

Present Bias Unconstrained: Consumption, Welfare, and the Present-Bias Dilemma*

Peter Maxted
Berkeley Haas

May 2023

Abstract

By augmenting the continuous-time specification of [Harris and Laibson \(2013\)](#) with the assumption that hard borrowing constraints do not bind in equilibrium, present bias can be tractably incorporated into rich consumption-saving models featuring stochastic income, risky and illiquid assets, and costly borrowing. I present closed-form expressions characterizing how present bias affects consumption, illiquid asset demand, and welfare. This welfare analysis specifies the channels through which present bias can matter for policy, and uncovers “the present-bias dilemma”: present bias can have large welfare costs, but individuals have little ability to alleviate these costs using financial commitment devices like illiquid assets.

Keywords: Present-Bias Dilemma, Illiquid Assets, Credit Card Borrowing, Instantaneous Gratification

*This paper was formerly titled “Present Bias in Consumption-Saving Models: A Tractable Continuous-Time Approach.” I would like to thank Hunt Allcott, Ned Augenblick, John Campbell, Taha Choukhmane (discussant), Christopher Clayton, Stefano DellaVigna, Xavier Gabaix, Thomas Graeber, Cosmin Ilut (discussant), David Laibson, Sean Lee, Chen Lian, Lu Liu, Omeed Maghziian, Ulrike Malmendier, Benjamin Moll, Andrew Perry, Matthew Rabin, Andrei Shleifer, Ludwig Straub, and participants at BEAM, the AEA Annual Meeting, the NBER Behavioral Finance Conference, and various institutions for insightful discussion and comments. I would like to thank the NBER Fellowship Program in Behavioral Macroeconomics for financial support.

1 Introduction

There is widespread evidence that consumers exhibit “present bias” across a variety of decision-making contexts. This evidence exists in lab settings (Frederick et al., 2002; Cohen et al., 2020), and in field settings ranging from credit card and payday loan usage (Meier and Sprenger, 2010; Allcott et al., 2022) to consumption choices during unemployment spells (Ganong and Noel, 2019) to retirement savings decisions (Madrian and Shea, 2001). Despite the evidence that consumers exhibit present bias, the modeling of consumption-saving behavior has been slow to incorporate these insights. This is because the literature is stuck at an impasse: while present bias can introduce a variety of novel and economically relevant behaviors, models with present bias are often difficult to characterize in practice.¹

This paper makes two contributions, one methodological and one analytical. The methodological contribution is the use of continuous-time methods to break this impasse and forge a new path forward. I develop a continuous-time toolbox for tractably characterizing the consumption-saving behavior of present-biased agents. This methodological innovation enables the analytical contribution of this paper. I present a novel set of results that provide closed-form answers to key open questions about the effects of present bias, including: How does present bias affect consumption? How does present bias affect the demand for illiquid assets? How does naivete affect these choices? How does present bias affect welfare? How does present bias affect the desirability of policy interventions?

In discrete time, present-biased preferences are characterized by the quasi-hyperbolic discount function: $1, \beta\delta, \beta\delta^2, \beta\delta^3, \dots$. Short-run discount factor β creates a disproportionate focus on the present period by driving a wedge between utility experienced “now” and utility experienced “later.” Whenever $\beta < 1$, preferences are time inconsistent. In the context of consumption-saving models, present bias implies that each

¹Present bias generates strategic interactions between selves, making consumption-saving decisions the equilibrium outcome of a dynamic intrapersonal game. Such strategic behavior often produces equilibrium non-uniqueness and consumption pathologies (i.e., highly sensitive policy functions that feature non-monotonicities and downward discontinuities). There is a large literature documenting that these issues can make models with present bias difficult to solve (e.g., Harris and Laibson, 2001, 2003; Krusell and Smith, 2003; Chatterjee and Eyigungor, 2016; Cao and Werning, 2018; Laibson and Maxted, 2022).

self overconsumes relative to the preferences of any other self.

The modeling of present bias also requires an assumption about the extent to which agents are aware of their self-control problems (O’Donoghue and Rabin, 1999, 2001). “Sophisticated” agents are fully aware of their time inconsistency. “Partially naive” agents underestimate the magnitude of their self-control problems, and instead expect (incorrectly) that all future selves will behave according to the discount function: $1, \beta^E \delta, \beta^E \delta^2, \beta^E \delta^3, \dots$, where $\beta^E \in (\beta, 1)$. In the limiting $\beta^E = 1$ case, “fully naive” agents completely fail to foresee their future present bias. Beliefs in this paper will generally be specified flexibly with β^E as a free parameter, thus allowing for an analysis of how naivete affects choices and welfare.

As in Harris and Laibson (2013), this paper studies present bias in the limiting continuous-time specification that results when the length of each period is taken to zero. This continuous-time specification of present bias is referred to as *Instantaneous Gratification* (IG), because each self lives for a vanishingly short period of time and discounts all future selves discretely by β . While the assumption that each self lives for a single instant is made for mathematical convenience, Laibson and Maxted (2022) show that IG preferences closely approximate discrete-time models with period lengths that are psychologically appropriate.²

I study IG preferences in a rich model of household balance sheets that allows for stochastic income, liquid and illiquid assets, and high-cost borrowing.³ Even in this general environment, the tractability of IG preferences allows me to derive closed-form expressions characterizing the effect of present bias on choices and welfare.

Before presenting the results, I emphasize that they rely on two main assumptions: (i) individuals have constant relative risk aversion (CRRA) utility; and (ii) borrowing limits do not bind in equilibrium. This second assumption is a key insight of this paper, and it is pivotal in unlocking many of the novel (and perhaps controversial) results that follow. Though the assumption of non-binding borrowing limits

²As detailed in Section 2, laboratory studies find that the temporal division between “now” and “later” is less than one week. However, discrete-time consumption-saving models typically use either quarterly or annual time-steps that are inconsistent with the high frequency at which present bias generally operates.

