Optimal Life Cycle Liquidity

Francisco Cabezón & Sebastián Guarda, Princeton University

October, 2022

Abstract

We study optimal household liquidity throughout the life cycle in the presence of incomplete markets and self-control issues in consumption/savings, e.g. due to present bias or costly temptation. In a deterministic setting, we provide a simple theory using age-dependent taxes/transfers with zero present value. Under just incomplete markets households always prefer front-loaded disposable income profiles, but if they additionally have overconsumption issues a paternalistic government will not want to provide young households too much liquidity. We characterize the unique optimal policy: the Euler equation holds with equality but households are Hand-to-Mouth. In a stochastic setting, we build a quantitative life-cycle model, match the age profile of borrowing-constrained households in the data, and calculate welfare gains from a policy experiment following a simple rule. We find that front-loading disposable income can increase average welfare by 0.04% in consumption-equivalent terms, and 0.06% for Hand-to-Mouth households.

1 Introduction

A great deal of policies have heterogenous effects on the income of households of different age: progressive income taxation hits older people harder when income is increasing through the life cycle, unemployment payments will benefit younger workers more if they spend a longer time without a job than older workers, and child tax benefits will mostly improve the situation of middle-aged households. While all these policies may serve different purposes and alter incentives in their own way, at some level of abstraction -for example, a stationary environment and a representative household- they all imply redistribution of resources across the life cycle. This means that, on average, if a household is benefitted while it is young then it will pay for these benefits when it is older, changing the profile of disposable income across age.

This paper asks: assuming a constant present value, what is the optimal disposable income profile? To answer this question in a unified way across policies, we assume that the age-dependent taxes and transfers that attain this profile are lump-sum, therefore they do not change incentives or wealth, they just redistribute liquidity throughout the life cycle.[1] In this sense, all the policies above will imply a change in household welfare through this liquidity channel, in addition to their own particular distortions or corrections.

[1] We consider the empirical or underlying income profile as a fixed feature of the environment. Then, we can speak indistinctly about the optimal disposable income profile, or the set of optimal age-dependent taxes and transfers, one implies the other. The liquidity profile -amount of liquid assets across the life cycle- will emerge endogenously from consumption-savings decisions.

With complete financial markets the answer is trivial: any change does not matter because households will save or borrow optimally no matter the income profile. With incomplete markets in the form of borrowing constraints, the answer is a corner solution: give as much liquidity to the young as is necessary to ensure that their borrowing constraints are not binding, this is equivalent to granting them a much-needed loan. However, many policies implemented around the world do quite the opposite and institute compulsory savings for working-age households. They do so on the basis of paternalistic motives which assume some kind of undersaving issue in preferences, for example present bias as in Laibson (1997) or a costly temptation to splurge as in Gul and Pesendorfer (2004). These models would predict that young households with too much liquidity would go on a spending spree that is undesirable from a long run point of view, as they will have little savings to confront later stages of life. There is also ample support for the predictions of these kind of preferences in laboratory and field settings.[2]

Once we consider these self-control issues as well as incomplete markets, we actually have a well-defined optimum: a paternalistic government will give households just enough liquidity at each age such that they can smooth lifetime consumption, but no more than that. For deterministic problems, we show that this implies a unique optimal income profile where an Euler equation holds with equality at all times but they have no liquidity, meaning they could not borrow to increase their consumption if they wanted to. This Euler equation is that of a household with no self-control issues but otherwise identical to the households in the economy, yet the optimal tax profile maximizes the welfare of the household **with** self-control issues. The result is robust to the size of the overconsumption problems, as long as they exist the optimal income profile is still the same and unique. The result is also independent of the degree of awareness that present-biased households may have, that is whether they are sophisticated or naive about their time-inconsistency.

A deterministic setting is convenient to derive this theoretical result, but one of the main uses of liquidity is to smooth consumption in the presence of uninsurable income risk. To study the optimal income profile in this setting, we build a quantitative model with borrowing constraints and uninsurable income risk in continuous-time that extends the Achdou et al. (2022) framework to include a life cycle and overconsumption preferences.[3] A stochastic setting lets us use a key statistic to tie our model to the data: the percentage of borrowing-constrained households across age, which Kaplan, Violante, and Weidner (2014) documented as Hand to Mouth (HtM) households. Henceforth we refer to this statistic as the empirical profile of liquidity, and we use their data and approach in Figure 1. We plot two profiles: the "Poor HtM" refer to households who are poor in both liquid and illiquid assets, while "All HtM" includes what Kaplan, Violante, and Weidner (2014) coined the "Wealthy HtM": households with little liquid assets but some illiquid assets. We can see that both profiles are decreasing in age, through the lens of our theory this can suggest that younger households are disproportionately restricted in liquidity. We discuss this interpretation in Section 5.

In this quantitative framework, for the time being we provide a simple exercise that changes the

[2]Present bias: in lab settings Thaler (1981); Frederick et al. (2002); Cohen et al. (2020), in field settings Meier and Sprenger (2010); Ganong and Noel (2019); Allcott et al. (2022). Costly temptation: in lab setting Toussaert (2018), in field setting Schilbach (2019), in structural estimation Kovacs et al. (2021).

[3]For the time being, we have incorporated naive present bias. Future work will include temptation preferences and/or sophisticated present bias.

Figure 1: Empirical Liquidity Profile

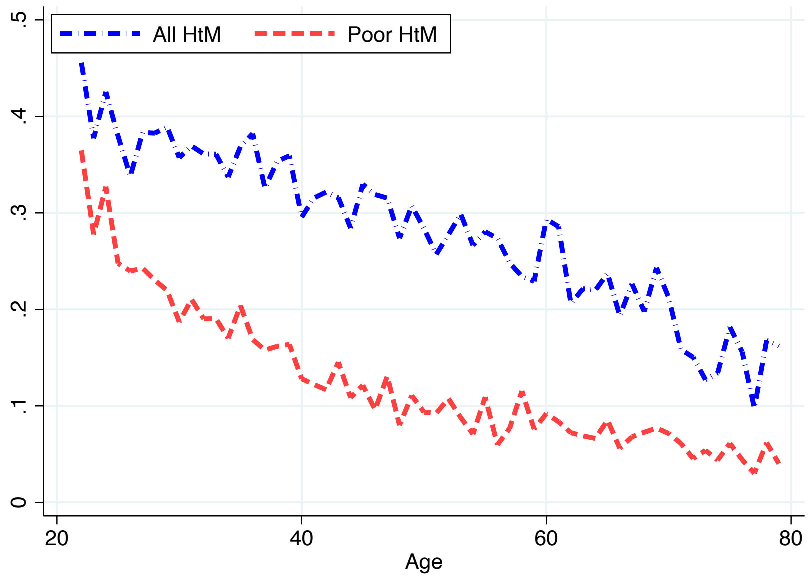

Source: Kaplan, Violante, and Weidner (2014) based on the U.S. Survey of Consumer Finances. Vertical axis corresponds to the fraction of households in the category, given the age of the head of the household.

slope of the income profile, that is we calculate the optimal slope of taxes/transfers.[4] We find that if we increase disposable income of the youngest households by 14% and linearly phase it out until they pay it in taxes at later stages, we can increase their welfare by an average of 0.04% in consumption-equivalent terms. This is due to a steeply increasing income profile which, when coupled with borrowing constraints, prevents consumption smoothing through the life cycle, a problem that this simple policy can help with. Self-control issues give us a well-defined optimum for this policy, we find welfare to be concave in the slope of the policy. To support this claim, we run the same exercise in a calibration where self-control issues are absent and we find the results from the deterministic setting carry over: households always prefer more front-loaded income profiles, because liquidity is always harmless in this environment.

A secondary contribution of this paper will be a byproduct of the calibration: we want to characterize how the forces in the model shape its liquidity profile, and pin down the key ingredients needed to reproduce the empirical profile. The strength of self-control issues is relevant in this case, because it influences consumption/savings decisions in the absence of an optimal tax profile, which is true for both deterministic and stochastic settings. This contribution is not quite ready yet, as we don't match the liquidity profile as well as we would like to, but we provide a discussion with some promising alternatives for future work.

A final comment: this paper is very much a work in progress, and especially so in the quantitative exercise, which we acknowledge requires a lot more work to be convincing, and where many more exercises could be carried out. However, our hope is that this paper poses an interesting, relevant and novel research question, and a promising approach to answer it.

[4]Later work will solve the full Ramsey optimal policy, which requires finding a nonparametric function of time.

The paper proceeds as follows. Section 2 reviews the related literature on life cycle liquidity, self-control in consumption/savings, and age-dependent taxation. Section 3 studies our question in a deterministic setting to allow theoretical results and build intuition. Section 4 builds and calibrates the quantitative model with stochastic income, and Section 5 analyzes the results of some policy experiments. Finally, Section 6 concludes.

2 Literature

To the best of our knowledge, the literature hasn't directly addressed our question of an optimal disposable income profile when taxes and transfers do not affect incentives (the age-dependent tax literature is discussed below), so this is our main contribution. This could be because one needs a particular environment to have a well-formulated problem, indeed that is what we establish in Section 3. However, in that section we also argue that as long as there is some lack of self-control, no matter how small, the optimal tax profile is unique. An additional problem is that lump-sum taxes are not feasible for policymakers, so it is hard to think of a good application, but we discuss the case for age-dependent pension contributions in the conclusion.

Nevertheless, our paper also contributes to three strands of the literature. The first one would be household liquidity across the life cycle. We are motivated to explain the empirical profile of HtM households by age outlined in Kaplan, Violante, and Weidner (2014), and covered briefly by Aguiar, Bils, and Boar (2020). An extensive literature features models that have theoretical and quantitative implications for life cycle liquidity, some notable examples are Huggett (1996), İmrohoroğlu (1998), Conesa, Kitao, and Krueger (2009) and Kaplan and Violante (2014), this last one including a portfolio decision between liquid and illiquid assets. However, none of them ask which features of the environment can explain the age pattern of HtM mentioned above. In the context of the household finance literature[5] Campanale, Fugazza, and Gomes (2015) study life cycle liquidity, but with a portfolio decision between liquid cash/bonds vs illiquid stocks. In doing so, they omit an explicit modeling of housing and retirement accounts, and they mix illiquidity with risky returns in a single asset (stocks), which is not the focus of our paper.

The second strand refers to issues of self control in consumption/savings. Two approaches have become popular for structurally modeling this phenomenon: the present bias (or quasi-hyperbolic discounting) framework featured in Laibson (1997) posits a high short-run discount rate coupled with a low long-run discount rate, generating time-inconsistency and over-consumption relative to the long-run discount rate. The second approach was developed by Gul and Pesendorfer (2001, 2004) who axiomatize preferences over menus and show a utility representation that has a commitment and a temptation component, hence we refer to this as the "temptation" framework. Either approach yields a commitment value to tighter borrowing constraints or illiquid assets: they help households to address their self-control issues by tying their own hands. Naturally then, this literature shares a large overlap with the study of liquidity, where Angeletos et al. (2001), Laibson et al. (2003) and Laibson, Repetto, and Tobacman (2007) are some prominent papers featuring a life cycle for the case of present bias, and they match some aspects of the age profile of liquidity, but we contribute by matching the age of

[5] See Gomes (2020) for a theoretical review of life cycle portfolio choice and Gomes, Haliassos, and Ramadorai (2021) for an empirical review of household finance, most of the literature focuses on the risky vs. risk-less asset portfolio choice.

profile of HtM households, which are the protagonists in our quantitative exercise. For temptation, Attanasio, Kovacs, and Moran (2020) match the age profile of wealthy HtM, so our exercise is closest to theirs, but we match the profile of poor HtM. Additionally, Attanasio, Kovacs, and Moran (2021) and Kovacs, Low, and Moran (2021) evaluate mortgage policies and estimate the strength of temptation respectively. To these two first strands of the literature, we contribute the study of age-dependent taxes as a way to reach the commitment solution.

