth is increasing in liquid resources, they estimate the degree of temptation to be around $\lambda = 0.2$. But this result hinges on the assumption of sophistication. If people are at all naive, then there is an omitted variable, the MPC, that is correlated with liquid resources. As a result, the estimates from these papers will be biased if people are at all naive.[17]

[17] The papers also structurally estimate the model parameter and find similar estimate as in their Euler equation empirical estimate. The results in this paper do not invalidate that approach.

If we had an ideal dataset, we could use the Euler equation approach to identify between the two models. We would require individual measures of (a) consumption growth, (b) tempting resources (well proxied by liquid resources, wealth and income), and (c) MPC. If we found that consumption growth was increasing in tempting resources, conditional on MPC, this would be evidence in favour of the temptation model against the null of IG. The data challenge for such an approach is finding a measure of the MPC across individuals.

The effect of sophistication on consumption An alternative identification strategy could exploit theoretical differences in the marginal impact of sophistication on consumption. Recalling that the consumption functions in the two types have the following relationship to the rational benchmark, there is a predicted sign difference in the marginal impact of sophistication itself.

$$u'(\tilde{c}^{\beta}(x)) = \beta \frac{\partial v^{\beta^E}(x)}{\partial v(x)} u'(c(x))$$
$$u'(\tilde{c}^{\lambda}(x)) = \frac{1}{1+\lambda} \frac{\partial v^{\lambda^E}(x)}{\partial v(x)} u'(c(x))$$

For IG consumers, greater sophistication (holding actual bias fixed) reduces current consumption at all points in the state–space because $\frac{\partial v^{\beta^E}(x)}{\partial v(x)}$ is greater than one, and increasing as β^E declines toward β.[18] By contrast, for tempted consumers, greater sophistication leads to an increase in consumption, as they over–consume in the present as a commitment device.

This raises the possibility that we could discern between the two preferences in a cross–sectional design if we had a measure of sophistication, separate from actual bias. What would this look like? Sophistication means the ability to appreciate one's future self without rose–tinted glasses. Measuring this could mean collecting (a) proxies that correlate with this faculty e.g. cognitive capacity, as in Zhang and Greiner (2021), or the various correlates in Cobb-Clark et al. (2024), (b) self–reported measures of sophistication from surveys[19] (as in Cobb-Clark et al., 2024), or (c) data capturing consumers' forecasting errors about their own behaviour (similar to the experiments in Augenblick and Rabin (2019) or Fedyk (2024)).

An ideal dataset would therefore contain measures of sophistication, consumption, and individual states like wealth and income. The directional predictions for the impact

[18] For reasonable calibrations of the elasticity of inter–temporal substitution $\sigma > 1$.

[19] Measures of self–reported self–control issues, such as the one employed in Chapter 2 likely confound both bias and sophistication: a low self–control score could mean either a sophisticated person with ample self–control, or a tempted naif.

of sophistication on conditional consumption could be used to build a test distinguishing the two approaches to modelling present–bias.

Commitment seeking Ever since models of present–bias have been written people have been exploring whether one of the predictions—commitment seeking by sophisticated agents—bears out. A high–level summary of this literature is that while people do like the idea of commitment devices, they seek them out less than we'd expect, and this behaviour is context dependent e.g. people adopt them in lab environments a lot, but not so much in the real world (Bernheim and Taubinsky, 2018; Laibson, 2015; Maxted, 2024).

In the context of our model, commitment–seeking appears in the voluntary contribution policy—deposits into the illiquid account tie both the agents' hands, and also their demons'. If we uphold the Maxted (2024) assumptions, then *any* commitment–seeking behaviour is evidence in favour of the temptation model rather than IG. Recall that in IG under these assumptions, actual behaviour can never be constrained and so commitment is pointless. By contrast, tempted consumers will seek commitment to rob their demons of ammunition even if they continue to over–consume. Under these assumptions, this is the simplest test of the models' predictions.

But the Maxted (2024) assumptions are strong. There may be no actual limit to the resources one *could* get, if willing to take on increasingly onerous financial and social costs. But some of these costs may be more fixed than marginal, and that would be enough to replicate a hard borrowing constraint by inducing procrastination (Maxted et al., 2024). In such a world, both models predict commitment seeking behaviour, but there would still be a difference. As identified in Toussaert (2018), *tempted consumers seek to tie their demon's hands*, not necessarily their own. As a result, they may adopt commitment devices that aren't particularly constraining except by reducing temptation. Finding examples of these outside of the lab would be an interesting opportunity to test between the two theories, and bolster the evidence found there.

3.7 Conclusion

There is substantial evidence that consumers are subject to some degree of present–bias. Building this into macro models is helpful for matching aggregate moments that matter for policy–making. I have covered the two main approaches to modelling present–bias, quasi-hyperbolic discounting (QHD) and temptation preferences, and shown how to implement both in continuous time, allowing for the first direct comparison between the two.

I show that temptation preferences are the more flexible tool. They nest IG preferences as a special case (either under naivete or when the temptation is to adopt IG) and they also require fewer structural assumptions. Unlike IG, temptation preferences can be used with any utility specification, making them practical in richer macro settings with e.g. labour supply.

While the two models are behaviourally equivalent for naive agents, they have very different welfare implications for the agents, and the behavioural equivalence breaks if agents are sophisticated. That first point should give us pause before using one or the other of these preferences in policy settings. The second presents opportunities for identification between the two.