³These are common and important features in modern consumption-saving models; see e.g. Kaplan and Violante (2014), Berger et al. (2018), Kaplan et al. (2018), and Auclert et al. (2018).

may seem strong at first glance, the model allows for flexibly specified interest rates on borrowing, effectively replacing binding borrowing limits with arbitrarily onerous interest rate schedules in order to curb borrowing. Methodologically, when the above two assumptions hold I show that the IG agent's equilibrium behavior can be characterized directly from the behavior of a standard exponential agent (i.e., an agent with $\beta = 1$ and hence without present bias). This simple but powerful observation allows me to express the *relative* effect of present bias in closed form, even in highly complex and cutting-edge models that must be solved numerically.

I start by characterizing the consumption decisions of present-biased agents. Though the effect of present bias on consumption has been characterized in simplified models that can be solved by hand (e.g., cake-eating models), an open question is how present bias affects consumption decisions in more realistic environments. With IG preferences – even in models featuring stochastic income, costly borrowing, and multiple assets of varying return and liquidity properties – the effect of present bias can be characterized in closed form. Let β denote the agent's true short-run discount factor, let $\beta^E \in [\beta, 1]$ denote their perceived present bias, and let γ denote the coefficient of relative risk aversion. If a standard exponential agent consumes \check{c} , a present-biased agent will consume $\left(\frac{\beta^E}{\beta}\right)^{\frac{1}{\gamma}} \left(\frac{\gamma}{\gamma - (1 - \beta^E)}\right) \times \check{c}$.

I also present an Euler equation for the IG agent to highlight the implications of this consumption rule. As in [Harris and Laibson \(2001\)](#), the Euler equation shows that the IG agent acts relatively more impatiently when their MPC is large, and relatively more patiently when their MPC is small. This state-dependent discounting is an endogenous outcome of dynamic disagreement. Intuitively, high MPCs discourage saving because high MPCs imply that a marginal dollar of savings will be more rapidly consumed by future selves. When the consumption function is concave, as is typical in incomplete markets models, present bias will therefore cause consumers to act relatively more impatiently when liquidity is low, and relatively more patiently as they accumulate liquidity.

This closed-form consumption equation can also be used to characterize the effect of naivete on consumption decisions. If a sophisticated agent consumes c^S , a naive agent with $\beta^E \in (\beta, 1]$ will consume $\left(\frac{\beta^E}{\beta}\right)^{\frac{1}{\gamma}} \frac{\gamma - (1 - \beta)}{\gamma - (1 - \beta^E)} \times c^S$. Moreover, it also implies

an observational equivalence between sophisticates and naifs. A sophisticate with short-run discount factor β will consume identically to a (partial) naif with perceived short-run discount factor β^E and true short-run discount factor $\beta' = \beta^E \left[\frac{\gamma - (1 - \beta)}{\gamma - (1 - \beta^E)} \right]^\gamma$. An important takeaway from this observational equivalence is that it will be difficult to identify sophistication versus naivete using field data on realized consumption choices.

Next, I study how present bias affects the demand for illiquid assets. Asset illiquidity is a prominent focus of research on present bias. Since the seminal papers of [Strotz \(1956\)](#) and [Laibson \(1997\)](#), much of the literature has argued that present-biased consumers seek out illiquidity as a commitment device to increase savings and limit overconsumption.

In sharp contrast to this typical conclusion, I show that present bias does not necessarily engender a demand for illiquid assets. Provided that the borrowing constraint does not bind in equilibrium, present-biased consumers do not seek out illiquidity because illiquid assets do not actually limit overconsumption. Intuitively, generating commitment is like playing a game of Whack-a-Mole — the *illiquid* asset is never needed to fund current consumption, because the agent can always increase their consumption by adjusting their holdings of the *liquid* asset instead. Indeed, the existence of a liquid asset completely undoes any commitment benefits from the illiquid asset. Retirement systems around the world rely on illiquidity to incentivize retirement savings ([Beshears et al., 2015](#)). However, the results in this paper cast doubt on the benefits of such policies.

Turning to normative considerations, I derive a closed-form expression characterizing the welfare cost of present bias. Again, this closed-form result holds even in rich environments with stochastic income, high-cost borrowing, and multiple assets. To present this welfare metric I consider the following experiment. Suppose that there exists a perfect commitment device that forces all future selves to behave with complete self-control ($\beta = 1$), but this device costs a perpetual consumption tax of τ . The welfare cost of present bias is equivalent to a perpetual consumption tax of $\tau = 1 - \left(\frac{\alpha^\gamma}{1 - \gamma + \gamma\alpha} \right)^{\frac{1}{1-\gamma}}$, where $\alpha = \left(\frac{\gamma - (1 - \beta^E)}{\gamma} \right) \left(\frac{\beta}{\beta^E} \right)^{\frac{1}{\gamma}}$.

This welfare cost can be large. For example, if $\beta = \beta^E = 0.75$ and $\gamma = 2$, the

welfare cost of present bias is equivalent to a perpetual 2% consumption tax. Under full naivete ($\beta^E = 1$), this cost rises to 2.4%. If $\beta = 0.5$ and $\beta^E = 1$, as estimated in Laibson et al. (2023), the welfare cost of present bias is equivalent to a perpetual consumption tax of 17.2%. For context, these costs are at least an order of magnitude larger than back-of-the-envelope estimates of the welfare cost of transitory business cycles (Lucas, 1987), and sit at the upper end of calculations in the literature (e.g., Krusell et al., 2009).