There is also some literature on the optimal degree of commitment in the face of self-control issues from a principal-agent perspective: Amador, Werning, and Angeletos (2006), Beshears et al. (2020) and Moser and Olea de Souza e Silva (2019) characterize conditions under which mandatory savings are optimal. This relates to our paper because age-dependent taxes that keep the present value of disposable income constant are isomorphic to forced savings/borrowing. This literature deals mostly with the optimal allocation in stylized two-period models, while we contribute a rich quantitative framework to quantify the welfare gains of simple policy instruments, which complements their findings. Additionally, their focus is on consumption in working life vs retirement, while our quantitative exercise is carried out entirely in the former period.

The third strand entails age-dependent taxes, which are a natural tool to attain the optimal disposable income profile. Most of the literature on this policy follows the Mirrlees (1971) approach, for example Weinzierl (2011), Farhi and Werning (2013), Golosov, Troshkin, and Tsyvinski (2016) and Stantcheva (2017). These papers are concerned with the age-specific profile of productivity and labor supply elasticity, with stylized frameworks for consumption/savings. A Ramsey approach allows us to consider a richer problem which includes HtM households. This is the approach that Karabarbounis (2016), da Costa and Santos (2018) and Heathcote, Storesletten, and Violante (2020) follow, featuring models with uninsurable income shocks, increasing income profiles and relevant borrowing constraints while also studying age-dependent taxes, and so they consider a pressure to increase liquidity for the young HtM households.[6] However, these papers focus on the changing incentives that these taxes imply, primarily age-varying labor supply elasticity and human capital accumulation. Krusell et al. (2010) and Pavoni and Yazici (2017) do consider the role of temptation and present bias (respectively) in age-dependent capital taxation, but under a completely different mechanism: they consider deterministic OLG models where each cohort has a representative agent and perfect capital markets, and distortionary taxes correct the under-saving distortion caused by self-control issues. We can conclude that our paper contributes a novel channel in which age-dependent taxes affect welfare, and that optimal liquidity redistribution is only well-framed in our setting, where incomplete markets are coupled with self-control issues.

3 Deterministic Setting

We will first study the environment without overconsumption issues, which we add later. Consider a household that lives for T discrete periods and maximizes its discounted sum of utility derived from consumption.

[6] Findeisen and Sachs (2017) and Ndiaye (2017) have the first two features, but they use the natural borrowing limit and thus don't feature HtM households.

$$\max_{\{c_t\}_{t=1}^T} \sum_{t=1}^{T} \delta^t u(c_t) \tag{1}$$

$$\text{s.t.} \quad c_t + a_{t+1} = Ra_t + y_t + \tau_t \quad \forall t \in \mathbb{T} \tag{2}$$

$$a_{t+1} \geq \underline{a} \quad \forall t \in \mathbb{T} \tag{3}$$

Where $\mathbb{T} = \{1, ..., T\}$. We shall henceforth call (1) the *stoic* preferences, because they are devoid of self-control issues. The utility function is well behaved, meaning that $u : \mathbb{R}_+ \to \mathbb{R}$ is twice continuously differentiable, strictly increasing and strictly concave. This household discounts future periods with a factor $\delta \geq 0$ and has access to saving and borrowing in an asset a_{t+1} at gross rate of return $R \geq 0$, with a borrowing constraint at $\underline{a} \leq 0$ every period.[7] It also starts off with some exogenous level of assets $a_1 \geq \underline{a}$. Its income profile $\{y_t\}_{t=1}^T \geq 0$ is exogenous, and a benevolent government determines the the lump-sum stream of taxes and transfers $\{\tau_t\}_{t=1}^T$, which are restricted to have zero present value:[8][9]

$$\sum_{t=1}^{T} \frac{\tau_t}{R^t} = 0 \tag{4}$$

This completes the definition of the environment, and now we ask ourselves the optimal $\{\tau_t\}_{t=1}^T$ policy, which we shall subsequently call the tax profile. Consider first the case of complete financial markets, which removes the borrowing constraints. In such a case, we can use the intertemporal budget constraint to characterize the feasible consumption sequences:

$$\sum_{t=1}^{T} \frac{c_t}{R^t} = Ra_1 + \sum_{t=1}^{T} \frac{y_t}{R^t} + \underbrace{\sum_{t=1}^{T} \frac{\tau_t}{R^t}}_{=0}$$

Trivially, the tax profile has no effect on allocations because it does not affect the choice set. Henceforth, the *optimal allocation* will be that which is attained by households with stoic preferences in the complete markets setting.

Consider now the case of incomplete markets, where $\underline{a}$ is some value higher than the natural borrowing limit. We must now consider every period independently because the borrowing constraint may bind. The solution is characterized by the stoic Euler equation:

$$u'(c_t) \geq \delta R u'(c_{t+1}) \tag{5}$$

[7] The results can be generalized to the case where the discount factor, interest rate and borrowing constraint vary across age, but this doesn't add more intuition at the cost of heavier notation, therefore we cover the simpler case where they are constant. Also, the current environment allows households to die with some debt if $\underline{a} < 0$, but the results are trivial to extend to the case where this is ruled out.

[8] One could be inclined instead to think of restricting tax profile to be revenue-neutral: $\sum_{t=1}^{T} \tau_t = 0$. However, notice that front-loading subsidies and back-loading taxes would increase the present value of disposable income of households, so this would not merely be a liquidity redistribution, which is our object of interest.

[9] The results generalize easily to the case where the present value of taxes is a scalar, with no restriction on its sign. We choose to cover the zero present value case because then we think of these taxes as purely redistributing liquidity.

Where $u'(c) \equiv \frac{\partial u(c)}{\partial c}$. This is a weak inequality because if the borrowing constraint is binding, the household cannot increase present consumption at the expense of future consumption, so current marginal utility might be strictly higher than the right hand side. However, there exists a knife-edge case in which the borrowing constraint is binding, but the household still achieves equality of the Euler. It will be useful to define these cases.

Definition 1. A household is *Hand-to-Mouth (HtM)* at age t if it chooses $a_{t+1} = \underline{a}$. Additionally, it is *constrained HtM* if the Euler equation holds with a strict inequality. If it is HtM but the Euler holds with equality, it is an *unconstrained HtM* household.

Intuitively, a benevolent government will try to allow households to smooth consumption as much as possible, which is given by the equality of the Euler equation. This leads us to the first result.

Proposition 1. *With incomplete markets, any $\{\tau_t\}_{t=1}^{T}$ for which there are no constrained HtM households achieves the optimal allocation.*

Proof. The optimal allocation is characterized by Euler equations (5) for $t \in \{1, ..., T-1\}$ holding with equality and the intertemporal budget constraint. If there are no unconstrained HtM in the incomplete markets setting, then all Eulers hold with equality. Additionally, the intertemporal budget constraint must always hold, because it is a discounted sum of the period budget constraints. ∎

The reason we care about Proposition 1 is because there is no unique tax profile that achieves the optimal allocation, indeed the set could be very large. Additionally, the presence of unconstrained HtM is not problematic, it's the constrained HtM that are not smoothing consumption as much as they could. Also, notice that it suggests a corner solution: giving subsidies to younger households that they pay in the form of taxes when they are older may benefit them in terms of lifetime utility, and can't possibly hurt them. This is because they can always save to undo the tax profile's liquidity redistribution, but sometimes they can't borrow to solve the opposite problem. Finally, it is apparent here that given an exogenous income profile $\{y_t\}_{t=1}^{T}$, the choice of a tax profile $\{\tau_t\}_{t=1}^{T}$ is isomorphic to the choice of a disposable income profile $\{y_t + \tau_t\}_{t=1}^{T}$, as argued in the introduction. Therefore we can say that a benevolent government wants to front-load the disposable income profile in this setting.

3.1 Present Bias

Now we consider the liquidity redistribution problem in the presence of quasi-hyperbolic discounting as in Laibson (1997), popularly known as present bias. Consider a household that is still subject to (2) and (3), but whose preferences instead maximize, for every $t \in \mathbb{T}$:

$$u(c_t) + \beta \sum_{s=t+1}^{T} \delta^{s-t} u(c_s)$$

For $\beta \in (0,1)$. We can think of this as the problem of the t self of a given household, which is potentially conflicting with the decisions that a $t-1$ self of that same household might take, because preferences "shift" as it ages: discount rates between any two given time periods are not constant over time, as in the case of exponential discounting (we called them stoic preferences earlier, they are nested by this framework when $\beta = 1$). We have to take a stance on the awareness that different selves have regarding the preference shift: as is standard in the literature we denote fully aware households as sophisticated, and unaware households as naive.[10]

3.1.1 Sophisticated Households

These households know that their preferences will shift in the future, therefore in a deterministic environment consumption sequences are anticipated correctly. We extend Laibson (1996)'s result of a generalized Euler equation to an incomplete markets setting, which we derive in Appendix A.1.[11] We will call it the sophisticated Euler. The t self of a sophisticated household will choose consumption at t according to:

$$u'(c_t) \geq \underbrace{\delta \left[1 - \frac{\partial c_{t+1}}{\partial y_{t+1}} (1-\beta)\right]}_{\equiv \hat{\delta}_{t+1}} Ru'(c_{t+1}) \tag{6}$$

(6) pins down the actual decision that present-biased sophisticated households would take, as it maximizes their own welfare. We can think of $\hat{\delta}_{t+1}$ as the effective discount factor: it is the δ under which a stoic would make the same consumption/savings decision. Notice that $\hat{\delta}_{t+1} < \delta$ for $t \in \{1, ..., T-1\}$, because $\beta \in (0,1)$ and the MPC is strictly positive, meaning that this household will act more impatiently than if $\beta = 1$, which would produce the stoic Euler in (5).

However, a $t = 0$ self for this household would only have a discount factor of δ when making intertemporal decisions between $\{t, t+1\} \in \mathbb{T}$, because they are all in the future. Therefore, this 0 self would prefer, if it could, to choose consumption sequences according to the stoic Euler. Indeed, the discounted sum of utility of the 0 self is an affine transformation of the sum of utility discounted exponentially, given by (1). Therefore they represent the same preference ordering over consumption sequences as stoic preferences. We argue that it is natural to consider a benevolent government that wants to maximize only the welfare of these 0 selves, also known as the long-run criterion.[12] We'll call this a paternalistic government, because its choices might conflict with the desires of the $t > 0$ selves. We are now ready to state the main result of this section.