3.A Appendix

3.A.1 Rational benchmark HJB equation

The following details the steps to derive a standard HJB equation for the rational benchmark. The consumer's value general is

$$v(x_t) = \max_{\mathbf{c},\mathbf{d}} \mathbb{E} \int_0^\infty e^{-\rho s} u(c_s(x_{t+s}))ds$$

Where the state transitions are governed by the policy functions and risk processes. Suppose we're dealing with the stationary solution, and so the stream of consumption and deposit plans resolve to two policy functions $\mathbf{c} = c(x)$ and $\mathbf{d} = d(x)$. To make the HJB out of this, first separate out the present Δ from the future

$$v(x_t) = \lim_{\Delta \to 0} \max_{c,d} u(c)\Delta + e^{-\rho\Delta} v(x_{t+\Delta}(c, d))$$

Approximate the continuation value with a second order Taylor expansion around $\Delta = 0$

$$v(x_{t+\Delta}) \approx v(x_t) + \left[\partial_b v(x_t) \cdot \dot{b}_t + \partial_a v(x_t) \cdot \dot{a}_t + \mathcal{A}[v](x)\right] \Delta$$

Where $\mathcal{A}[v](x) = -\theta_z \ln z_t \partial_z v(x) + \frac{1}{2}\sigma_z^2 \partial_{zz} v(x)$ is the infinitesimal generator of the idiosyncratic risk process. Substitute this back in and also use the approximation $e^{-\rho\Delta} \approx 1 - \rho\Delta$ (that gets better in the limit as $\Delta \to 0$

$$v(x_t) = \lim_{\Delta \to 0} \max_{c,d} u(c)\Delta + (1 - \rho\Delta) \left(v(x_t) + \left[\partial_b v(x_t) \cdot \dot{b}_t + \partial_a v(x_t) \cdot \dot{a}_t + \mathcal{A}[v](x_t)\right] \Delta\right)$$

Finally, rearranging and taking the limit yields the HJB equation

$$\rho v(x) = \max_{c,d} u(c) + \partial_b v(x) \cdot \dot{b} + \partial_a v(x) \cdot \dot{a} + \mathcal{A}[v](x)$$

3.A.2 Tempted HJB equation

An HJB equation can be derived for each case of temptation preferences introduced in the main text. In each case we must do a little work to show how the tempting discount function affects the value function, and how this can be re–arranged into a form appropriate for deriving an HJB equation.

Tempted by greater discounting Consumers are tempted to adopt the alternative discount function $D(s) = e^{-\hat{\rho}s}$. Their tempting consumption plan $\hat{\mathbf{c}}$ will be the solution to the model with that discount rate, where bold case captures that this is a stream of consumption plans into the future. Their long–run preferences for a consumption and deposit plan $(\mathbf{c}, \mathbf{d})$ are

$$h(\mathbf{c}, \mathbf{d}, x) = \mathbb{E} \int_0^\infty e^{-\rho s} u(c(x_{t+s}, s)) ds$$

And the tempting utility from the same plan is

$$k(\mathbf{c}, \mathbf{d}, x) = \mathbb{E} \int_0^\infty \lambda e^{-\hat{\rho} s} u(c(x_{t+s}, s)) ds$$

Based on these, the sophisticated tempted value is

$$v^\lambda(x) = \max_{\mathbf{c},\mathbf{d}} h(\mathbf{c}, \mathbf{d}, x) - \left[k(\hat{\mathbf{c}}, \hat{\mathbf{d}}, x) - k(\mathbf{c}, \mathbf{d}, x) \right]$$

Substituting the definitions, and separating present from future

$$\begin{aligned}
v^\lambda(x) &= \max_{\mathbf{c},\mathbf{d}} \mathbb{E} \int_0^\infty e^{-\rho s} u(c(x_{t+s}, s)) ds - \mathbb{E} \int_0^\infty \lambda e^{-\hat{\rho} s} \left[u(\hat{c}(x_{t+s}, s)) - u(c(x_{t+s}, s)) \right] ds \\
v^\lambda(x) &= \lim_{\Delta \to 0} \max_{\mathbf{c},\mathbf{d}} \left[(1+\lambda) u(c(x, 0)) - \lambda u(\hat{c}(x, 0)) \right] \Delta \\
&\quad + e^{-\rho \Delta} \mathbb{E} \int_\Delta^\infty e^{-\rho s} u(c(x_{t+s}, s)) ds - e^{-\hat{\rho} \Delta} \mathbb{E} \int_\Delta^\infty \lambda e^{-\hat{\rho} s} \left[u(\hat{c}(x_{t+s}, s)) - u(c(x_{t+s}, s)) \right] ds \\
v^\lambda(x) &= \lim_{\Delta \to 0} \max_{\mathbf{c},\mathbf{d}} \left[(1+\lambda) u(c(x, 0)) - \lambda u(\hat{c}(x, 0)) \right] \Delta \\
&\quad + e^{-\rho \Delta} \left[h(\mathbf{c}_\Delta, \mathbf{d}_\Delta, x_{t+\Delta}) - \left[k(\hat{\mathbf{c}}_\Delta, \hat{\mathbf{d}}_\Delta, x_{t+\Delta}) - k(\mathbf{c}_\Delta, \mathbf{d}_\Delta, x_{t+\Delta}) \right] \right] \\
&\quad + \left[e^{-\rho \Delta} - e^{-\hat{\rho} \Delta} \right] \left[k(\hat{\mathbf{c}}_\Delta, \hat{\mathbf{d}}_\Delta, x_{t+\Delta}) - k(\mathbf{c}_\Delta, \mathbf{d}_\Delta, x_{t+\Delta}) \right]
\end{aligned}$$

And rearranging

$$\begin{aligned}
v^\lambda(x) &= \lim_{\Delta \to 0} \max_{c,d} \left[(1+\lambda) u(c) - \lambda u(\hat{c}(x)) \right] \Delta + e^{-\rho \Delta} v^\lambda(x_{t+\Delta}) \\
&\quad + (e^{-\rho \Delta} - e^{-\hat{\rho} \Delta}) \left(k(\hat{\mathbf{c}}, \hat{\mathbf{d}}, x_{t+\Delta}) - k(\mathbf{c}, \mathbf{d}, x_{t+\Delta}) \right)
\end{aligned}$$

This has the form of a standard Bellman with a distorted utility function, where the final term makes an adjustment for the discounting of the temptation cost going forward, which is not felt as strongly the continuation value suggests. To find the HJB we take a second order approximation around $\Delta = 0$ and rearrange, as in Appendix 3.A.1.

$$\begin{aligned}
\rho v^\lambda(x) &= \max_{c,d} \, (1+\lambda) u(c) - \lambda u(\hat{c}) + (\hat{\rho} - \rho) \left(k(\hat{\mathbf{c}}, \hat{\mathbf{d}}, x_{t+\Delta}) - k(\mathbf{c}, \mathbf{d}, x_{t+\Delta}) \right) \\
&\quad + \partial_b v^\lambda(x) \dot{b} + \partial_a v^\lambda(x) \dot{a} + \mathcal{A}[v^\lambda](x)
\end{aligned}$$