Given the magnitude of this welfare cost, it is natural to ask how the welfare of present-biased agents can be improved. To answer this question, note first that the welfare cost of present bias depends on only three parameters: β , β^E , and γ . Looked at the other way around, this highlights all of the variables that the welfare cost of present bias does not depend on: wealth levels, the income process, interest rates, and illiquidity. Accordingly, any policy intervention that alters these variables will improve the welfare of a present-biased agent if and only if it also improves the welfare of a standard exponential agent. This is a key policy takeaway — it implies that a policymaker does not need to consider present bias when determining whether or not a given policy is welfare improving.

From the perspective of an individual consumer, this same conclusion leads to the quandary that I call the *present-bias dilemma*: the welfare cost of present bias is large, but it is difficult for an individual to reduce. Self-imposed financial commitment devices, such as penalty borrowing rates or asset illiquidity, will not improve the welfare of a present-biased agent because they do not improve the welfare of a standard exponential agent. Intuitively, though self-imposed financial penalties can improve incentives they may not generate perfect commitment, in which case the benefit of improved incentives can be dominated by the added financial cost of the penalized behavior that still occurs. That is, such commitment devices are undermined by “slips twixt the cup and the lip.”

Commitment has been perhaps *the* central theme of research on self-control problems ever since Strotz (1956). However, the present-bias dilemma suggests that such a focus may have been largely misguided. Despite the present-bias dilemma thus being a conclusion that flies in the face of conventional wisdom, it is well-known that

“commitment is a problematic prediction [in the first place], since we see so little of it in the economy” (Laibson, 2015, p. 267). The present-bias dilemma provides a novel remedy to this long-standing tension.

The present-bias dilemma implies, regrettably, that the welfare cost of present bias is both large and difficult for an individual to mitigate. On a brighter note, I end the paper by outlining one potential resolution to this dilemma. Government interventions differ from the sorts of financial commitments that any individual can self-impose, because governments can not only impose corrective taxes (which alone will not improve welfare), but can also redistribute revenues back to consumers. Unlike financial penalties alone, I show in a simple model that the combination of penalties plus redistribution can be welfare-improving.⁴

Related Literature. The methodological contribution of this paper is to show that IG preferences are an essential tool for modeling present bias in frontier consumption-saving models. The methods developed here are applied in Laibson et al. (2021) and Lee and Maxted (2023). While Laibson et al. (2021) and Lee and Maxted (2023) thus provide a more quantitative assessment of the channels through which present bias can matter for household behavior and macroeconomic policy, my focus here is instead on analytically characterizing the decisions of present-biased agents and the welfare consequences of those decisions.

The IG model was first developed in Harris and Laibson (2013). Laibson and Maxted (2022) show that discrete-time models with short period lengths (e.g., 1 week) are closely approximated by continuous-time IG models. The IG framework builds on the foundational work of Barro (1999) and Luttmer and Mariotti (2003), and more recent continuous-time implementations include Grenadier and Wang (2007), Cao and Werning (2016), Shigeta (2020), Acharya et al. (2022), Beshears et al. (2022a), and Rivera (2022).

The consumption-saving model that I study is cast in continuous time and features stochastic income, costly borrowing, and liquid and illiquid assets, along the lines of Kaplan et al. (2018). I show how – both analytically and numerically – to tractably

⁴See e.g. Moser and Olea de Souza e Silva (2019) and Beshears et al. (2022b) for richer analyses related to government interventions in economies populated by present-biased consumers.

incorporate present-biased preferences into these sorts of frontier incomplete markets models.

Using these novel methods, I provide theoretical results on overconsumption and high-cost borrowing that add to a large literature studying how present bias encourages short-term borrowing on unsecured accounts such as credit cards (Heidhues and Köszegi, 2010; Meier and Sprenger, 2010; Gathergood, 2012; Kuchler and Pagel, 2020) and payday loans (Skiba and Tobacman, 2018; Allcott et al., 2022). I also study the interaction of present bias with asset illiquidity, building on papers such as Strotz (1956), Laibson (1997), Angeletos et al. (2001), Amador et al. (2006), Galperti (2015), Bond and Sigurdsson (2018), Moser and Olea de Souza e Silva (2019), and Beshears et al. (2022b).

Finally, I use IG preferences to characterize the welfare cost of present bias. For a discussion of welfare in models with time-inconsistent preferences, see Bernheim and Rangel (2009) and Bernheim and Taubinsky (2018). This analysis also relates to the more general literature studying present-biased agents' demand for commitment. For overviews, see DellaVigna (2009), Bryan et al. (2010), Laibson (2015), and Carrera et al. (2020).

2 Instantaneous Gratification: A Summary

I begin by summarizing the Harris and Laibson (2013) specification of Instantaneous Gratification (IG) time preferences. In discrete time, the quasi-hyperbolic discount function is given by: $1, \beta\delta, \beta\delta^2, \beta\delta^3, \dots$. IG preferences are the continuous-time limit of this discount function, where each self lives for a vanishingly short length of time.

Let the current period be denoted t . Taking the limit of the discrete-time discount function, IG preferences are described by the limiting discount function $D(s)$ for $s \geq t$:

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

Parameter ρ is the exponential discount rate. Parameter $\beta \leq 1$ is the short-run

discount factor, which drives a wedge between utility “now” and utility “later.” When $\beta < 1$, discount function $D(s)$ features a discontinuity at $s = t$. This is because the current self lives for only a single instant, and discounts all future selves by β . For reference, Figure 1 below plots an IG discount function for $\beta = 0.75$ and $\rho = 2\%$.