Proposition 2. *With incomplete markets and sophisticated present bias, there is a unique* $\{\tau_t\}_{t=1}^T$ *that maximizes* 0 *self welfare, and it achieves the optimal allocation. The condition that characterizes this unique tax profile is that all households are unconstrained HtM from the perspective of* 0 *selves.*

Proof. First, as argued above, whatever tax profile maximizes 0 self welfare will also maximize stoic welfare because the utility functions represent the same preferences, therefore it will achieve the optimal allocation.

Next we will show existence of a tax profile that achieves it. Define $\tilde{c}_1 \equiv y_1 + Ra_1 - \underline{a}$ and $\tilde{c}_t \equiv y_t + (R-1)\underline{a}$ for $t \in \{2, ..., T\}$, this is the consumption sequence of a household that lives in an environment with no taxes and decides to max out its borrowing limit every period. Consider the tax

[10] See O'Donoghue and Rabin (1999) for a discussion.

[11] We have not yet checked whether this case exhibits the multiple equilibria that may arise with sophisticated present bias, which Cao and Werning (2016) discuss, so this subsection should be considered as a heuristic statement. However, our argument holds rigorously with naive present bias and temptation preferences, and that accomplishes our goal for this section, which is discussed in Section 4.

[12] For a general discussion, see Bernheim and Taubinsky (2018) section 2.2.5. We also discuss this below in section 3.3

profile that solves the following conditions:

$$u'(\tilde{c}_t + \tau_t) \quad = \quad \delta R u'(\tilde{c}_{t+1} + \tau_{t+1}) \quad \forall t \in \{1, ..., T-1\} \tag{7}$$

Plus zero present value, as in (4). Under this tax profile, conjecture that the household chooses $c_t = \tilde{c}_t + \tau_t \ \forall t \in \mathbb{T}$. To verify, just notice that the t self of this household is constrained HtM every period, because $\hat{\delta}_{t+1} < \delta$ so the sophisticated Euler holds with strict inequality, and borrowing is maxed out, therefore it effectively chooses this consumption sequence. Then, this tax profile implements the 0 self Euler equation with equality at all times, so they are unconstrained HtM from the perspective of the 0 self. This Euler is equal to that of the optimal allocation, therefore this tax profile achieves it, using the same arguments as in Proposition 1.

To show that this tax profile is unique, consider a different tax profile $\{\tau_t'\}_{t=1}^T$. If there is any period t for which the present self is unconstrained, then it cannot achieve the optimal allocation because the sophisticated Euler would hold with equality, preventing the exponential Euler to hold with equality. Therefore, we should only consider allocations in which $\{\tau_t'\}_{t=1}^T$ implies $c_t = \tilde{c}_t + \tau_t$ as above. Consider now $\tau_s' > \tau_s$ for some period $s \in \mathbb{T}$. Notice that an optimal allocation should have $u'(\tilde{c}_t + \tau_t') = \delta R u'\left(\tilde{c}_{t+1} + \tau_{t+1}'\right)$ for $t \in \{1, ..., T-1\}$. Because $u'()$ is a strictly decreasing function this implies $\tau_t' > \tau_t$ for all $t \in \mathbb{T}$, which violates zero present value. ■

Proposition 2 gives us a sharper result than its stoic counterpart: the optimal tax profile is unique and it is an interior solution. The intuition is that, in the presence of sophisticated present bias, excessive present liquidity (relative to future liquidity) can be harmful from the perspective of the long-run criterion, because households might overconsume. Therefore, the paternalistic government wants to provide just the right amount of liquidity such that the 0 self can smooth consumption, while future selves cannot undo the desired consumptions sequence because they are constrained. This is the same allocation that would be achieved in the presence of commitment devices (such as some illiquid assets), so we can view borrowing constraints as a feature that allows for commitment in the presence of a benevolent government that can redistribute liquidity, even in the absence of explicit commitment devices. Finally, note that the tax profile pinned down by (7) does not depend on β, so this result holds for any strength of the present bias.

3.1.2 Naive Households

We will show that Proposition 2 also applies to a setting with naive agents. For these households, consumption sequences are not correctly anticipated, even in a deterministic setting, because they are surprised by their preference shift every period. That is, they believe they will not be present-biased in the future, which turns out to be false. In particular, let c_{t+1}^e be the consumption that would be chosen by a stoic household at time $t+1$, which the naive households expects at t. Then, the naive Euler for deciding is given by:

$$u'(c_t) \geq \beta \delta R u'\left(c_{t+1}^e\right) \tag{8}$$

Notice that the 0 self Euler is the same as in the sophisticated case, which is the stoic Euler. By the same arguments as above, we consider a government that wants to maximize 0 self welfare, which leads us to the same result.

Proposition 3. *With incomplete markets and naive present bias, there is a unique* $\{\tau_t\}_{t=1}^T$ *that maximizes* 0 *self welfare, and it achieves the optimal allocation. It is the same one as in Proposition 2: all households are unconstrained HtM from the perspective of* 0 *selves.*

Proof. Consider the same tax profile as in Proposition 2, given by (7) and zero present value. Conjecture that the household chooses $c_t = \tilde{c}_t + \tau_t \ \forall t \in \mathbb{T}$, and expects to do so, meaning $c_t^e = c_t \ \forall t \in \mathbb{T}$. To verify, first notice that the t self of this household is constrained HtM every period, because $\beta < 1$ so the naive Euler holds with strict inequality, and because borrowing is maxed out, it effectively chooses this consumption sequence. Second, notice that the 0 self expects all futures selves to be unconstrained, because the stoic Euler holds with equality, so the conjecture is verified. Optimality and uniqueness follow using the same arguments as in Proposition 2. ■

The intuition is similar, the paternalistic government wants to give just enough liquidity and no more, because it knows that households have a self-control problem. However, expected effects of the policy differ among levels of awareness: the sophisticated 0 self appreciates that liquidity will never be excessive, because it knows that it helps it to prevent overconsumption. The naive 0 self doesn't think it will ever suffer from excess liquidity, so it is indifferent toward back-loading of disposable income profiles that the government could want to implement, which would never constrain a stoic household in any case. Nevertheless, Propositions 2 and 3 jointly let us conclude that with present bias, independently of the degree of awareness, the 0 selves of these households could appreciate and would never be contrary to the optimal tax profile.

3.2 Temptation and Self-Control

We introduce another kind of preferences that can induce overconsumption: dynamic self-control preferences à la Gul and Pesendorfer (2004). Under this setting households are exponential discounters and thus time-consistent, but they are not stoic because if the household has a "tempting" option in the choice set, it must exert costly self-control to refrain itself from choosing it. An example would be going on a shopping spree and maxing out the credit card limit. This is what we will model: the limiting formulation in which the tempting option is consuming as much as the borrowing constraint allows, which is what an infinitely impatient household would do. We additionally choose the temptation utility as Kaplan and Violante (2022) do, which is parametrized by a cost $\varphi > 0$. Thus, subject to the same constraints (2) and (3), households maximize:

$$\max_{\{c_t\}_{t=1}^T} \sum_{t=1}^{T} \delta^t \left(u\left(c_t\right) - \varphi\left[u\left(\hat{c}_t\right) - u\left(c_t\right)\right]\right) \tag{9}$$

$$\hat{c}_t = Ra_t + y_t + \tau_t - \underline{a}$$

The term in square brackets is the cost of self-control, because it's the difference in utility between the tempting alternative ($\hat{c}_t$ is the consumption that comes from $a_{t+1} = \underline{a}$) and the chosen level of consumption. When the household chooses to max out its borrowing limit, this term is equal to zero. Also, when $\varphi = 0$ we get the stoic preferences back.

A standard first order condition gets us the temptation Euler:

$$u'(c_t) \geq \underbrace{\delta\left[1-\left(\frac{\varphi}{1+\varphi}\right)\frac{u'(\hat{c}_{t+1})}{u'(c_{t+1})}\right]}_{\equiv\hat{\delta}_{t+1}} Ru'(c_{t+1}) \tag{10}$$

In this case, the effective discount factor is guided by φ and the ratio of marginal utilities.[13] As before, $\hat{\delta}_{t+1} < \delta$, which comes immediately from $\gamma > 0$ and an increasing utility function. Unlike the case with present bias, here we do not need to make any assumption about whose welfare we want to maximize because there is no shift in preferences.

Proposition 4. *With incomplete markets and costly temptation, there is a unique* $\{\tau_t\}_{t=1}^{T}$ *that maximizes* (9)*, and it achieves the optimal allocation. It is the same one as in Proposition 2: Households are unconstrained HtM from the perspective of a stoic household.*

Proof. Denote the discounted sum of utility of a household afflicted by costly temptation as:

$$\hat{U}(c_1, ..., c_T) = \sum_{t=1}^{T} \delta^t \left(u(c_t) - \varphi\left[u(\hat{c}_t) - u(c_t)\right]\right)$$

Analogously, denote the case of $\varphi = 0$ as:

$$U(c_1, ..., c_T) = \sum_{t=1}^{T} \delta^t u(c_t)$$

Then, note that because of the borrowing constraint, $\hat{c}_t \geq c_t \ \forall t \in \mathbb{T}$ and it must be that:

$$U(c_1, ..., c_T) \geq \hat{U}(c_1, ..., c_T) \quad \forall \{c_1, ..., c_T\} > 0 \tag{11}$$

This means that the optimal allocation, which maximizes U, gives an upper bound for $\hat{U}$. Now, consider the same tax profile as in Proposition 2, given by (7) and zero present value. Again because $\hat{\delta}_{t+1} < \delta$ in (10), the $\varphi > 0$ household chooses $c_t = \tilde{c}_t + \tau_t \ \forall t \in \mathbb{T}$ and is constrained every period. Because it is constrained, $\hat{c}_t = c_t$ and

$$\hat{U}(c_1, ..., c_T) = U(c_1, ..., c_T)$$

Because $\{c_t\}_{t=1}^{T}$ makes all $\varphi = 0$ Euler equations hold with equality (making households unconstrained HtM from this perspective), we know that it maximizes U. Since this is the upper bound, the optimal allocation is achieved and it maximizes $\hat{U}$. Uniqueness follows using the same argument as in Proposition 2. ■

This result holds for any $\varphi > 0$, so again the strength of temptation does not matter. Another similarity with the present bias case is that a paternalistic government will want to give just enough liquidity and no more, but it's a different kind of paternalism. In the case of costly temptation,

[13] This ratio of marginal utilities is a monotonic function of an average propensity to consume under CRRA preferences.

the tension is not between different selves of the same household, but it comes from the fact that households directly prefer to be constrained from being able to choose a tempting alternative, and the borrowing constraint coupled with the tax profile provide the desired constraint, while still allowing for consumption-smoothing. Therefore, a household will appreciate the optimal tax profile at any time during the life cycle, because it is a form of commitment, and the government is paternalistic because it provides that commitment by constraining the household's choice set. In contrast, with present bias, the 0 self appreciates the optimal tax profile, but the $t = 1$ self would prefer a different one, because preferences have shifted. In this sense, present bias presents a conflict between the government and the household that is absent in the case of costly temptation, which we interpret as a harder form of paternalism.

3.3 Discussion

Propositions 2 - 4 can be interpreted according to a general principle: when households have overconsumption issues, excessive liquidity can be harmful and an optimal tax profile, coupled with borrowing constraints, can overcome this. These results can probably be extended more generally,[14] but these specific cases are illustrative of the forces that justify an optimal tax profile that induces liquidity redistribution. They are also illustrative of how households might react to the policy, especially in the case of present bias, as we have discussed above.