Tempted by quasi–hyperbolic discounting Recall the IG discount function

$$D(s) = \begin{cases} 1 & \text{if } s = 0 \\ \beta e^{-\rho s} & \text{else} \end{cases}$$

And note the tempting consumption function $\hat{c}$ for a given degree of present bias will be a scale multiple of the rational equivalent. Their long–run preferences for a consumption and deposit plan $(\mathbf{c}, \mathbf{d})$ are

$$h(\mathbf{c}, \mathbf{d}, x)) = \mathbb{E} \int_0^\infty e^{-\rho s} u(c(x_{t+s}, s)) ds$$

and note the tempting version of these satisfies

$$k(\mathbf{c}, \mathbf{d}, x) = \mathbb{E} \int_0^\infty \lambda D(s) u(c(x_{t+s}, s)) ds = \lambda \beta h(\mathbf{c}, \mathbf{d}, x))$$

Now define the sophisticated tempted value as follows

$$v^\lambda(x_t) = \max_{\mathbf{c},\mathbf{d}} \mathbb{E} \int_0^\infty e^{-\rho s} u(c(x_{t+s}, s)) ds - \mathbb{E} \int_0^\infty \lambda D(s) \left[u(\hat{c}(x_{t+s}, s)) - u(c(x_{t+s}, s)) \right] ds$$

Separate the present from the future

$$\begin{aligned} v^\lambda(x_t) = \lim_{\Delta \to 0} \max_{\mathbf{c},\mathbf{d}} & \left(u(c(x_t, 0)) - \lambda \left[u(\hat{c}(x_t, 0)) - u(c(x_t, 0)) \right] \right) \Delta \\ & + e^{-\rho \Delta} \left(\mathbb{E} \int_\Delta^\infty e^{-\rho s} u(c(x_{t+s}, s)) ds - \lambda \beta \, \mathbb{E} \int_\Delta^\infty e^{-\rho s} \left[u(\hat{c}(x_{t+s}, s)) - u(c(x_{t+s}, s)) \right] ds \right) \end{aligned}$$

Use the definitions of $h(\cdot)$ for the future parts, and let $\mathbf{c}_\Delta$ denote the consumption plan starting from $t + \Delta$

$$\begin{aligned} v^\lambda(x_t) = \lim_{\Delta \to 0} \max_{\mathbf{c},\mathbf{d}} & \left(u(c) - \lambda \left[u(\hat{c}(x_t, 0)) - u(c) \right] \right) \Delta \\ & + e^{-\rho \Delta} \, \mathbb{E} \left(h(\mathbf{c}_\Delta, \mathbf{d}_\Delta, x_{t+\Delta}) - \lambda \beta \left[h(\hat{\mathbf{c}}_\Delta, \hat{\mathbf{d}}_\Delta, x_{t+\Delta}) - h(\mathbf{c}_\Delta, \mathbf{d}_\Delta, x_{t+\Delta}) \right] \right) \end{aligned}$$

Recognise that $\lambda \beta \cdot h(\mathbf{c}, \mathbf{d}, x)) = k(\mathbf{c}, \mathbf{d}, x))$ and substitute

$$\begin{aligned} v^\lambda(x_t) = \lim_{\Delta \to 0} \max_{\mathbf{c},\mathbf{d}} & \left(u(c(x_t, 0)) - \lambda \left[u(\hat{c}(x_t, 0)) - u(c(x_t, 0)) \right] \right) \Delta \\ & + e^{-\rho \Delta} \, \mathbb{E} \left(h(\mathbf{c}_\Delta, \mathbf{d}_\Delta, x_{t+\Delta}) - \left[k(\hat{\mathbf{c}}_\Delta, \hat{\mathbf{d}}_\Delta, x_{t+\Delta}) - k(\mathbf{c}_\Delta, \mathbf{d}_\Delta, x_{t+\Delta}) \right] \right) \end{aligned}$$

And simplify to arrive at the continuation value function

$$v^{\lambda}(x_t) = \lim_{\Delta \to 0} \max_{c,d} \left(u(c) - \lambda \left[u(\hat{c}(x_t, 0)) - u(c)\right]\right) \Delta + e^{-\rho \Delta} \, \mathbb{E} \, v^{\lambda}(x_{t+\Delta})$$

This has a standard form, and so the HJB does as well. Taking a second–order approximation around $\Delta = 0$, as in Appendix 3.A.1, yields

$$\rho v^{\lambda}(x) = \max_{c,d} \, (1+\lambda) u(c) - \lambda u(\hat{c}) + \partial_b v^{\lambda}(x) \dot{b} + \partial_a v^{\lambda}(x) \dot{a} + \mathcal{A}[v^{\lambda}](x)$$

Tempted by ignoring the future This case uses the piecewise tempting discount function detailed in Section 3.5.1. The steps to derive the HJB are identical to those with quasi–hyperbolic discounting, and the result is the same except that the tempting consumption plan $\hat{c}$ solves a different problem—cake–eating over a limited span, compared with the IG consumption policy.

3.A.3 Tempted Euler equation

Here I derive an Euler equation for a tempted consumer with tempting alternative consumption $\hat{c}$ and an arbitrary level of sophistication where $\lambda^E \in [0, \lambda]$. To start, note that these consumers define their choices based on the value function they expect to prevail in future. This is defined by the HJB below

$$\begin{aligned} \rho v^{\lambda^E}(x) &= (1+\lambda^E) u(c^E(x)) - \lambda^E u(\hat{c}^E(x)) \\ &\quad + \partial_b v^{\lambda^E}(x) \dot{b}^E(x) + \partial_a v^{\lambda^E}(x) \dot{a}^E(x) + \mathcal{A}\left[v^{\lambda^E}\right](x) \end{aligned}$$