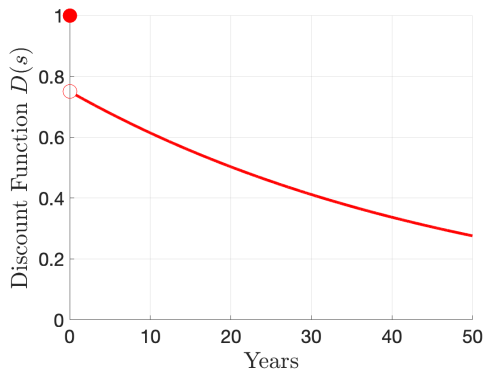


Figure 1: **IG Discount Function.** This figure plots an IG discount function for $\beta = 0.75$ and $\rho = 2\%$. This discount function features a discontinuity between “now” and “later.”

IG preferences should be thought of as a mathematically tractable limit case, since the temporal division between “now” and “later” is certainly longer than a single instant dt . However, and in contrast to the quarterly or annual horizons that discrete-time consumption-saving models have often used in the past, the “psychological present” is estimated to be short: [Augenblick \(2018\)](#) estimates that the division between “now” and “later” is approximately 2 hours, [Augenblick and Rabin \(2019\)](#) find that essentially all discounting occurs within one week, and [McClure et al. \(2007\)](#) use fMRI data to estimate that food rewards are discounted by 50% over a one-hour horizon.⁵ In light of this evidence, [Laibson and Maxted \(2022\)](#) substantiate the IG specification by showing that it closely approximates models in which the duration of “now” is psychologically appropriate. Accordingly, one of my goals in the current paper is to illustrate just how tractable the IG specification can be for incorporating present bias into frontier consumption-saving models.⁶

⁵See also [DellaVigna \(2018\)](#) and [Gottlieb and Zhang \(2021\)](#) for related discussions.

⁶The results of [Laibson and Maxted \(2022\)](#) imply that one could analogously study present bias in discrete time, but with short (e.g. daily) time-steps. In practice, there are two drawbacks to this approach relative to IG. First, the IG model allows for closed-form expressions characterizing

Remark. *IG time preferences are a generalization of standard time-consistent preferences. Exponential discounting is recovered by setting $\beta = 1$.*

Expectations and the Intrapersonal Equilibrium. The modeling of present bias requires an assumption about the extent to which each self foresees the present bias of future selves (O’Donoghue and Rabin, 1999, 2001). As discussed in the introduction, I use β^E to denote the short-run discount factor that the current self expects all future selves to have.

For all but full naivete, present-biased agents don’t share the perceived preferences of future selves. This complicates the analysis of consumption-saving models, because it means that decisions need to be modeled as a dynamic intrapersonal game played by different “selves” of the agent (Strotz, 1956; Laibson, 1997). Taking prices as given, an equilibrium to this intrapersonal game will be referred to as an *intrapersonal equilibrium*.

I follow Harris and Laibson (2013) in studying stationary Markov-perfect equilibria to the intrapersonal game (Maskin and Tirole, 2001).⁷ For the model analyzed in this paper, a critical property of IG preferences is that the intrapersonal equilibrium is unique under this refinement. Moreover, the intrapersonal equilibrium satisfies a partial differential equation, and Section 7.2 discusses how to use well-developed numerical methods to characterize this equilibrium. This is in contrast to discrete time, where equilibria may not be unique and the equilibrium identified by numerical methods often contains pathological properties.⁸

Economists have differing views on the extent to which agents are aware of their self-control problems (see e.g. DellaVigna (2018) and Allcott et al. (2022) for discussions). I do not take a stand on this issue: for all levels of naivete, this paper provides a general method for solving models with present-biased agents that is analytically

the effect of present bias. Second, discrete-time models with short time-steps can be slow to solve numerically, whereas this paper extends the fast continuous-time methods discussed in Achdou et al. (2022) to IG preferences (see Section 7.2).

⁷Though the weaker subgame-perfect refinement can introduce a variety of interesting equilibria (Laibson, 1994; Bernheim et al., 2015), an analysis of non-Markov equilibria is beyond the scope of this paper.

⁸See Harris and Laibson (2001, 2003), Krusell and Smith (2003), Cao and Werning (2016), and Laibson and Maxted (2022) for discussions of equilibrium pathologies and non-uniqueness.

and numerically tractable, robust to pathologies, and features a unique equilibrium.

3 Consumption-Saving Model

I now present the consumption-saving model that I study in this paper. The model is cast in continuous time following [Achdou et al. \(2022\)](#), and features a generalized income process, costly borrowing, and illiquid assets in order to capture key innovations of frontier models in the literature (e.g., [Kaplan et al., 2018](#)). Despite these enrichments, the model presented below is intentionally streamlined along various dimensions for simplicity and to maintain a closer connection to existing papers in the literature. The tractability of IG preferences implies that many of the results in this paper will continue to hold in even richer economic environments. [Section 7.1](#) discusses some relevant extensions.

Throughout this paper I study consumption-saving decisions in partial equilibrium, and I focus on the behavior of a single agent. The reason that I focus on a single agent in partial equilibrium is that this is where the issues with present bias arise. Once I show how to characterize the intrapersonal equilibrium of a present-biased agent, expansion to heterogeneous agents and aggregation to general equilibrium follow standard practices.⁹

3.1 The Household Balance Sheet

My model of the household balance sheet is similar to [Kaplan et al. \(2018\)](#). I adopt their notation when possible.

The household faces idiosyncratic income risk. The household’s income flow at time t is denoted y_t , and income y_t can follow an arbitrary finite-state Poisson process subject to the restriction that the minimum income state is weakly positive.¹⁰ Let $\lambda^{y \rightarrow y'}$ denote the switching intensity from income state y to income state y' .

⁹For more, the Supplementary Material Appendix presents a general equilibrium Aiyagari-Bewley-Huggett model augmented with present-biased agents. Similar to [Achdou et al. \(2022\)](#), the solution takes the form of two coupled PDEs.

¹⁰Many continuous processes can be approximated arbitrarily closely by such finite-state processes.