There is one interpretation of our propositions that we want readers to avoid and is worth mentioning: it's true that every household is Hand-to-Mouth with an optimal tax profile in our deterministic setting, but our results do **not** suggest that we should increase the number of HtM households in the economy. Rather, the message is that we should distinguish between constrained and unconstrained HtM from the perspective of a household without self-control issues. In our framework, policy can help with the former HtM, yet not with the latter. And it does so by redistributing available liquidity to the stage of the life cycle where it is most needed. This can reduce the overall level of liquidity in the economy by increasing the amount of HtM households, but because consumption is smooth through the life cycle according to the stoic Euler, a carefully chosen optimal policy manages to do no damage with this.

We have emphasized that the optimal tax profile does not depend on the strength of the self-control problems. The reason for both models is that self-control does not matter in a commitment solution (which would be chosen by the 0 self in the case of present bias, and all selves in the case of costly temptation), and the tax profile coupled with borrowing constraints provide such commitment. We view this as a desirable result because both models nest stoic preferences as a limiting case, therefore they can be viewed as a generalization, and if there is a even a small amount of self-control issues, the results kick in. Further, a policymaker does not need accurate measures of the degree of self-control problems to know the optimal policy.[15]

We interpret the fact that the tax profile is the same in Propositions 2 - 4 as support for our

[14] Obvious directions are partial naiveté in present bias, and a tempting option that is moderately more impatient, instead of infinitely so for the case of costly temptation. We believe both are very likely to produce the same results, these can be studied in future work.

[15] This statement does not carry over rigorously to the stochastic setting, we discuss this below. However, we believe the broad intuition still applies.

choice of maximizing the welfare of 0 selves in the present bias setting. Other choices could be made, such as maximizing the welfare of a weighted sum of the different selves, or characterizing the Pareto efficient allocations. However, these choices would likely yield different tax profiles. Our choice gives us a consistent policy recommendation that is independent of the underlying issue causing overconsumption, be it present bias or costly temptation, and connects it with stoic preferences.

Finally, as we said in the introduction, the key object of interest in this paper, liquidity, is chosen in great part to be able to cope with idiosyncratic uninsurable income risk, as the large heterogenous-agent macroeconomics literature has shown. Therefore, the next section presents a model that incorporates this feature. However, we view the current section as providing the theoretical background and the intuition for one of the key concepts that will operate in the rest of this paper: we want to give each stage of the life cycle just enough liquidity, but no more than that.

4 Quantitative Model

4.1 Environment

The individual households are the protagonists of this analysis, therefore we will describe their life cycle problem in detail. We only consider a simple partial equilibrium for the calibration of the model, and we will describe it then. We choose to model their life cycle in continuous time, mostly because the instantaneous-gratification framework of Harris and Laibson (2013) operationalizes the continuous time limit of quasi-hyperbolic discounting in a way that guarantees a unique equilibrium and features a single welfare criterion.[16] In addition, we use standard techniques for solving heterogeneous-agent models in continuous time, which have many advantages over their discrete-time counterparts, as documented by Achdou et al. (2022) in an infinite horizon setting, and used by Weil (2018) in a life cycle application.

Finally, for the time being this quantitative exercise will consider just the case of naive present bias under two arguments. First, Section 3 shows us that the main mechanism that we want to quantify works identically (when compared to sophistication or temptation) in the deterministic setting and this makes us hopeful that it works similarly in a stochastic one. Second, it is the technically easiest to implement as it just entails solving a standard Hamilton-Jacobi-Bellman equation.[17] We now proceed to describe the environment in detail.

Households are born at $t = 0$ and die at rate $\theta(t)$ before T, and then deterministically at T. They have stochastic income $y(t)$ over their life, the specific formulation for which will be described in Section 4.3, and they face age-dependent lump sum taxes and transfers $\tau(t)$. They can consume $c(t)$ and get flow utility $u(c)$, which has the same properties as in Section 3: twice continuously differentiable, strictly increasing and strictly concave. Alternatively, they can save their income in a risk-free and liquid asset $a(t)$, with interest rate r. They are born with a_0 assets and face a borrowing

[16] Present bias models in discrete time are known to feature multiple equilibria, see Cao and Werning (2016) for a discussion and more references.

[17] For the case of sophisticated present bias, see Maxted (2021) for a discussion of the violation of the Barles and Souganidis (1991) conditions that guarantee convergence. Laibson, Maxted, and Moll (2021) study an application of naive present bias, our setup is similar to theirs but the absence of a discrete decision (mortgage refinancing in their case) makes our formulation much simpler.

constraint of $\underline{a}(t)$. Thus, they are subject to the following constraints:

$$\begin{aligned} c(t) + \dot{a}(t) &= ra(t) + y(t) + \tau(t) \\ a(t) &\geq \underline{a}(t) \\ a(0) &= a_0 \geq \underline{a}(0) \\ y(0) &= y_0 \end{aligned}$$

We model households having naive instantaneous gratification, where naiveté has the same meaning as in Section 3. A household of age t discounts the future by a factor $D(s)$ for $t' \geq t$ where:

$$D(t') = \begin{cases} 1 & \text{if } t' = t \\ \beta e^{-\rho t'} & \text{if } t' > t \end{cases}$$

The factor D is the continuous-time limit of the $\beta - \delta$ discounting structure that we covered in Section 3.1 under quasi-hyperbolic discounting. Indeed, now the present only lasts for the instant in which $t' = t$, and for any future -no matter how close- we discount with the factor β as before. The discount rate ρ is meant to capture discounting between two future periods, same as δ. It's easy to notice that exponential discounting is nested by the case of $\beta = 1$.

Naiveté implies that households believe they will be exponential discounters from the next instant onwards, so the time-consistent problem is important to formulate. Let $s \equiv (a, y, t)$ be the state vector. We can represent the $\beta = 1$ problem recursively with a Hamilton-Jacobi-Bellman (HJB) equation for $t \in (0, T)$:

$$(\rho + \theta(t)) v(s) = \max_c \{u(c) + \partial_a v(s) \dot{a} + \mathcal{A}_y v(s) + \dot{v}(s)\} \tag{12}$$

$$\begin{aligned} \text{s.t.} \quad c + \dot{a} &= ra + y + \tau(t) \qquad (13) \\ a &\geq \underline{a}(t) \end{aligned}$$

Where ∂_x refers to the partial derivative with respect to some variable x, $\dot{x} \equiv \partial_t x$, and $\mathcal{A}_y$ is the infinitesimal generator of the stochastic process for income.[18] Because $v(s)$ is a value function with exponential discounting, it is the continuation-value function for naive households. Their current-value function, which characterizes their expected utility flows into the future, is simply given by:

$$\hat{v}(s) = \beta v(s) \tag{14}$$

However, one must be careful with the interpretation of (14), it does not mean that the present-biased agent differs in just affine transformations of all utility flows (which would yield the same policy functions). It means that the continuation value is weighted by a factor β. This leads to the first order condition that pins down consumption/savings when the borrowing constraint does not bind:

$$u'(c(s)) = \beta \partial_a v(s)$$

[18]See Achdou et al. (2022) for some discussion on these kind of operators, and Section 4.3 for the HJB with a specific income process.

Where $u'(c) = \partial_c u(c)$. With exponential discounting, this condition would be identical except for the absence of β, which implies that present-biased households will overconsume due to $u''(c) < 0$. This is proven by Laibson, Maxted, and Moll (2021) heuristically and by Harris and Laibson (2013) rigorously, and if we follow their approach and assume $u'(c)$ is invertible, we can fully characterize the naive consumption policy:

$$c(s) = \begin{cases} u'^{-1}\left(\beta \partial_a v(s)\right) & \text{if } a > \underline{a}(t) \\ \min\left\{u'^{-1}\left(\beta \partial_a v(s)\right), r\underline{a}(t) + y + \tau(t) - \dot{\underline{a}}(t)\right\} & \text{if } a = \underline{a}(t) \end{cases} \tag{15}$$

This is all consistent with standard continuous-time consumption/savings with incomplete markets, except for the last term in the case of $a = \underline{a}(t)$: if the borrowing constraint is time-varying, one needs to account for this variation to specify the consumption of persistently HtM households, which will be able to consume less (more) if the borrowing constraint is becoming more (less) stringent over time.

Once we have the consumption policy function we can calculate the welfare of the present-biased household. We need to notice that this agent is naive, and so $\hat{v}(s)$ is not the correct measure of welfare, because this value function assumes exponential discounting in the future. Instead, one can define implicitly the welfare value function as:

$$(\rho + \theta(t))\, w(s) = u(c(s)) + \partial_a w(s)\, \dot{a}(s) + \mathcal{A}_y w(s) + \dot{w}(s) \tag{16}$$

Where $c(s)$ is the effective consumption policy, and $\dot{a}(s)$ is its implied savings policy, which follows from (13). Notice this functional equation is devoid of any optimization, it just captures an expected present value, where the dynamics are given by the exogenous processes of time and income and the endogenous choice (15) that solves another problem. This is analogous to the 0 self welfare that we studied in Section 3.1, indeed notice that (16) only features the exponential discount rate ρ and no present bias, except through the dynamics of consumption and savings.

4.2 Functional Forms

We assume households have CRRA utility functions with relative risk aversion γ over consumption flows:

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

We assume an income process with a deterministic and a stochastic component:

$$\log y(t) = x(t) + z(t)$$

The deterministic component $x(t)$ is meant to capture the expected income profile during the life cycle, and so it is an exogenous and perfectly known function to the household. For now, we use a quadratic polynomial:

$$x(t) = 1 + \phi_1 t + \phi_2 t^2$$

The stochastic component $z(t)$ captures unexpected idiosyncratic and uninsurable shocks to income. We describe it in discrete time and implement it in continuous time. Suppose that we observe periodic snapshots of productivity at age t given by z_t, then for all $t > 0$:

$$z_t = \rho_z z_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim N(0, \sigma_z)$$

And for $t = 0$, $\varepsilon_0 \sim N(0, \sigma_{z_0})$. We discretize this process through a Rouwenhorst approximation with N_z states, which gives us a Markov transition matrix $1 + \Lambda$. In the continuous time limit, this becomes a Poisson process where the density of the income states g_z evolves according to $\dot{g}_z(t) = \Lambda' g_z(t)$. In particular, Λ is conformed by the transition rates, where $\lambda_{ii'}$ is the rate from state z_i to $z_{i'}$.

For the time being we will assume a constant no-borrowing constraint $\underline{a}(t) = 0$, future work will relax this assumption. With this structure we can define the new vector of states as $s = (a, z, t)$ and rewrite (12) to expand the infinitesimal generator:

$$(\rho + \theta(t)) v(s) = \max_c \left\{ \frac{c^{1-\gamma}}{1-\gamma} + \partial_a v(s) \dot{a} + \sum_{z_{i'} \neq z_i} \lambda_{ii'} \left[v(a, z', t) - v(a, z, t) \right] + \dot{v}(s) \right\} \tag{17}$$

$$\text{s.t.} \quad c + \dot{a} = ra + \exp(x(t) + z) + \tau(t)$$

$$a \geq 0$$

Appendix B outlines the algorithm to solve (17) numerically. A simplified version of this same algorithm is used to solve for (16).