To derive the Euler we first take the derivative with respect to the liquid asset

$$\begin{aligned} \rho \partial_b v^{\lambda^E}(x) &= (1+\lambda^E) u'(c^E(x)) \partial_b c^E(x) - \lambda^E u'(\hat{c}^E(x)) \partial_b \hat{c}^E(x) \\ &\quad + \partial_b v^{\lambda^E}(x) (r(b) + r_b(b) b - \partial_b c^E(x)) \\ &\quad + \partial_{bb} v^{\lambda^E}(x) \dot{b}^E(x) + \partial_{ab} v^{\lambda^E}(x) \dot{a}^E(x) + \mathcal{A}\left[\partial_b v^{\lambda^E}\right](x) \end{aligned}$$

Substituting the *realised* FOC $\partial_b v^{\lambda^E}(x) = (1+\lambda) u'(c(x))$ to eliminate $\partial_b v^{\lambda^E}(x)$, and collecting terms

$$\begin{aligned} \left(\rho - r(b) - r_b(b) b\right) u'(c(x)) &= \frac{1+\lambda^E}{1+\lambda} u'(c^E(x)) \partial_b c^E(x) - \frac{\lambda^E}{1+\lambda} u'(\hat{c}^E(x)) \partial_b \hat{c}^E(x) \\ &\quad - u'(c(x)) \partial_b c^E(x) \\ &\quad + \partial_b u'(c(x)) \dot{b}^E(x) + \partial_a u'(c(x)) \dot{a}^E(x) + \mathcal{A}\left[u'(c(x))\right](x) \end{aligned}$$

Now recognise that the consumer expects their future selves to choose consumption to meet the FOC $(1+\lambda^E)u'(c^E(x)) = \partial_b v^{\lambda^E}(x)$. Combining this expectation with the present consumer's FOC we know $(1+\lambda^E)u'(c^E(x)) = (1+\lambda)u'(c(x))$. Substituting and collecting terms

$$\begin{aligned}(\rho - r(b) - r_b(b)b)\, u'(c(x)) = &-\frac{\lambda^E}{1+\lambda}u'(\hat{c}^E(x))\partial_b \hat{c}^E(x) \\ &+ \partial_b u'(c(x))\dot{b}^E(x) + \partial_a u'(c(x))\dot{a}^E(x) + \mathcal{A}\left[u'(c(x))\right](x)\end{aligned}$$

The next step is to track the expected path of marginal utility for the consumer in the present. The drift terms at the moment are framed in terms of their expected future policies, rather than those employed right now. We need to add and subtract $\partial_b u'(c(x))\dot{b}(x) + \partial_a u'(c(x))\dot{a}(x)$ and collect terms, substituting the definition for the time derivative of marginal utility $\mathbb{E}\left[du'(c(x))/dt\right] = \partial_b u'(c(x))\dot{b}(x) + \partial_a u'(c(x))\dot{a}(x) + \mathcal{A}\left[u'(c(x))\right](x)$. This yields

$$\begin{aligned}(\rho - r(b) - r_b(b)b)\, u'(c(x)) = &-\frac{\lambda^E}{1+\lambda}u'(\hat{c}^E(x))\partial_b \hat{c}^E(x) + \mathbb{E}\left[du'(c(x))/dt\right] \\ &+ \partial_b u'(c(x))\left[\dot{b}^E(x) - \dot{b}(x)\right] + \partial_a u'(c(x))\left[\dot{a}^E(x) - \dot{a}(x)\right]\end{aligned}$$

And now substituting for the state transition equations, and recognising that $d^E(x) = d(x)$

$$\begin{aligned}(\rho - r(b) - r_b(b)b)\, u'(c(x)) = &-\frac{\lambda^E}{1+\lambda}u'(\hat{c}^E(x))\partial_b \hat{c}^E(x) + \mathbb{E}\left[du'(c(x))/dt\right] \\ &+ \partial_b u'(c(x))\left[c^E(x) - c(x)\right]\end{aligned}$$

And then rearranging we have our general Euler equation

$$\mathbb{E}\left[\frac{du'(c(x))/dt}{u'(c(x))}\right] = \rho - r(b) - r_b(b)b + \frac{\lambda^E}{1+\lambda}\frac{u'(\hat{c}^E(x))}{u'(c(x))}\partial_b \hat{c}^E(x) - \frac{\partial_b u'(c(x))}{u'(c(x))}\left[c^E(x) - c(x)\right]$$

Here we can see the general Euler equation includes two separate distortions to discounting stemming from temptation and sophistication about that temptation. In the limit where consumers are fully sophisticated, the final term drops out and we're left with just the one distortion. Similarly, when consumers are fully naive the second to last term drops out, and we're left with the other.

The Euler simplifies further when we assume CRRA utility such that $u'(c) = c^{-\sigma}$ and we know from combining FOC that $c^E(x) = \left(\frac{1+\lambda}{1+\lambda^E}\right)^{-\frac{1}{\sigma}} c(x)$ and also substituting for the coefficient of relative risk aversion $-\sigma = \frac{u''(c(x))c(x)}{u'(c(x))}$

$$\mathbb{E}\left[\frac{\dot{c}(x)}{c(x)}\right] = -\frac{1}{\sigma}\left(\rho - r(b) - r_b(b)b + \frac{\lambda^E}{1+\lambda}\left(\frac{\hat{c}^E(x)}{c(x)}\right)^{-\sigma}\partial_b\hat{c}^E(x) + \sigma\left[1 - \left(\frac{1+\lambda^E}{1+\lambda}\right)^{\frac{1}{\sigma}}\right]\partial_b c(x)\right)$$

This Euler simplifies to two special cases. With complete sophistication, we recover the continuous–time equivalent of the tempting Euler equation derived in Attanasio et al. (2024). With complete naivete, we recover the equivalent of the naive IG Euler equation in Maxted (2024), noting hat $1/(1+\lambda)$ takes the place of β.

Bibliography

Achdou, Y., J. Han, J.-M. Lasry, P.-L. Lions, and B. Moll (2022): "Income and wealth distribution in macroeconomics: A continuous-time approach," *The review of economic studies*, 89, 45–86.

Adams, J. and D. Nettle (2009): "Time perspective, personality and smoking, body mass, and physical activity: An empirical study," *British journal of health psychology*, 14, 83–105.

Aguiar, M., M. Bils, and C. Boar (2024): "Who are the Hand-to-Mouth?" *Review of Economic Studies*.