The household has access to both a liquid asset b and an illiquid asset a . For the liquid asset, when $b > 0$ the household earns a constant return of r on their liquid wealth. The household can also borrow up to a borrowing limit of \underline{b} in the liquid asset. However, borrowing is (potentially) costly. Specifically, I assume that a *marginal* dollar of borrowing requires the household to pay an interest rate wedge of $\omega(b) \geq 0$ above the risk-free rate r .¹¹ This implies that when $b < 0$ the household pays an *average* interest rate on their debt of $r + \mathcal{W}(b)$, where $\mathcal{W}(b) = \frac{\int_b^0 \omega(q) dq}{-b}$ denotes the average borrowing wedge.

In this model it will be notationally easier to work with average interest rates. Let $r(b)$ denote the average wealth-varying interest rate, which is given by:

$$r(b) = \begin{cases} r & \text{if } b \geq 0 \\ r + \mathcal{W}(b) & \text{if } b < 0 \end{cases}. \quad (2)$$

When $\mathcal{W}(0) > 0$, equation (2) introduces what is known as a “soft borrowing constraint.” The soft constraint discourages borrowing by setting a higher interest rate on borrowing than on saving. Additionally, the flexible specification of borrowing wedge $\mathcal{W}(b)$ allows for the average interest rate on borrowing to increase as the household borrows more. Once the household hits the hard borrowing limit of \underline{b} , additional borrowing is completely restricted.

The illiquid asset a has an expected return of r^a and a volatility of σ^a . Short positions against the illiquid asset are restricted (i.e., $a_t \geq 0$). Let d_t denote the household’s flow of deposits to (or withdrawals from) the illiquid asset at time t . The asset is illiquid because these deposits/withdrawals are subject to a flow transaction cost of $\chi(d, a)$. I assume that transaction cost function $\chi(d, a)$ is everywhere differentiable with $\chi(0, a) = 0$, $\chi(d, a) \geq 0$, and $\frac{\partial^2 \chi(d, a)}{\partial d^2} > 0$.¹² Two relevant benchmarks are: (i) perfect liquidity, with no transaction costs on asset a ; and (ii) a 401(k)/IRA plan, with no cost to contributions and a constant 10% penalty on withdrawals. While these

¹¹ $\omega(b)$ is defined on $[\underline{b}, 0]$ and can have a finite number of discontinuities that capture, for example, the household maxing out its credit card and switching to borrowing using payday loans.

¹²Convexity prevents a_t from jumping (this is for simplicity, and is later extended in Section 7.1). $\chi(d, a)$ could also feature a kink at $d = 0$ to produce an inaction region (Kaplan et al., 2018).

two benchmarks do not satisfy the restrictions placed on χ , they can be approximated arbitrarily closely by allowable functions.

Finally, I assume that the household is not allowed to deposit/withdraw from the illiquid asset when they are at the borrowing constraint (i.e., $d_t = 0$ when $b_t = \underline{b}$). This simplifies the exposition by providing the simple boundary condition that $c_t \leq y_t + r(\underline{b})\underline{b}$ when the household is at the borrowing limit \underline{b} . This assumption is a minimal restriction in practice, because a household with liquid wealth of $b_t = \underline{b}$ can choose to consume ϵ less for a vanishingly short period of time in order to move away from \underline{b} and regain access to the illiquid asset.¹³

To summarize, the household's balance sheet evolves as follows:

$$db_t = (y_t + r(b_t)b_t - d_t - \chi(d_t, a_t) - c_t) dt \quad (3)$$

$$\frac{da_t}{a_t} = \left(r^a + \frac{d_t}{a_t} \right) dt + \sigma^a dZ_t, \quad (4)$$

subject to the constraints $b_t \geq \underline{b}$ and $a_t \geq 0$. Z_t is a standard Brownian motion.

3.2 Utility and Value

The household accrues CRRA utility over consumption:

$$u(c) = \begin{cases} \frac{c^{1-\gamma}-1}{1-\gamma} & \text{if } \gamma \neq 1 \\ \ln(c) & \text{if } \gamma = 1 \end{cases}. \quad (5)$$

Under IG time preferences, the *actual* continuation-value function is given by:

$$v_t = \mathbb{E}_t \left[\int_t^\infty e^{-\rho(s-t)} u(c_s) ds \right], \quad (6)$$

and the *actual* current-value function is given by:

$$w_t = \beta v_t. \quad (7)$$

¹³Moreover, this paper will focus on the case where \underline{b} never binds in the first place (see Assumption 1 below). Even without the simplifying assumption that $d_t = 0$ when $b_t = \underline{b}$, note that there is already a maximum flow amount that the household will choose to withdraw because χ is convex.

The intuition for equation (7) is as follows. The current self discounts the utility of all future selves by β . But in continuous time the current self lives for just a single instant dt , and therefore the utility accrued by the current self has no measurable impact on the overall value function. So, $w_t = \beta v_t$.¹⁴

I emphasize the term *actual* for equations (6) and (7) because the expectation operator in those equations denotes the expectation of the modeler. This will not necessarily equal the household's own expectation if they are partially or fully naive (i.e., $\beta^E \in (\beta, 1]$). The household's *perceived* continuation-value function is:

$$v_t^E = \mathbb{E}_t \left[\int_t^\infty e^{-\rho(s-t)} u(c_s^E) ds \right], \quad (8)$$

where c^E denotes the consumption rate that the household would adopt if they were sophisticated with short-run discount factor β^E . The *perceived* current-value function is:

$$w_t^E = \beta v_t^E. \quad (9)$$

The true short-run discount factor β is still used in equation (9), since the current self discounts the utility of future selves by β regardless of naivete.