4.3 Calibration

We set the initial time $t = 0$ as 22 years old, consistent with the Kaplan, Violante, and Weidner (2014) data on HtM households, and we normalize the size of a cohort to 1 at this age. Demographics follow Weil (2018), who uses the Human Mortality Database to calibrate mortality rates $\theta(t)$ according to 2000 U.S. data. We set the time period to be one year, and choose $T = 82$ as the terminal time for households. This means we allow households to be up to 104 years old, but this happens with only a very low probability given demographics data, in particular less than 0.1% of the cohort is alive by 104. Figure 2 shows the evolution of $\theta(t)$ and the cohort size. In this way, we think of the distribution of households in this economy as overlapping generations which share no intergenerational linkages, and operate in partial equilibrium. We will only use this stylized definition of demographics to calculate a % of HtM which we target below.

We choose a relative risk aversion coefficient of $\gamma = 2$, a standard value in the literature. For purposes of the calibration we set $\tau(t) = 0$, so our policy can be read as a difference over the prevailing tax scheme. The rest of the calibration is **very** preliminary, we expect much of it change, but it will get us some results for the time being.

The coefficients for the deterministic component of income are taken from Abbott et al. (2019), Figure 3 shows the deterministic profile. We use their coefficients throughout the life cycle (including retirement) because the resulting income profile does not look too different from our preliminary

Figure 2: Demographics

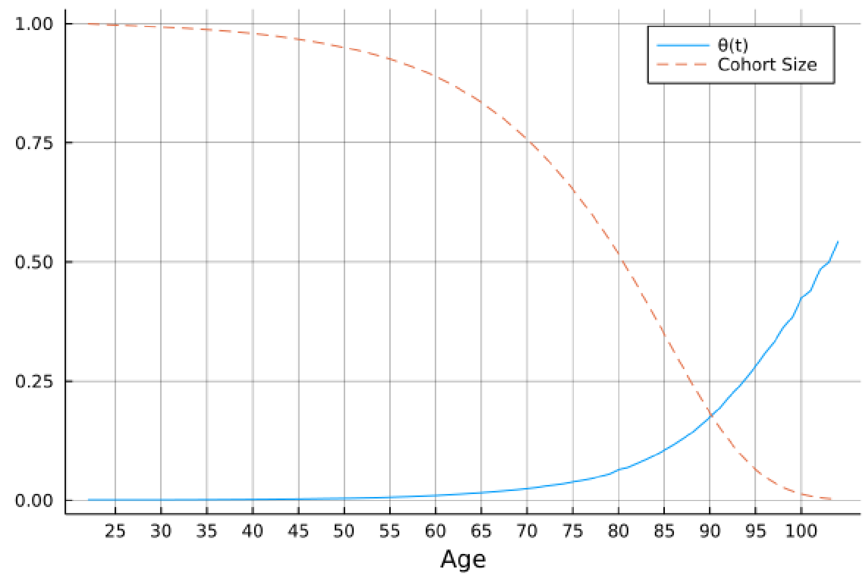

Source: Mortality rate $\theta(t)$ is taken from the Human Mortality Database. Cohort size is normalized to 1 at 22 years old, and its evolution is calculated with $\theta(t)$.

estimations with SCF data. The parameters for the stochastic income are taken from Attanasio, Kovacs, and Moran (2020) who assume the same process, they choose $\rho_z = 0.9$ and $\sigma_z = 0.05$. We choose σ_{z_0} by the ergodic distribution of the z process given ρ_z and σ_z. The initial distribution of assets is chosen to match the distribution of wealth in the SCF data that Kaplan, Violante, and Weidner (2014) use: we first set $a_0 = 0$ for 36% of the cohort to match the poor HtM that we observe at this age. For the other 64%, we assume their assets follow an exponential distribution where the mean is taken from the data.

In the absence of a credible way to structurally estimate β,[19] we set $\beta = 0.75$ following Maxted (2021) who argues that this is a conservative choice among the estimates in the literature. We set $r = 5\%$ as a standard value in the literature and $\rho = -0.0243$ is chosen to match the % of Poor HtM from the SCF data, which is 14.2% and we match exactly.[20] We choose to match the Poor HtM instead of all HtM (which includes those that are wealthy in illiquid assets) because we believe that they are closest to our object of analysis in a model without illiquid assets. An asset that is illiquid in a 1-year horizon -and is therefore relevant for the effect of some policy in this horizon- might be relatively liquid in a 10-year horizon, so it might prove somewhat useful in smoothing life cycle consumption. Therefore we are interested in households that don't even have this alternative: the Poor HtM.

[19]This could be done by allowing credit card debt in the model, as it is known since Laibson et al. (2003) that present bias allows simultaneous matching of data on illiquid assets and credit card debt.

[20]A negative ρ might be striking, and $r = 5\%$ might be considered relatively high as well. We expect these parameters to change as we work in the calibration, but for now note that the death rate is also an important source of discounting. Also, note that in models with a finite horizon, a positive discount rate is not strictly necessary.

Figure 3: Deterministic Income Profile

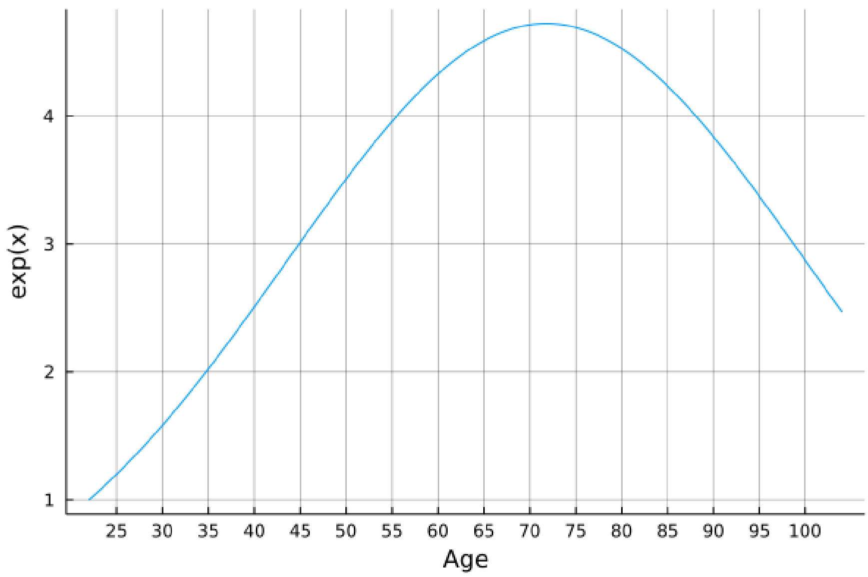

Source: Abbott et al. (2019).

4.4 Benchmark Economy

In Figure 4 we present the key statistic of our analysis: the fraction of HtM households by age, which we call the liquidity profile. We believe the model doesn't do a terrible job of matching it, but is still leaves much to be desired. In particular, the empirical profile seems to have a decreasing trend at all times, while the model-implied profile is increasing at the very beginning of working age, and some years after retirement age. In the latter case, the explanation could be directly tied to the death rate, which is a form of discounting that is more present at later stages in life. In this case, we believe a bequest motive might help to ensure a decreasing fraction of HtM in later stages of the life cycle, as it will give these households an extra saving motive. The former case might require more creative features, such as having children with some probability, in which case savings will come in handy (this can be operationalized with a stochastic discount rate).

Future work in this subsection will include more analysis into what are the essential ingredients to bring the model-implied profile closer to the data.

5 Results

5.1 A Simple Policy

We will now consider a tax profile $\tau(t)$ for $t \in [0, T_\tau]$, where $T_\tau < T$ will be the retirement age, at 65 years old. The reason why we only consider policy changes over working age households is twofold: first and foremost, we want to avoid any overlap with the very important question of the size of pensions and how much we should expect households to save for retirement. If our tax profile covered retirement age it would partly be answering that question, and while it could give an answer, the model is not

Figure 4: Liquidity Profile

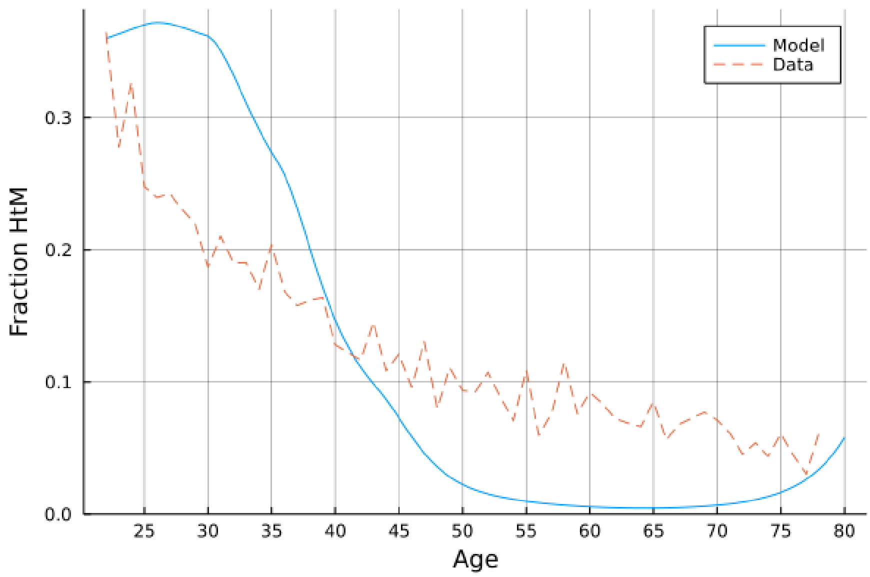

Source of data: Kaplan, Violante, and Weidner (2014)'s Poor HtM, based on the U.S. Survey of Consumer Finances.

tailored to do so, unlike a vast literature better suited for it. Therefore, we restrict our object of analysis to purely liquidity redistribution within working age. Second, it is closest to the application that we have in mind for future work, which entails age-dependent pension contributions. Therefore, for any policy experiment, we will have

$$\begin{aligned} \int_0^{T_\tau} e^{-rt}\tau(t)\,dt &= 0 \\ \forall t > T_\tau,\ \tau(t) &= 0 \end{aligned} \tag{18}$$

Where the first equation also ensures zero present value. Ideally we would like to calculate the full Ramsey optimal policy, which would optimize over the nonparametric function $\tau : [0, T_\tau] \to \mathbb{R}$.[21] We leave this exercise for future work and, for the time being we present a simpler exercise: a linear tax profile. We believe this is still an interesting exercise, as policymakers often prefer simple rules over complex fully-optimal policy. In this sense, it also has the advantage of being a simple comparison of liquidity needs between younger and "older" households, where the latter are still of working age but closer to retirement. This simple linear profile follows:

$$\tau(t) = \tau_0 + \tau_1 t \tag{19}$$

The combination of (18) and (19) mean that we only have one degree of freedom in choosing the policy parameters τ_0 and τ_1. The former is interpreted as the lump sum net transfers to 22-year olds, while the latter is the change in these payments per year. They are connected through the following

[21] We believe that the continuous time limit of the tax profile outlined in Section 3 could serve as guidance, or as an initial guess to be perturbed.