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

——— (2013): "Deconstructing life cycle expenditure," *Journal of Political Economy*, 121, 437–492.

Aguila, E., O. Attanasio, and C. Meghir (2011): "Changes in consumption at retirement: evidence from panel data," *Review of Economics and Statistics*, 93, 1094–1099.

Aiyagari, S. R. (1994): "Uninsured idiosyncratic risk and aggregate saving," *The Quarterly Journal of Economics*, 109, 659–684.

Almlund, M., A. L. Duckworth, J. Heckman, and T. Kautz (2011): "Personality Psychology and Economics," in *Handbook of The Economics of Education*, ed. by E. A. Hanushek, S. Machin, and L. Woessmann, Elsevier, vol. 4 of *Handbook of the Economics of Education*, 1–181.

Amador, M., I. Werning, and G.-M. Angeletos (2006): "Commitment vs. flexibility," *Econometrica*, 74, 365–396.

Ameriks, J., A. Caplin, and J. Leahy (2003): "Wealth accumulation and the propensity to plan," *The Quarterly Journal of Economics*, 118, 1007–1047.

Amromin, G. and P. Smith (2003): "What explains early withdrawals from retirement accounts? Evidence from a panel of taxpayers," *National Tax Journal*, 56, 595–612.

Andersen, H. Y. (2020): "Spending when illiquid savings become liquid: Evidence from Danish wage earners," *Unpublished Working Paper*.

Andersen, H. Y., A. Bartscher, S. Leth-Petersen, and P. Moran (2024): "Early Withdrawals and Optimal Liquidity," *Unpublished Working Paper*.

Andreoni, J., M. A. Kuhn, and C. Sprenger (2015): "Measuring time preferences: A comparison of experimental methods," *Journal of Economic Behavior & Organization*, 116, 451–464.

Andreoni, J. and C. Sprenger (2012): "Estimating time preferences from convex budgets," *American Economic Review*, 102, 3333–3356.

Angeletos, G.-M., C. Lian, and C. K. Wolf (2023): "Can Deficits Finance Themselves?" *NBER Working Paper*.

Argento, R., V. L. Bryant, and J. Sabelhaus (2015): "Early withdrawals from retirement accounts during the Great Recession," *Contemporary Economic Policy*, 33, 1–16.

Attanasio, O., A. Kovacs, and P. Moran (2024): "Temptation and Commitment: a Model of Hand-to-Mouth Behavior," *Journal of the European Economic Association*, 2025–2073.

Auclert, A., M. Rognlie, and L. Straub (2023): "The trickling up of excess savings," *AEA Papers and Proceedings*, 113, 70–75.

——— (2024a): "Fiscal and Monetary Policy with Heterogeneous Agents," *Unpublished Working Paper*.

——— (2024b): "The Intertemporal Keynesian Cross," *Journal of Political Economy*, 132, 4068–4121.

Augenblick, N., M. Niederle, and C. Sprenger (2015): "Working over time: Dynamic inconsistency in real effort tasks," *The Quarterly Journal of Economics*, 130, 1067–1115.

Augenblick, N. and M. Rabin (2019): "An experiment on time preference and misprediction in unpleasant tasks," *Review of Economic Studies*, 86, 941–975.

Baker, S. R., R. A. Farrokhnia, S. Meyer, M. Pagel, and C. Yannelis (2023): "Income, liquidity, and the consumption response to the 2020 economic stimulus payments," *Review of Finance*, 27, 2271–2304.

Banks, J., R. Blundell, and S. Tanner (1998): "Is there a retirement-savings puzzle?" *American Economic Review*, 769–788.

Banks, J., C. o'Dea, and Z. Oldfield (2010): "Cognitive function, numeracy and retirement saving trajectories," *The Economic Journal*, 120, F381–F410.

Barles, G. and P. E. Souganidis (1991): "Convergence of approximation schemes for fully nonlinear second order equations," *Asymptotic analysis*, 4, 271–283.

Barsky, R. B., F. T. Juster, M. S. Kimball, and M. D. Shapiro (1997): "Preference parameters and behavioral heterogeneity: An experimental approach in the health and retirement study," *The Quarterly Journal of Economics*, 112, 537–579.

Bateman, H., L. I. Dobrescu, J. Liu, B. R. Newell, and S. Thorp (2023): "Determinants of early-access to retirement savings: Lessons from the COVID-19 pandemic," *The Journal of the Economics of Ageing*, 24, 100441.

Bateman, H., G. Kingston, and J. Piggott (2001): *Forced saving: Mandating private retirement incomes*, Cambridge University Press.

Baugh, B., I. Ben-David, H. Park, and J. A. Parker (2021): "Asymmetric Consumption Smoothing," *American Economic Review*, 111, 192–230.

Bayer, C., B. Born, and R. Luetticke (2023): "The liquidity channel of fiscal policy," *Journal of Monetary Economics*, 134, 86–117.

Bayer, C., R. Luetticke, L. Pham-Dao, and V. Tjaden (2019): "Precautionary Savings, Illiquid Assets, and the Aggregate Consequences of Shocks to Household Income Risk," *Econometrica*, 87, 255–290.

Bernheim, B. D., J. Skinner, and S. Weinberg (2001): "What accounts for the variation in retirement wealth among US households?" *American Economic Review*, 91, 832–857.

Bernheim, B. D. and D. Taubinsky (2018): "Chapter 5 - Behavioral Public Economics," in *Handbook of Behavioral Economics - Foundations and Applications 1*, ed. by B. D. Bernheim, S. DellaVigna, and D. Laibson, North-Holland, vol. 1 of *Handbook of Behavioral Economics: Applications and Foundations 1*, 381–516.

Bertrams, A. and O. Dickhäuser (2009): "Messung dispositioneller selbstkontrollkapazität: eine deutsche adaptation der kurzform der self-control scale (SCS-KD)," *Diagnostica*, 55, 2–10.

Beshears, J., J. J. Choi, C. Clayton, C. Harris, D. Laibson, and B. C. Madrian (2025): "Optimal illiquidity," *Journal of Financial Economics*, 165, 103996.

Beshears, J., J. J. Choi, C. Harris, D. Laibson, B. C. Madrian, and J. Sakong (2020): "Which early withdrawal penalty attracts the most deposits to a commitment savings account?" *Journal of Public Economics*, 183, 104144.