Throughout this paper I impose the restriction that $\gamma > 1 - \beta^E$, similar to [Harris and Laibson \(2013\)](#). This ensures that the desire to smooth consumption (γ) is greater than perceived time inconsistency ($1 - \beta^E$).

3.3 Intrapersonal Equilibrium: Definition and Uniqueness

I consider stationary Markov-perfect equilibria in three state variables: liquid wealth $b \in [\underline{b}, \infty)$, illiquid wealth $a \in [0, \infty)$, and income $y \in \{y_1, y_2, \dots, y_N\}$. To simplify notation, let $x = (b, a, y)$ denote the vector of state variables that characterize the household's balance sheet position.

¹⁴To show this as a limiting argument, let Δ denote the time-step of a discrete-time model. Heuristically, $w_t = \beta v_t$ results from $w_t = \lim_{\Delta \rightarrow 0} u(c_t)\Delta + \beta e^{-\rho\Delta} \mathbb{E}_t v_{t+\Delta}$.

Equilibrium Under Sophistication. I begin by defining the intrapersonal equilibrium for a sophisticated agent. I then generalize the definition to allow for naivete.

Starting with the sophisticate's problem, a stationary Markov-perfect equilibrium to the sophisticated IG agent's intrapersonal problem is characterized by the following Bellman equation, which consists of a differential equation defined on x .¹⁵

$$\begin{aligned} \rho v(x) = & u(c(x)) + v_b(x) (y + r(b)b - d(x) - \chi(d(x), a) - c(x)) \\ & + v_a(x) (r^a a + d(x)) + \frac{1}{2} v_{aa}(x) (a\sigma^a)^2 \\ & + \sum_{y' \neq y} \lambda^{y \rightarrow y'} (v(b, a, y') - v(b, a, y)), \end{aligned} \quad (10)$$

subject to the optimality conditions:

$$u'(c(x)) = \begin{cases} \beta v_b(x) & \text{if } b > \underline{b} \\ \max \left\{ \beta v_b(x), u'(y + r(\underline{b})\underline{b}) \right\} & \text{if } b = \underline{b} \end{cases}, \quad \text{and} \quad (11)$$

$$\chi_d(d(x), a) = \begin{cases} \frac{v_a(x)}{v_b(x)} - 1 & \text{if } b > \underline{b} \\ 0 & \text{if } b = \underline{b} \end{cases}. \quad (12)$$

Equation (10) defines the continuation-value function v of the IG agent. It says that the instantaneous change in value due to discounting (ρv) must equal the current utility flow ($u(c)$) plus the expected instantaneous change in the value function ($\mathbb{E}dv/dt$).¹⁶

Equation (11) defines the IG agent's consumption choice. In continuous time, consumption is unconstrained for all $b > \underline{b}$. Whenever consumption is unconstrained, the IG agent sets the marginal utility of consumption equal to the marginal value of current liquid wealth: $u'(c(x)) = w_b(x) = \beta v_b(x)$. At the borrowing constraint \underline{b} , the optimality condition is refined to ensure that the agent does not violate the constraint: $c(x) \leq y + r(\underline{b})\underline{b}$. If $\beta v_b(x) \geq u'(y + r(\underline{b})\underline{b})$ then the agent will choose to

¹⁵Under certain calibrations v will have a convex kink, so viscosity solutions are used. See [Harris and Laibson \(2013\)](#) and [Achdou et al. \(2022\)](#) for details.

¹⁶See [Harris and Laibson \(2013\)](#), [Laibson et al. \(2021\)](#), or [Laibson and Maxted \(2022\)](#) for similar Bellman equations.

set $c(x) \leq y + r(\underline{b})\underline{b}$. Otherwise, consumption is restricted to $y + r(\underline{b})\underline{b}$.

Equation (12) defines the asset allocation decision $d(x)$. For $b > \underline{b}$ the IG agent chooses $d(x)$ to equate the marginal value of illiquid wealth to the marginal value of liquid wealth, adjusted for transaction costs: $w_a(x) = w_b(x)(1 + \chi_d(d(x), a))$. Using the property that $w(x) = \beta v(x)$ and rearranging gives the top row of equation (12). The bottom row of equation (12) imposes the restriction that the agent must set $d = 0$ when $b = \underline{b}$.

Comparing the consumption decision in equation (11) to the asset allocation decision in equation (12), the key difference between the two decisions is that the β discount factor only has a direct effect on the consumption decision. Intuitively, β does not directly impact the asset allocation decision because this decision only affects the consumption of future selves.¹⁷ However, $\beta < 1$ can still indirectly affect the asset allocation decision if it alters v_a or v_b .

Equations (10) through (12) look similar to the Hamilton-Jacobi-Bellman (HJB) equation that would arise for a standard exponential agent with $\beta = 1$.¹⁸ The key difference is that present bias alters the consumption optimality condition (11). The IG agent sets $u'(c(x)) = \beta v_b(x)$, whereas a standard exponential agent would instead set $u'(c(x)) = v_b(x)$. In both cases the current self sets the marginal utility of consumption equal to the marginal value of liquid wealth. However, under IG preferences the marginal value of wealth is discounted by β , since wealth is consumed by future selves whose utility is discounted by β .

Value Function Uniqueness. As in Harris and Laibson (2013), an important property of IG preferences is that value function v is unique. This is in contrast to discrete-time models of present bias, where non-uniqueness is known to exist in some

¹⁷For $b > \underline{b}$, the current (instantaneous) self always has enough liquidity to fund their consumption.