Figure 5: Welfare Change by τ_0

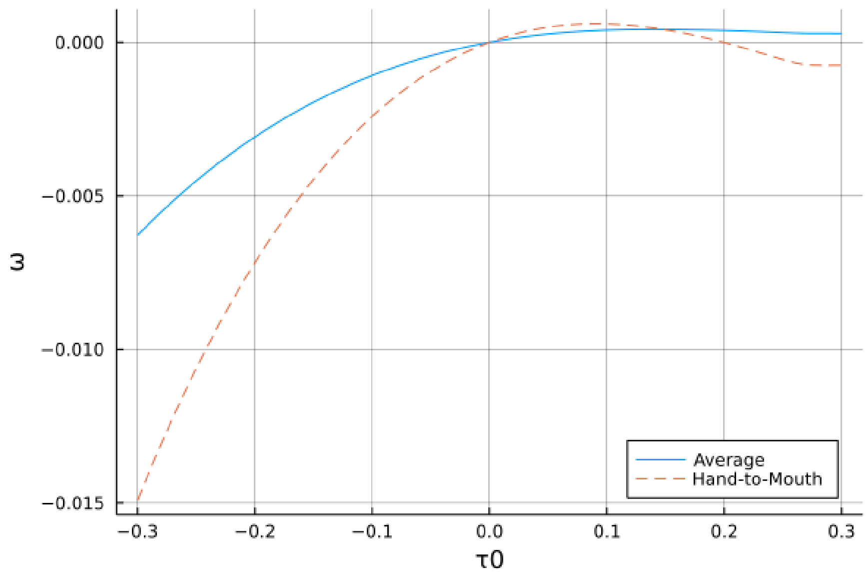

condition

$$\tau_1 = \frac{\tau_0 r \left(1 - e^{-rT_\tau}\right)}{\left(1 + rT_\tau\right) e^{-rT_\tau} - 1} \tag{20}$$

Which we derive in Appendix A.2. We will use equation (20) by exploring the space of τ_0 and its effects on welfare, while adjusting τ_1 to be consistent with zero present value. This policy can then be evaluated in terms of a consumption-equivalent gain with:

$$\omega\left(a,z\right) \quad = \quad \left(\frac{w^*\left(a,z,0\right)}{w\left(a,z,0\right)}\right)^{\frac{1}{1-\gamma}} - 1$$

Where $w\left(s\right)$ is the value function calculated with (16) of Section 4.4 which assumes $\tau\left(t\right) = 0$, while $w^*\left(s\right)$ is the same value function with the policy. This measure of welfare change is calculated only at the beginning of the life cycle, as we intend it to, and it can be computed for any starting level of assets or income, which is useful to consider the heterogeneous impact of this policy.

Figure 5 exhibits our main result (albeit preliminary): liquidity redistribution can be beneficial, up to a point. In the case of the distribution-weighted average of $\omega\left(a,z\right)$ exhibited in Figure 5, a 14% percent increase in disposable income at the beginning of the life cycle problem maximizes welfare by achieving a gain of 0.04% in consumption equivalent terms. This welfare gain is not very large, however it is an average across the population, where there might be heterogeneity in terms of what is exactly the optimal τ_0 depending on the initial states. Note also that while welfare gains are very small, there is an asymmetry in the sense that welfare losses from decreasing τ_0 are considerably larger, something that is relevant for policies which might decrease liquidity. This comes from being "to the left" of the optimal τ_0, as welfare seems to be concave in this parameter. Lastly, notice that this is just a simple linear policy, a fully Ramsey optimal function $\tau\left(t\right)$ will be able to achieve higher welfare (however, just how much is a quantitative question).

Figure 6: Welfare Change by τ_0, $\beta = 1$

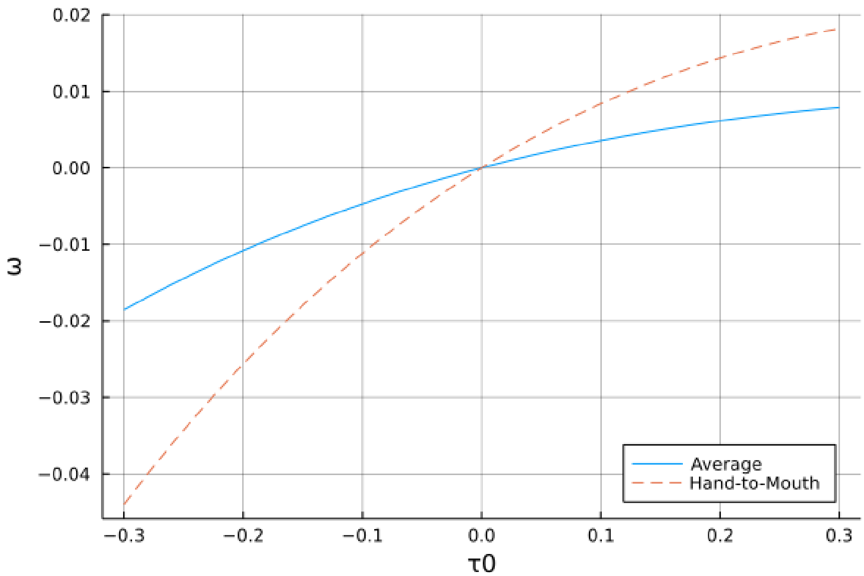

Figure 5 also shows the change in welfare from this same exercise for HtM households: those who have 0 assets when they are 22 years old (we further take the case of those whose z implies they have median income). We can see that welfare changes are substantially larger, and welfare as a function of τ_0 is more clearly concave. Indeed, the welfare-maximizing τ_0 increases disposable income to young households by 10% achieving a welfare gain of 0.06%. However, this is still quite small.

Finally, a note on these results: the simple intuition from Section 3 carries over naturally to this setting because it's still quite stylized, the dynamic problem is fairly similar across ages. There's much more variation of the drivers of liquidity through the life cycle that we could implement, e.g. it could be that younger households face more income risk, or that their discount rate ρ or their degree of present bias $1 - \beta$ are higher, or that borrowing constraints are relaxed as households age, etc...[22] This could easily overturn the present results, meaning that it could be optimal to actually decrease disposable income for young households in exchange for more liquidity when they are older.

5.2 The case without Present Bias

This section briefly shows that the message of Proposition 1 regarding liquidity redistribution without self-control issues also holds in our quantitative exercise. This is not hard proof that the proposition carries over to stochastic settings, but it is suggestive evidence and therefore a reassuring result. Figure 6 shows that a simple policy can only increase welfare by redistributing liquidity towards younger households, even in large amounts such as increasing their current income by 30%, and probably even more. This is true both in average and especially so for HtM households. Repeating the message of Proposition 1: without self-control issues, the incomplete markets setting is not enough to have a well-posed question of what is the optimal liquidity redistribution across the life cycle.

[22] Pavoni and Yazici (2017) discuss the evidence on age-dependent self-control problems, and implement them in a life cycle model to study capital taxation.

6 Conclusion

Our paper explores the idea of optimal liquidity redistribution across the life cycle in a setting with incomplete markets and self-control issues in consumption/savings. The main thought experiment is a set of age-dependent lump-sum taxes and transfers with zero present value, which shape household's disposable income profile. In a deterministic setting we derive sharp results characterizing the optimal tax profile, which achieves the complete-markets allocation by using borrowing constraints to prevent households from overconsuming, while at the same time smoothing their consumption across the life cycle. Then we move on to a stochastic setting where we build and calibrate a quantitative model to quantify welfare gains from a simple policy: linear tax profiles. The results and intuition from the deterministic setting carry over, however we find relatively small welfare gains from this policy experiment.

One of the reasons that we might be the first to study this specific thought experiment is because it's hard to come up with both relevant and direct applications that do not introduce complicated issues into the question. Lump-sum taxes are not feasible for policymakers, and with age-dependent labor income taxes, one needs to consider age-dependent labor supply elasticity and incentives to human capital accumulation. However, we believe a very natural application is age-dependent pension contributions. As long as agents internalize that a higher pension contribution still contributes to their wealth in a present value sense, labor distortions should not be as large as with labor income taxes, so most of the action should be in terms of liquidity.[23] We intend to work on this application in future work.

Empirically, our paper highlights the importance of leveraging the profile of HtM across age as a relevant statistic to infer whether there are potential welfare gains for life cycle redistribution of liquidity. Therefore, it calls for a deeper empirical analysis of this statistic, for example with regards to ex-ante heterogeneity. It's also interesting to note that the decreasing profile of HtM seen in the U.S. doesn't seem to be an immutable characteristic of modern economies, if we look at Germany or France in Figure 10 of Kaplan, Violante, and Weidner (2014) they seem quite flat. As discussed before, this does not directly imply that there are gains to liquidity redistribution in the U.S. but not in these other countries, deeper research into the drivers of liquidity needs across the life cycle is warranted to make this assessment.

We could write at great lengths about the future work to be done with this paper, but we will just mention some aspects we believe to be of significant economic importance. First and foremost, we have not mentioned illiquid assets in this paper. This is something we have wanted to incorporate from the beginning of this project, but we believe that the exact definition of illiquidity is paramount and not an easy choice. For example, Laibson (1997) shows how the commitment solution can be attained with quasi-hyperbolic preferences in the face of borrowing constraints and assets which are completely illiquid during the current period, but completely liquid when considering decisions between future periods. This yields our tax profile useless. However, we consider that a different formulation of illiquidity could be both more relevant to applications and yield more interesting results. But exactly

[23]Depending on the pension system some of the contributions might not benefit the taxpayer's pension directly, such as the Social Security contributions for higher incomes. This should obviously be considered when searching for a good application. The case of Chile is promising: it's mostly comprised of mandatory contributions of 10% of covered earnings, which are invested in individual accounts, so there is full internalization of contributions.

which formulation remains an important question.

Second, we should explore heterogeneity at the beginning of the life cycle, what we called ex-ante heterogeneity. We assumed a unique deterministic component of income for households, and we only considered the heterogeneity in the initial conditions and the subsequent evolution of states. However, we can observe in the SCF data how both the income profile and the HtM profile over age vary across level of education. This naturally implies different optimal tax profiles for these groups, which we could explore. We have some work in progress on this area, but much remains to be done.

Our third extension is a point made before: we need a better understanding of how the factors driving liquidity vary across the life cycle. This could mean variation in uninsurable income uncertainty, the discount rate, the strength of self-control issues, risk aversion, borrowing constraints, and also illiquid asset availability when we incorporate them. This should help in matching the empirical liquidity profile, while also shaping the optimal tax profile.

References

Abbott, B., G. Gallipoli, C. Meghir, and G. L. Violante (2019, December). Education Policy and Intergenerational Transfers in Equilibrium. *Journal of Political Economy 127*(6), 2569–2624. Publisher: The University of Chicago Press.

Achdou, Y., J. Han, J.-M. Lasry, P.-L. Lions, and B. Moll (2022, January). Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach. *The Review of Economic Studies 89*(1), 45–86.

Aguiar, M. A., M. Bils, and C. Boar (2020, January). Who Are the Hand-to-Mouth? Working Paper 26643, National Bureau of Economic Research. Series: Working Paper Series.

Allcott, H., J. Kim, D. Taubinsky, and J. Zinman (2022, May). Are High-Interest Loans Predatory? Theory and Evidence from Payday Lending. *The Review of Economic Studies 89*(3), 1041–1084.