Beshears, J., J. J. Choi, J. Hurwitz, D. Laibson, and B. C. Madrian (2015): "Liquidity in Retirement Savings Systems: An International Comparison," *American Economic Review*, 105, 420–25.

Blanchard, O. J. (1985): "Debt, deficits, and finite horizons," *Journal of political economy*, 93, 223–247.

Blundell, R., E. French, and G. Tetlow (2016): "Retirement incentives and labor supply," in *Handbook of the economics of population aging*, Elsevier, vol. 1, 457–566.

Boehm, J., E. Fize, and X. Jaravel (2025): "Five facts about MPCs: Evidence from a randomized experiment," *American Economic Review*, 115, 1–42.

Borghans, L., A. L. Duckworth, J. J. Heckman, and B. Ter Weel (2008): "The economics and psychology of personality traits," *Journal of human Resources*, 43, 972–1059.

Bosch, M., C. C. Felix, M. García-Huitrón, and M. T. S. Porto (2020): "Access to mandatory retirement savings in times of COVID-19: public policy considerations," *Inter-American Development Bank Working Paper*.

Brown, H. and M. Van der Pol (2015): "Intergenerational transfer of time and risk preferences," *Journal of Economic Psychology*, 49, 187–204.

Campbell, J. Y. and N. G. Mankiw (1989): "Consumption, income, and interest rates: Reinterpreting the time series evidence," *NBER Macroeconomics Annual*, 4, 185–216.

Carroll, C., J. Slacalek, K. Tokuoka, and M. N. White (2017): "The distribution of wealth and the marginal propensity to consume," *Quantitative Economics*, 8, 977–1020.

Carroll, C. D., M. B. Holm, and M. S. Kimball (2021): "Liquidity constraints and precautionary saving," *Journal of Economic Theory*, 195, 105276.

Carroll, C. D. and M. S. Kimball (1996): "On the Concavity of the Consumption Function," *Econometrica*, 64, 981–992.

Chang, Y., S.-B. Kim, and F. Schorfheide (2013): "Labor-market heterogeneity, aggregation, and policy (in) variance of DSGE model parameters," *Journal of the European Economic Association*, 11, 193–220.

Chetty, R., J. N. Friedman, S. Leth-Petersen, T. H. Nielsen, and T. Olsen (2014): "Active vs. Passive Decisions and Crowd-Out in Retirement Savings Accounts: Evidence from Denmark," *The Quarterly Journal of Economics*, 129, 1141–1220.

Choukhmane, T., J. Colmenares, C. O'Dea, J. Rothbaum, and L. Schmidt (2024): "Who benefits from retirement saving incentives in the US? Evidence on racial gaps in retirement wealth accumulation," *NBER Working Paper*.

Choukhmane, T. and C. Palmer (2024): "How do consumers finance increased retirement savings," *Unpublished Working Paper*.

Cobb-Clark, D. A., S. C. Dahmann, D. A. Kamhöfer, and H. Schildberg-Hörisch (2024): "Sophistication about self-control," *Journal of Public Economics*, 238, 105196.

Cobb-Clark, D. A., S. C. Dahmann, D. A. Kamhöfer, and H. Schildberg-Hörisch (2022): "The Predictive Power of Self-Control for Life Outcomes," *Journal of Economic Behavior & Organization*, 197, 725–744.

Coyne, D., I. Fadlon, and T. Porzio (2022): "Measuring valuation of liquidity with penalized withdrawals," .

Deaton, A. (1989): "Saving and liquidity constraints," *NBER Working Paper*.

Epper, T., E. Fehr, H. Fehr-Duda, C. T. Kreiner, D. D. Lassen, S. Leth-Petersen, and G. N. Rasmussen (2020): "Time discounting and wealth inequality," *American Economic Review*, 110, 1177–1205.

Fagereng, A., M. B. Holm, B. Moll, and G. Natvik (2019): "Saving behavior across the wealth distribution: The importance of capital gains," .

FAGERENG, A., M. B. HOLM, AND G. J. NATVIK (2021): "MPC heterogeneity and household balance sheets," *American Economic Journal: Macroeconomics*, 13, 1–54.

FEDYK, A. (2024): "Asymmetric naivete: Beliefs about self-control," *Management Science*.

FELDSTEIN, M. (1985): "The optimal level of social security benefits," *The Quarterly Journal of Economics*, 100, 303–320.

FLAVIN, M. A. (1981): "The adjustment of consumption to changing expectations about future income," *Journal of political economy*, 89, 974–1009.

FLINN, C. J., P. E. TODD, AND W. ZHANG (2018): "Personality traits, intra-household allocation and the gender wage gap," *European Economic Review*, 109, 191–220.

FLODEN, M. AND J. LINDÉ (2001): "Idiosyncratic risk in the United States and Sweden: Is there a role for government insurance?" *Review of Economic dynamics*, 4, 406–437.

FRIEDMAN, M. (1957): "The Permanent Income Hypothesis," in *A Theory of the Consumption Function*, Princeton University Press, 20–37.

FUDENBERG, D. AND D. K. LEVINE (2006): "A dual-self model of impulse control," *American Economic Review*, 96, 1449–1476.

GALÍ, J. (2020): "The effects of a money-financed fiscal stimulus," *Journal of Monetary Economics*, 115, 1–19.

GANONG, P., D. JONES, P. J. NOEL, F. E. GREIG, D. FARRELL, AND C. WHEAT (2023): "Wealth, race, and consumption smoothing of typical income shocks," *Unpublished Working Paper*.

GANONG, P. AND P. NOEL (2020): "Liquidity versus wealth in household debt obligations: Evidence from housing policy in the great recession," *American Economic Review*, 110, 3100–3138.

GELMAN, M. (2022): "The Self-Constrained Hand-to-Mouth," *The Review of Economics and Statistics*, 104, 1096–1109.

GELMAN, M., S. KARIV, M. D. SHAPIRO, AND D. SILVERMAN (2022): "Rational Illiquidity and Consumption: Theory and Evidence from Income Tax Withholding and Refunds," *American Economic Review*, 112, 2959–91.