¹⁸The $\beta = 1$ agent's HJB equation is (ignoring boundary conditions; see also Kaplan et al., 2018):

$$\begin{aligned} \rho v(x) = \max_{c,d} & u(c) + v_b(x) (y + r(b)b - d - \chi(d, a) - c) \\ & + v_a(x) (r^a a + d) + \frac{1}{2} v_{aa}(x) (a\sigma^a)^2 \\ & + \sum_{y' \neq y} \lambda^{y \rightarrow y'} (v(b, a, y') - v(b, a, y)). \end{aligned}$$

deterministic models (Krusell and Smith, 2003; Cao and Werning, 2016; Laibson and Maxted, 2022), and uniqueness in stochastic models has only been proven for β close to 1 (Harris and Laibson, 2003).

Proposition 1. *The IG agent's value function $v(x)$ is unique.*

Proof. Unless stated in the main text, all proofs are provided in Appendix A. The proof here relies on methods that are not introduced until Section 4. \square

One-Step Extension to Naivete. I now extend the equilibrium definition to allow for naivete. Recall that a naive agent believes that all future selves will be sophisticated with short-run discount factor $\beta^E \in (\beta, 1]$. Thus, a naive agent believes that equations (10) through (12) characterize the equilibrium that all future selves will follow (except that β is replaced by β^E in equation (11)).

Let $v^E(x)$ denote the value function that solves equations (10) through (12) for a sophisticated agent with short-run discount factor β^E . Using equation (9), the naive agent's actual consumption and asset allocation decisions are given by:

$$u'(c(x)) = \begin{cases} \beta v_b^E(x) & \text{if } b > \underline{b} \\ \max\{\beta v_b^E(x), u'(y + r(\underline{b})\underline{b})\} & \text{if } b = \underline{b} \end{cases}, \text{ and} \quad (13)$$

$$\chi_d(d(x), a) = \begin{cases} \frac{v_a^E(x)}{v_b^E(x)} - 1 & \text{if } b > \underline{b} \\ 0 & \text{if } b = \underline{b} \end{cases}. \quad (14)$$

Comparing the naif's actual behavior to their beliefs, naivete creates incorrect expectations about the consumption decision but not the asset allocation decision. For consumption, the naif expects that future selves' consumption choices will depend on β^E , but equation (13) shows that actual consumption decisions depend on β .

4 Tractability in Continuous Time: The \hat{u} Agent

A key benefit of IG preferences is that the problem of the dynamically inconsistent IG agent can be recast as a dynamically consistent optimization problem. Specifically,

Harris and Laibson (2013) show that the value function of the sophisticated IG agent is equivalent to the value function of an agent who discounts exponentially ($\beta = 1$), but has a reverse-engineered utility function denoted \hat{u} . Thus, one does not need to solve directly for the intrapersonal equilibrium of the IG agent. Instead, one can solve the time-consistent \hat{u} agent’s problem, and then recover the IG agent’s intrapersonal equilibrium from the \hat{u} agent.

While the \hat{u} agent is thus an important mathematical tool, there is no inherent economic content to their behavior. In particular, the \hat{u} agent has a nonstandard utility function, \hat{u} , that is reverse-engineered for the sole purpose of characterizing the IG agent’s equilibrium. This creates a problem: while the IG agent’s behavior can be characterized from the \hat{u} agent’s behavior, the \hat{u} agent’s behavior is itself nonstandard (and economically immaterial).¹⁹

So, the critical next step in the current paper is to relate the behavior of this nonstandard \hat{u} agent to the behavior of a “standard” exponential agent ($\beta = 1$) that also has standard CRRA utility. Then, one can use the \hat{u} agent as a conduit to link the policy and value functions of the IG agent to those of the standard exponential agent.

4.1 Key Insight: Non-Binding Borrowing Limits

To establish this mapping between the standard exponential agent and the \hat{u} agent, most of the results below exploit the following observation:

Remark. *If the borrowing limit does not bind in equilibrium then the \hat{u} utility function is a positive affine transformation of standard CRRA utility. Accordingly, if the borrowing limit does not bind in equilibrium then the \hat{u} agent’s policy functions are identical to those of the standard exponential agent.*

This simple observation is the key to unlocking the power of IG preferences. It implies that whenever the borrowing limit does not bind (further discussed in Section

¹⁹Indeed, Harris and Laibson (2013, p. 207) write: “The nonstandard optimization problem [of the \hat{u} agent] is interesting, not because we think it is psychologically relevant, but because its partial equivalence enables us to use the machinery of optimization to study the value function of a dynamically inconsistent problem.”

5.1), the nonstandard \hat{u} agent actually behaves identically to the standard exponential agent. When this is the case, the IG agent’s behavior can then be characterized directly from that of the standard exponential agent, even in complex models that must be solved numerically.

4.2 Introducing Two Additional Types of Agents

Filling in the above discussion, I now introduce the two additional types of agents that will be used for characterizing the IG agent: the \hat{u} agent and the standard exponential agent.

Definition (\hat{u} Agent). *The first agent is referred to as the “ \hat{u} agent.” The \hat{u} agent discounts exponentially ($\beta = 1$) but has a modified utility function, denoted \hat{u} (defined below). The value and policy functions of the \hat{u} agent will be denoted with a hat (e.g., $\hat{v}(x)$ and $\hat{c}(x)$).*

Following Harris and Laibson (2013), the \hat{u} agent is reverse-engineered so that the value function of the \hat{u} agent is equivalent to the continuation-value function of the sophisticated IG agent. The \hat{u} agent allows for the IG agent’s problem to be recast as a time-consistent optimization problem.

Definition (Standard Exponential Agent). *The second agent is referred to as the “standard exponential agent.” The standard exponential agent discounts exponentially ($\beta = 1$) and has standard CRRA utility $u(c)$. The value and policy functions of the standard exponential agent will be denoted with an upside-down hat (e.g., $\check{v}(x)$ and $\check{c}(x)$).*

I call the standard exponential agent “standard” because they have the CRRA utility and exponential time preferences that economists typically work with. Moreover, the standard exponential agent is what the IG agent would become if they had $\beta = 1$. Using the \hat{u} agent as a conduit, the contribution of the analysis in Sections 5 and 6 is to relate the policy and value functions of the IG agent to those of the standard exponential agent.