Amador, M., I. Werning, and G.-M. Angeletos (2006). Commitment vs. Flexibility. *Econometrica 74*(2), 365–396. Publisher: [Wiley, Econometric Society].

Angeletos, G.-M., D. Laibson, A. Repetto, J. Tobacman, and S. Weinberg (2001, September). The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation. *Journal of Economic Perspectives 15*(3), 47–68.

Attanasio, O., A. Kovacs, and P. Moran (2020, October). Temptation and Commitment: Understanding Hand-to-Mouth Behavior. Working Paper 27944, National Bureau of Economic Research. Series: Working Paper Series.

Attanasio, O., A. Kovacs, and P. Moran (2021, June). Temptation and Incentives to Wealth Accumulation. Working Paper 28938, National Bureau of Economic Research. Series: Working Paper Series.

Barles, G. and P. E. Souganidis (1991, January). Convergence of approximation schemes for fully nonlinear second order equations. *Asymptotic Analysis 4*(3), 271–283. Publisher: IOS Press.

Bernheim, B. D. and D. Taubinsky (2018, January). Chapter 5 - Behavioral Public Economics. In B. D. Bernheim, S. DellaVigna, and D. Laibson (Eds.), *Handbook of Behavioral Economics: Applications and Foundations 1*, Volume 1 of *Handbook of Behavioral Economics - Foundations and Applications 1*, pp. 381–516. North-Holland.

Beshears, J., J. J. Choi, C. Clayton, C. Harris, D. Laibson, and B. C. Madrian (2020, July). Optimal Illiquidity. Working Paper 27459, National Bureau of Economic Research. Series: Working Paper Series.

Campanale, C., C. Fugazza, and F. Gomes (2015, April). Life-cycle portfolio choice with liquid and illiquid financial assets. *Journal of Monetary Economics 71*, 67–83.

Cao, D. and I. Werning (2016, February). Dynamic Savings Choices with Disagreements.

Cohen, J., K. M. Ericson, D. Laibson, and J. M. White (2020, June). Measuring Time Preferences. *Journal of Economic Literature 58*(2), 299–347.

Conesa, J. C., S. Kitao, and D. Krueger (2009, March). Taxing Capital? Not a Bad Idea after All! *American Economic Review 99*(1), 25–48.

da Costa, C. E. and M. R. Santos (2018). Age-Dependent Taxes with Endogenous Human Capital Formation. *International Economic Review 59*(2), 785–823. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/iere.12288.

Farhi, E. and I. Werning (2013). Insurance and Taxation over the Life Cycle. *The Review of Economic Studies 80*(2 (283)), 596–635. Publisher: [Oxford University Press, The Review of Economic Studies, Ltd.].

Findeisen, S. and D. Sachs (2017, February). Redistribution and insurance with simple tax instruments. *Journal of Public Economics 146*, 58–78.

Frederick, S., G. Loewenstein, and T. O'Donoghue (2002). Time Discounting and Time Preference: A Critical Review. *Journal of Economic Literature 40*(2), 351–401. Publisher: American Economic Association.

Ganong, P. and P. Noel (2019, July). Consumer Spending during Unemployment: Positive and Normative Implications. *American Economic Review 109*(7), 2383–2424.

Golosov, M., M. Troshkin, and A. Tsyvinski (2016, February). Redistribution and Social Insurance. *American Economic Review 106*(2), 359–386.

Gomes, F. (2020, November). Portfolio choice over the life-cycle: a survey. *Annual Review of Financial Economics 12*(1), 277–304. Number: 1 Publisher: Annual Reviews.

Gomes, F., M. Haliassos, and T. Ramadorai (2021, September). Household Finance. *Journal of Economic Literature 59*(3), 919–1000.

Gul, F. and W. Pesendorfer (2001). Temptation and Self-Control. *Econometrica 69*(6), 1403–1435. Publisher: [Wiley, Econometric Society].

Gul, F. and W. Pesendorfer (2004). Self-Control and the Theory of Consumption. *Econometrica 72*(1), 119–158. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1468-0262.2004.00480.x.

Harris, C. and D. Laibson (2013). INSTANTANEOUS GRATIFICATION. *The Quarterly Journal of Economics 128*(1), 205–248. Publisher: Oxford University Press.

Heathcote, J., K. Storesletten, and G. L. Violante (2020, September). Optimal progressivity with age-dependent taxation. *Journal of Public Economics 189*, 104074.

Huggett, M. (1996, December). Wealth distribution in life-cycle economies. *Journal of Monetary Economics 38*(3), 469–494.

Kaplan, G. and G. L. Violante (2014). A Model of the Consumption Response to Fiscal Stimulus Payments. *Econometrica 82*(4), 1199–1239. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA10528.

Kaplan, G. and G. L. Violante (2022). The Marginal Propensity to Consume in Heterogeneous Agent Models. *Annual Review of Economics 14*(1), 747–775. _eprint: https://doi.org/10.1146/annurev-economics-080217-053444.

Kaplan, G., G. L. Violante, and J. Weidner (2014). The Wealthy Hand-to-Mouth. *Brookings Papers on Economic Activity 2014*(1), 77–138. Publisher: Brookings Institution Press.

Karabarbounis, M. (2016, April). A Road Map for Efficiently Taxing Heterogeneous Agents. *American Economic Journal: Macroeconomics 8*(2), 182–214.

Kovacs, A., H. Low, and P. Moran (2021). Estimating Temptation and Commitment Over the Life Cycle. *International Economic Review 62*(1), 101–139. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/iere.12491.

Krusell, P., B. Kuruşçu, and A. A. Smith Jr. (2010). Temptation and Taxation. *Econometrica 78*(6), 2063–2084. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA8611.

Laibson, D. (1997). Golden Eggs and Hyperbolic Discounting. *The Quarterly Journal of Economics 112*(2), 443–477. Publisher: Oxford University Press.

Laibson, D., P. Maxted, and B. Moll (2021, July). Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy. Working Paper 29094, National Bureau of Economic Research. Series: Working Paper Series.

Laibson, D., A. Repetto, and J. Tobacman (2007, August). Estimating Discount Functions with Consumption Choices over the Lifecycle. Working Paper, National Bureau of Economic Research.

Laibson, D., A. Repetto, J. Tobacman, P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford (2003). *Knowledge, Information, and Expectations in Modern Economics: In Honor of Edmund S. Phelps.* Princeton: Princeton University Press. Pages: 228-266.

Laibson, D. I. (1996, June). Hyperbolic Discount Functions, Undersaving, and Savings Policy.

Maxted, P. (2021). Present Bias in Consumption-Saving Models: A Tractable Continuous-Time Approach. *Mimeo*.

Meier, S. and C. Sprenger (2010, January). Present-Biased Preferences and Credit Card Borrowing. *American Economic Journal: Applied Economics 2*(1), 193–210.

Mirrlees, J. A. (1971). An Exploration in the Theory of Optimum Income Taxation. *The Review of Economic Studies 38*(2), 175–208. Publisher: [Oxford University Press, Review of Economic Studies, Ltd.].

Moser, C. and P. Olea de Souza e Silva (2019, January). Optimal Paternalistic Savings Policies. SSRN Scholarly Paper 2959844, Social Science Research Network, Rochester, NY.

Ndiaye, A. (2017, November). Flexible Retirement and Optimal Taxation. SSRN Scholarly Paper 3301236, Social Science Research Network, Rochester, NY.

O'Donoghue, T. and M. Rabin (1999, March). Doing It Now or Later. *American Economic Review 89*(1), 103–124.

Pavoni, N. and H. Yazici (2017, June). Optimal Life-Cycle Capital Taxation Under Self-Control Problems. *The Economic Journal 127*(602), 1188–1216.

Schilbach, F. (2019, April). Alcohol and Self-Control: A Field Experiment in India. *American Economic Review 109*(4), 1290–1322.

Stantcheva, S. (2017, December). Optimal Taxation and Human Capital Policies over the Life Cycle. *Journal of Political Economy 125*(6), 1931–1990. Publisher: The University of Chicago Press.

Thaler, R. (1981, January). Some empirical evidence on dynamic inconsistency. *Economics Letters 8*(3), 201–207.

Toussaert, S. (2018). Eliciting Temptation and Self-Control Through Menu Choices: A Lab Experiment. *Econometrica 86*(3), 859–889. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.3982/ECTA14172.

Weil, P. E. (2018, December). Redistribution from the Cradle to the Grave: A Unified Approach to Heterogeneity in Age, Income and Wealth. Technical Report pwe433, Job Market Papers. Publication Title: 2018 Papers.

Weinzierl, M. (2011). The Surprising Power of Age-Dependent Taxes. *The Review of Economic Studies 78*(4), 1490–1518. Publisher: [Oxford University Press, Review of Economic Studies, Ltd.].

İmrohoroğlu, S. (1998). A Quantitative Analysis of Capital Income Taxation. *International Economic Review 39*(2), 307–328. Publisher: [Economics Department of the University of Pennsylvania, Wiley, Institute of Social and Economic Research, Osaka University].

Appendix

A Derivations

A.1 Deterministic setting: generalized Euler equation with incomplete markets

This derivation extends Laibson (1996)'s Proposition 2 to a setting with borrowing constraints. Let us first define the problem: the t self of a sophisticated present-biased household maximizes:

$$\begin{aligned} &\max_{c_t, a_{t+1}} u\left(c_t\right) + \beta \sum_{s=t+1}^{T} \delta^{s-t} u\left(c_s\left(a_{t+1}\right)\right) \\ &\text{s.t.} \quad c_t + a_{t+1} = R a_t + y_t + \tau_t \\ &\qquad\quad a_{t+1} \geq \underline{a} \end{aligned} \tag{21}$$

Where c_s $\forall s \in \{t+1, ..., T\}$ is some function of a_{t+1} because it is not directly chosen by the t self. Instead, the $t+1$ self will choose c_{t+1} by solving (21) for the period $t+1$, which will feature different preferences, as β will shift one period. Therefore, this is a sequential game between the different selves, and the only way the t self can influence the consumption of future selves is through a_{t+1}. We assume this game has a unique subgame-perfect equilibrium, and proceed to derive it. The first order condition, coupled with the borrowing constraint, yield an Euler equation in the following form:

$$\begin{aligned} u'\left(c_t\right) &\geq \beta\left[\delta u'\left(c_{t+1}\right)\frac{\partial c_{t+1}}{\partial a_{t+1}} + \delta^2 u'\left(c_{t+2}\right)\frac{\partial c_{t+2}}{\partial a_{t+1}} + ...\right] \\ u'\left(c_t\right) &\geq \beta\left[\delta u'\left(c_{t+1}\right)\frac{\partial c_{t+1}}{\partial a_{t+1}} + \delta^2 u'\left(c_{t+2}\right)\frac{\partial c_{t+2}}{\partial a_{t+2}}\left(R - \frac{\partial c_{t+1}}{\partial a_{t+1}}\right) + ...\right] \\ u'\left(c_t\right) &\geq \beta\left[\delta R u'\left(c_{t+1}\right)\frac{\partial c_{t+1}}{\partial y_{t+1}} + \left(\delta R\right)^2 u'\left(c_{t+2}\right)\frac{\partial c_{t+2}}{\partial y_{t+2}}\left(1 - \frac{\partial c_{t+1}}{\partial y_{t+1}}\right) + ...\right] \\ u'\left(c_t\right) &\geq \beta\sum_{s=t+1}^{T}\left(\delta R\right)^{s-t} u'\left(c_s\right)\frac{\partial c_s}{\partial y_s}\prod_{k=t+1}^{s-1}\left(1 - \frac{\partial c_k}{\partial y_k}\right) \end{aligned}$$