GERARD, F. AND J. NARITOMI (2021): "Job displacement insurance and (the lack of) consumption-smoothing," *American Economic Review*, 111, 899–942.

GODA, G. S., D. JONES, AND S. RAMNATH (2022): "Temporary and permanent effects of withdrawal penalties on retirement savings accounts," *Journal of Public Economics*, 215, 104734.

GODA, G. S., M. LEVY, C. F. MANCHESTER, A. SOJOURNER, AND J. TASOFF (2019): "Predicting Retirement Savings Using Survey Measures of Exponential-Growth Bias and Present Bias," *Economic Inquiry*, 57, 1636–1658.

Goodman, L., J. Mortenson, K. Mackie, and H. R. Schramm (2021): "Leakage from retirement savings accounts in the united states," *National Tax Journal*, 74, 689–719.

Graves, S. (2023): "Does Unemployment Risk Affect Business Cycle Dynamics?" *Unpublished Working Paper.*

Guerrieri, V. and G. Lorenzoni (2017): "Credit Crises, Precautionary Savings, and the Liquidity Trap," *The Quarterly Journal of Economics*, 132, 1427–1467.

Gul, F. and W. Pesendorfer (2001): "Temptation and self-control," *Econometrica*, 69, 1403–1435.

——— (2004): "Self-control, revealed preference and consumption choice," *Review of Economic Dynamics*, 7, 243–264.

Guvenen, F., G. Kambourov, B. Kuruscu, S. Ocampo, and D. Chen (2023): "Use it or lose it: Efficiency and redistributional effects of wealth taxation," *The Quarterly Journal of Economics*, 138, 835–894.

Hall, R. E. (1978): "Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence," *Journal of political economy*, 86, 971–987.

Hamilton, S., G. Liu, J. Miranda-Pinto, and T. Sainsbury (2024): "A $100,000 marshmallow experiment: Withdrawal and spending responses to early retirement-savings access," *Unpublished Working Paper.*

Harris, C. and D. Laibson (2013): "Instantaneous Gratification," *The Quarterly Journal of Economics*, 128, 205–248.

Heckman, J. J., T. Jagelka, and T. D. Kautz (2021): "Some contributions of economics to the study of personality," *Handbook of Personality - Theory and Research.*

Holzmann, R., R. P. Hinz, and M. Dorfman (2008): "Pension systems and reform conceptual framework," *World Bank Discussion Paper*, 824.

Huang, K. X., Z. Liu, and J. Q. Zhu (2015): "Temptation and Self-Control: Some Evidence and Applications," *Journal of Money, Credit and Banking*, 47, 581–615.

Jappelli, T. and L. Pistaferri (2010): "The consumption response to income changes," *Annu. Rev. Econ.*, 2, 479–506.

——— (2014): "Fiscal policy and MPC heterogeneity," *American Economic Journal: Macroeconomics*, 6, 107–136.

Kaplan, G., K. Mitman, and G. L. Violante (2020a): "The housing boom and bust: Model meets evidence," *Journal of Political Economy*, 128, 3285–3345.

Kaplan, G., B. Moll, and G. L. Violante (2018): "Monetary policy according to HANK," *American Economic Review*, 108, 697–743.

——— (2020b): "The great lockdown and the big stimulus: Tracing the pandemic possibility frontier for the US," *NBER Working Paper.*

Kaplan, G. and G. L. Violante (2014): "A model of the consumption response to fiscal stimulus payments," *Econometrica*, 82, 1199–1239.

——— (2022): "The marginal propensity to consume in heterogeneous agent models," *Annual Review of Economics*, 14, 747–775.

Kaplan, G., G. L. Violante, and J. Weidner (2014): "The wealthy hand-to-mouth," .

Karger, E. and A. Rajan (2020): "Heterogeneity in the marginal propensity to consume: evidence from Covid-19 stimulus payments," *FRB of Chicago Working Paper No. WP 2020-15*.

Kolsrud, J., C. Landais, D. Reck, and J. Spinnewijn (2024): "Retirement consumption and pension design," *American Economic Review*, 114, 89–133.

Kovacs, A., H. Low, and P. Moran (2021): "Estimating Temptation and Commitment Over the Life Cycle," *International Economic Review*, 62, 101–139.

Kreiner, C. T., D. D. Lassen, and S. Leth-Petersen (2019a): "Liquidity constraint tightness and consumer responses to fiscal stimulus policy," *American Economic Journal: Economic Policy*, 11, 351–379.

——— (2019b): "Liquidity constraint tightness and consumer responses to fiscal stimulus policy," *American Economic Journal: Economic Policy*, 11, 351–379.

Krusell, P., B. Kuruşçu, and A. A. Smith Jr (2010): "Temptation and taxation," *Econometrica*, 78, 2063–2084.

Kueng, L. (2018): "Excess Sensitivity of High-Income Consumers," *The Quarterly Journal of Economics*, 133, 1693–1751.

Laibson, D. (1997): "Golden Eggs and Hyperbolic Discounting," *The Quarterly Journal of Economics*, 112, 443–478.

——— (2015): "Why Don't Present-Biased Agents Make Commitments?" *American Economic Review*, 105, 267–72.

Laibson, D., S. C. Lee, P. Maxted, A. Repetto, and J. Tobacman (2024): "Estimating Discount Functions with Consumption Choices over the Lifecycle," *The Review of Financial Studies*.

Laibson, D. and P. Maxted (2023): "The Beta-Delta-DELTA Sweet Spot," .

Lee, S. C. and P. Maxted (2023): "Credit Card Borrowing in Heterogeneous-Agent Models: Reconciling Theory and Data," *Unpublished Working Paper*.

Love, D. A. (2017): "Countercyclical retirement accounts," *European Economic Review*, 98, 32–48.

Lusardi, A. and O. S. Mitchell (2014): "The economic importance of financial literacy: Theory and evidence," *American Economic Journal: Journal of Economic Literature*, 52, 5–44.

Madeira, C. (2022): "The impact of the Chilean pension withdrawals during the Covid pandemic on the future savings rate," *Journal of International Money and Finance*, 126, 102650.

——— (2024): "The effect of the Covid pension withdrawals and the Universal Guaranteed Pension on the income of the future retirees and its fiscal costs," *Latin American Journal of Central Banking*, 5, 100122.

Maloney, P. W., M. J. Grawitch, and L. K. Barber (2012): "The multi-factor structure of the Brief Self-Control Scale: Discriminant validity of restraint and impulsivity," *Journal of Research in Personality*, 46, 111–115.

Manapat, P. D., M. C. Edwards, D. P. MacKinnon, R. A. Poldrack, and L. A. Marsch (2021): "A psychometric analysis of the Brief Self-Control Scale," *Assessment*, 28, 395–412.

Maxted, P. (2024): "Present Bias Unconstrained: Consumption, Welfare, and the Present–Bias Dilemma," *Unpublished Working Paper*.

Maxted, P., D. Laibson, and B. Moll (2024): "Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy," *The Quarterly Journal of Economics*, 691–743.

Modigliani, F. and R. Brumberg (1954): "Utility analysis and the consumption function: An interpretation of cross-section data," *Franco Modigliani*, 1, 388–436.

Moran, P., M. O'Connell, C. O'Dea, F. Parodi, M. R. Submitter, et al. (2021): "Heterogeneity in household spending and well-being around retirement," *Heterogeneity in Household Spending and Well-being around Retirement (September 2021)*.

Moser, C. and P. Olea de Souza e Silva (2019): "Optimal paternalistic savings policies," *Columbia Business School Research Paper*.

Naik, C. and D. Reck (2024): "Intrapersonal Utility Comparisons as Interpersonal Utility Comparisons: Welfare, Ambiguity, and Robustness in Behavioral Policy Problems," .

Ni, S. and Y. Seol (2014): "New evidence on excess sensitivity of household consumption," *Journal of Monetary Economics*, 63, 80–94.

O'Donoghue, T. and M. Rabin (2006): "Optimal sin taxes," *Journal of Public Economics*, 90, 1825–1849.

OECD (2018): *Financial Incentives and Retirement Savings*, OECD Publishing, Paris.

——— (2021): *Pension Markets in Focus 2021*, OECD Publishing, Paris.

——— (2023): *Pensions at a Glance 2023*, OECD Publishing, Paris.

Oh, H. and R. Reis (2012): "Targeted transfers and the fiscal response to the great recession," *Journal of Monetary Economics*, 59, S50–S64.

Olafsson, A. and M. Pagel (2018): "The Liquid Hand-to-Mouth: Evidence from Personal Finance Management Software," *The Review of Financial Studies*, 31, 4398–4446.

O'Donoghue, T. and G. Loewenstein (2004): "Animal Spirits: Affective and Deliberative Processes in Economic Behavior," *Unpublished Working Paper*.

Parker, J. A. (2017): "Why Don't Households Smooth Consumption? Evidence from a $25 Million Experiment," *American Economic Journal: Macroeconomics*, 9, 153–83.

Phelps, E. S. and R. A. Pollak (1968): "On Second-Best National Saving and Game-Equilibrium Growth1," *The Review of Economic Studies*, 35, 185–199.

Poterba, J. M. (2014): "Retirement Security in an Aging Population," *American Economic Review*, 104, 1–30.

Poterba, J. M. and J. J. Rotemberg (1986): "Money in the utility function: An empirical implementation," .

Preston, A. (2022): "Financial fragility, financial literacy and the early withdrawal of retirement savings during COVID-19," *Australian Journal of Labour Economics*, 25, 127–147.

Sabet, S. and P. Schneider (2024): "The Nested–Drift Algorithm," *Unpublished Working Paper*.

Sahm, C. R., M. D. Shapiro, and J. Slemrod (2010): "Household response to the 2008 tax rebate: Survey evidence and aggregate implications," *Tax Policy and the Economy*, 24, 69–110.

——— (2012): "Check in the mail or more in the paycheck: does the effectiveness of fiscal stimulus depend on how it is delivered?" *American Economic Journal: Economic Policy*, 4, 216–250.

Samwick, A. A. (1998): "Discount rate heterogeneity and social security reform," *Journal of Development Economics*, 57, 117–146.

Schneider, P. and P. Moran (2024a): "Household Liquidity Policy," *Unpublished Working Paper*.

——— (2024b): "Situational and Behavioral Determinants of Early Withdrawal from Retirement Accounts," *Unpublished Working Paper*.

Shapiro, J. M. (2005): "Is there a daily discount rate? Evidence from the food stamp nutrition cycle," *Journal of Public Economics*, 89, 303–325.

Shapiro, M. D. and J. Slemrod (1995): "Consumer Response to the Timing of Income: Evidence from a Change in Tax Withholding," *The American Economic Review*, 85, 274–283.

Shimer, R. (2005): "The cyclical behavior of equilibrium unemployment and vacancies," *American Economic Review*, 95, 25–49.

Stango, V. and J. Zinman (2023): "We are all behavioural, more, or less: A taxonomy of consumer decision-making," *The Review of Economic Studies*, 90, 1470–1498.

Strotz, R. H. (1955): "Myopia and Inconsistency in Dynamic Utility Maximization," *The Review of Economic Studies*, 23, 165–180.

Tangney, J. P., R. F. Baumeister, and A. L. Boone (2004): "High Self-Control Predicts Good Adjustment, Less Pathology, Better Grades, and Interpersonal Success," *Journal of Personality*, 72, 271–324.

Toczynski, J. (2023): "Spend or invest? Analyzing MPC heterogeneity across three stimulus programs," *Unpublished Working Paper*.

Todd, P. E. and W. Zhang (2020): "A dynamic model of personality, schooling, and occupational choice," *Quantitative Economics*, 11, 231–275.

Toussaert, S. (2018): "Eliciting temptation and self-control through menu choices: A lab experiment," *Econometrica*, 86, 859–889.

Yaari, M. E. (1965): "Uncertain lifetime, life insurance, and the theory of the consumer," *The Review of Economic Studies*, 32, 137–150.

Zeldes, S. P. (1989): "Consumption and liquidity constraints: an empirical investigation," *Journal of political economy*, 97, 305–346.

Zhang, Q. and B. Greiner (2021): "Time inconsistency, sophistication, and commitment: An experimental study," *Economics Letters*, 206, 109982.