The steps above follow directly from the budget constraint for each period. Note that $\frac{\partial c_{t+1}}{\partial y_{t+1}}$ is the marginal propensity to consume (MPC) in period $t+1$, which is one minus the marginal propensity to save (MPS), which are multiplied in the product on the far right. Isolate the first term in the sum to get:

$$u'\left(c_t\right) \geq \beta\delta R u'\left(c_{t+1}\right)\frac{\partial c_{t+1}}{\partial y_{t+1}} + \beta\sum_{s=t+2}^{T}\left(\delta R\right)^{s-t} u'\left(c_s\right)\frac{\partial c_s}{\partial y_s}\prod_{k=t+1}^{s-1}\left(1 - \frac{\partial c_k}{\partial y_k}\right) \tag{22}$$

Due to incomplete markets, we have two possibilities for next period: either the household is a constrained HtM and the MPC is one, or it is not and the Euler equation holds with equality. In the former, the summation at the right becomes zero because the MPS of $t+1$ is zero, and the Euler equation reduces to

$$u'(c_t) \quad \geq \quad \beta\delta R u'(c_{t+1})$$

In the latter case, we can write the Euler for $t+1$ with equality:

$$u'(c_{t+1}) \quad = \quad \beta \sum_{s=t+2}^{T} (\delta R)^{s-(t+1)} u'(c_s) \frac{\partial c_s}{\partial y_s} \prod_{k=t+2}^{s-1} \left(1 - \frac{\partial c_k}{\partial y_k}\right)$$

Multiply by $\delta R \left(1 - \frac{\partial c_{t+1}}{\partial y_{t+1}}\right)$:

$$\delta R u'(c_{t+1}) \left(1 - \frac{\partial c_{t+1}}{\partial y_{t+1}}\right) \quad = \quad \beta \sum_{s=t+2}^{T} (\delta R)^{s-t} u'(c_s) \frac{\partial c_s}{\partial y_s} \prod_{k=t+1}^{s-1} \left(1 - \frac{\partial c_k}{\partial y_k}\right)$$

Substitute this term in (22) and rearrange to get:

$$u'(c_t) \quad \geq \quad \delta \left[1 - \frac{\partial c_{t+1}}{\partial y_{t+1}} (1-\beta)\right] R u'(c_{t+1})$$

Which is (6). Notice that this equation also fits the case where the household is constrained HtM at $t+1$, when simply $\frac{\partial c_{t+1}}{\partial y_{t+1}} = 1$.

A.2 Linear Tax Profile with Zero Present Value

Recall our two equations restricting the policy:

$$\begin{aligned} \int_0^{T_\tau} e^{-rt} \tau(t)\, dt &= 0 \\ \tau(t) &= \tau_0 + \tau_1 t \end{aligned}$$

Insert the latter into the former and solve the integral:

$$\begin{aligned} \int_0^{T_\tau} e^{-rt} (\tau_0 + \tau_1 t)\, dt &= 0 \\ \tau_0 \int_0^{T_\tau} e^{-rt} dt + \tau_1 \int_0^{T_\tau} e^{-rt} t dt &= 0 \end{aligned}$$

Integrating by parts on the second integral,

$$\begin{aligned} \frac{\tau_0 \left(1 - e^{-rT_\tau}\right)}{r} + \tau_1 \left[\frac{\left(-e^{-rT_\tau} T_\tau\right)}{r} - \int_0^{T_\tau} \frac{-e^{-rt}}{r} dt\right] &= 0 \\ \frac{\tau_0 \left(1 - e^{-rT_\tau}\right)}{r} + \tau_1 \left[\frac{\left(-e^{-rT_\tau} T_\tau\right)}{r} + \frac{\left(1 - e^{-rT_\tau}\right)}{r^2}\right] &= 0 \\ \tau_0 r \left(1 - e^{-rT_\tau}\right) + \tau_1 \left[-r e^{-rT_\tau} T_\tau + 1 - e^{-rT_\tau}\right] &= 0 \end{aligned}$$

Rearranging we get

$$\tau_1 = \frac{\tau_0 r \left(1 - e^{-rT_\tau}\right)}{\left(1 + rT_\tau\right) e^{-rT_\tau} - 1}$$

B Computational Algorithm

Starting from (17), we can discretize the problem by letting $V_{i,j}^k \equiv v(a,z,t)$ where i indexes a, j indexes z, k indexes t:

$$\left(\rho + \theta^k\right) V_{i,j}^k = u\left(c_{i,j}^k\right) + \partial_a V_{i,j}^k \dot{a}_{i,j}^k + \sum_{j' \neq j} \lambda_{j'j} \left[V_{i,j'}^k - V_{i,j}^k\right] + \dot{V}_{i,j}^k$$

We will use a finite difference approach, that is we will calculate backward and forward differences of V_l^k with respect to assets as:

$$\begin{aligned}
\partial_{a,B} V_{i,j}^k &= \frac{V_{i,j}^k - V_{i-1,j}^k}{\Delta_a} \\
\partial_{a,F} V_{i,j}^k &= \frac{V_{i+1,j}^k - V_{i,j}^k}{\Delta_a}
\end{aligned}$$

Where Δ_a is the length of a step in asset gridpoint. For the boundaries of the asset grid we must impose state constraints:

$$\begin{aligned}
\partial_{a,B} V_{1,j}^k &= u'\left(\underbrace{ra_1}_{0} + \exp\left(x_k + z_j\right) + \tau_k\right) \\
\partial_{a,F} V_{I,j}^k &= u'\left(ra_I + \exp\left(x_k + z_j\right) + \tau_k\right)
\end{aligned}$$

Where I is the number of asset gridpoints. Let $D = \{F, B\}$ .These finite differences will be useful because according to the first order condition and the upwinding method:

$$\begin{aligned}
u'(c) &= \partial_a V(a,z,t) \\
c_{i,j,D}^k &= u'^{-1}\left(\partial_{a,D} V_{i,j}^k\right)
\end{aligned}$$

And then we can calculate:

$$\dot{a}_{i,j,D}^k = ra_i + x_k z_j - c_{i,j,D}^k$$

Additionally, define:

$$\begin{aligned}
\left(\dot{a}_{i,j,D}^k\right)^+ &\equiv \max\left\{\dot{a}_{i,j,D}^k, 0\right\} \\
\left(\dot{a}_{i,j,D}^k\right)^- &\equiv \min\left\{\dot{a}_{i,j,D}^k, 0\right\}
\end{aligned}$$

Let Δ_t be length of the time step. Then we can work towards a matrix form of the HJB:

$$
\begin{aligned}
\left(\rho+\theta^k\right)V_{i,j}^k &= u\left(c_{i,j}^k\right)+\frac{V_{i+1,j}^k-V_{i,j}^k}{\Delta_a}\left(\dot{a}_{i,j,F}^k\right)^+ +\frac{V_{i,j}^k-V_{i-1,j}^k}{\Delta_a}\left(\dot{a}_{i,j,B}^k\right)^- +\sum_{j'\neq j}\lambda_{j'j}\left[V_{i,j'}^k-V_{i,j}^k\right]+\frac{V_{i,j}^{k+1}-V_{i,j}^k}{\Delta_t}\\
\left(\rho+\theta^k\right)V_{i,j}^k &= u\left(c_{i,j}^k\right)+V_{i+1,j}^k\underbrace{\frac{\left(\dot{a}_{i,j,F}^k\right)^+}{\Delta_a}}_{z_{i,j}^k}+V_{i,j}^k\underbrace{\left(\frac{\left(\dot{a}_{i,j,B}^k\right)^-}{\Delta_a}-\frac{\left(\dot{a}_{i,j,F}^k\right)^+}{\Delta_a}-\sum_{j'\neq j}\lambda_{j'j}\right)}_{y_{i,j}^k}\\
&+ V_{i-1,j}^k\underbrace{\frac{-\left(\dot{a}_{i,j,B}^k\right)^-}{\Delta_a}}_{x_{i,j}^k}+\sum_{j'\neq j}V_{i,j'}^k\lambda_{j'j}+\frac{V_{i,j}^{k+1}-V_{i,j}^k}{\Delta_t}\\
\left(\rho+\theta^k\right)V^k &= u^k+\mathbf{A}^kV^k+\mathbf{I}\frac{V^{k+1}-V^k}{\Delta_t}
\end{aligned}
$$

Here, $\mathbf{A}^k$ is the discretized infinitesimal generator for the asset and income process in matrix form:

$$
\mathbf{A}^kV^k=\begin{pmatrix}
y_{1,1} & z_{1,1} & \cdots & \lambda_{1,2} & \cdots & \lambda_{1,J} & \cdots\\
x_{2,1} & y_{2,1} & z_{2,1} & 0 & \ddots & \ddots & \ddots\\
0 & x_{3,1} & y_{3,1} & z_{3,1} & \ddots & \ddots & \ddots\\
\ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots\\
\ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots\\
\ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots\\
\cdots & \lambda_{J,1} & \cdots & \lambda_{J,J-1} & \cdots & x_{I,J} & y_{I,J}
\end{pmatrix}
\begin{pmatrix}
V_{1,1}^k\\ V_{2,1}^k\\ \vdots\\ V_{I,1}\\ V_{1,2}\\ \vdots\\ V_{I,J}
\end{pmatrix}
$$

Notice that $\mathbf{A}^k$ and u^k depend on V^k and vice-versa, so we need a fixed point. We can find it by updating with a semi-implicit method, letting l be an iteration:

$$
\begin{aligned}
\left(\rho+\theta^k\right)V_{l+1}^k &= u_l^k+\mathbf{A}_l^kV_{l+1}^k+\frac{V^{k+1}-V_{l+1}^k}{\Delta_t}\\
V_{l+1}^k\left[\mathbf{I}\left(\rho+\theta^k+\frac{1}{\Delta_t}\right)-\mathbf{A}_l^k\right] &= u_l^k+\frac{V^{k+1}}{\Delta_t}\\
V_{l+1}^k &= \left[\mathbf{I}\left(\rho+\theta^k+\frac{1}{\Delta_t}\right)-\mathbf{A}_l^k\right]^{-1}\left[u_l^k+\frac{V^{k+1}}{\Delta_t}\right]
\end{aligned}
$$

We update until convergence. This satisfies the Barles and Souganidis (1991) conditions, so convergence is guaranteed. In the simpler case of trying to solve (16), we do not need a fixed point because $\mathbf{A}^k$ and u^k are given from the exponential-discounting HJB, which we just solved. Therefore, the updating process follows a simple backwards recursion in time:

$$
W_{l+1}^k=\left[\mathbf{I}\left(\rho+\theta^k+\frac{1}{\Delta_t}\right)-\mathbf{A}^k\right]^{-1}\left[u^k+\frac{W^{k+1}}{\Delta_t}\right]
$$