The remainder of this section formalizes the above discussion. It can be skipped without loss of continuity, in which case the reader is directed to Section 5 for results.

4.3 The \hat{u} Construction

As in [Harris and Laibson \(2013\)](#), define

$$\hat{u}_+(\hat{c}) = \frac{\psi}{\beta} u\left(\frac{1}{\psi}\hat{c}\right) + \frac{\psi - 1}{\beta}, \quad \text{where} \quad \psi = \frac{\gamma - (1 - \beta)}{\gamma}. \quad (15)$$

$\hat{u}_+(\hat{c})$ is a positive affine transformation of CRRA utility function $u(c)$, and is constructed so that $\hat{u}_+(\hat{c}) < u(\hat{c})$ for all $\hat{c} > 0$. Additionally, note that $\psi \in (0, 1]$.²⁰

The complete \hat{u} utility function depends on whether or not the borrowing constraint binds. Fully, \hat{u} is defined as follows:²¹

$$\hat{u}(\hat{c}, x) = \begin{cases} \hat{u}_+(\hat{c}) & \text{if } b > \underline{b} \\ \hat{u}_+(\hat{c}) & \text{if } b = \underline{b} \text{ and } \hat{c} \leq \psi(y + r(\underline{b})\underline{b}) \\ -\infty & \text{if } b = \underline{b} \text{ and } \hat{c} \in (\psi(y + r(\underline{b})\underline{b}), y + r(\underline{b})\underline{b}) \\ u(\hat{c}) & \text{if } b = \underline{b} \text{ and } \hat{c} \geq y + r(\underline{b})\underline{b} \end{cases}. \quad (16)$$

The \hat{u} utility function can be split into two sub-cases: a case where the borrowing constraint does not currently bind, and a case where it does. When the constraint does not bind (first two lines), utility is given by $\hat{u}_+(\hat{c})$. The constraint binds when $b = \underline{b}$ and $\hat{c} = y + r(\underline{b})\underline{b}$ (fourth line). In this case the \hat{u} utility function is given by the standard CRRA utility function $u(\cdot)$. Since $u(\hat{c}) > \hat{u}_+(\hat{c})$, the \hat{u} agent can obtain a “utility boost” at $b = \underline{b}$ by setting $\hat{c} = y + r(\underline{b})\underline{b}$.

The third line imposes that the \hat{u} agent earns $-\infty$ utility whenever $b = \underline{b}$ and $\hat{c} \in (\psi(y + r(\underline{b})\underline{b}), y + r(\underline{b})\underline{b})$. Essentially, this forces the \hat{u} agent to make a choice at \underline{b} : they can either set $\hat{c} \leq \psi(y + r(\underline{b})\underline{b})$ or $\hat{c} = y + r(\underline{b})\underline{b}$. The former choice earns lower utility \hat{u}_+ but allows the agent to move off of the constraint. The latter choice earns the “utility boost,” but requires the agent to stay constrained at \underline{b} .

I again emphasize that the \hat{u} utility function is just a reverse-engineered mathematical tool to ensure that the value function of the \hat{u} agent, denoted $\hat{v}(x)$, is equivalent to the continuation-value function $v(x)$ of the sophisticated IG agent. To understand

²⁰This follows from the parameter restrictions of $\beta \in (0, 1]$ and $\gamma > 1 - \beta$ for a sophisticated agent.

²¹For pedagogy, this paper’s specification of the \hat{u} utility function differs slightly from [Harris and Laibson \(2013\)](#) in order to emphasize the effect of borrowing constraint \underline{b} on the \hat{u} agent’s behavior.

how the \hat{u} construction produces this value function equivalence, consider first the case where the \hat{u} agent is unconstrained (either $b > \underline{b}$ or $\hat{c}(\underline{b}, a, y) \leq \psi(y + r(\underline{b})\underline{b})$). Since the \hat{u} agent is time consistent they choose \hat{c} to maximize \hat{v} , whereas the time-inconsistent IG agent overconsumes. Thus, the \hat{u} agent's utility must be adjusted downward to ensure that $\hat{v}(x) = v(x)$, which is why $\hat{u}_+(c) < u(c)$. Next, consider the case where the \hat{u} agent is constrained at \underline{b} . Since this constraint also prevents the IG agent from overconsuming, the \hat{u} agent's utility no longer needs to be penalized, so $\hat{u}(\hat{c}, x) = u(\hat{c})$. That is, this “utility boost” captures that binding constraints benefit present-biased agents by restricting overconsumption at \underline{b} .

4.4 Recovering the IG Equilibrium from the \hat{u} Agent

The next proposition formalizes the tractability that the reverse-engineered \hat{u} agent provides.

Proposition 2. *Let $\hat{v}(x)$ denote the value function of the \hat{u} agent:*

$$\hat{v}(x_t) = \max_{\{\hat{c}_s\}} \mathbb{E}_t \int_t^\infty e^{-\rho(s-t)} \hat{u}(\hat{c}_s, x_s) ds. \quad (17)$$

$v(x)$ is the continuation-value function of the sophisticated IG agent (i.e., it solves (10) – (12)) if and only if $v(x)$ is the value function of the \hat{u} agent.

Proposition 2 shows that the sophisticated IG agent's intrapersonal equilibrium can be recast as an optimization problem for the \hat{u} agent. This property is used to prove the uniqueness of the IG agent's value function (Proposition 1). Having characterized the value function $v(x)$ of the sophisticated IG agent, equations (11) and (12) then determine the sophisticate's policy functions $c(x)$ and $d(x)$. Similarly, equations (13) and (14) can be used to determine the naif's policy functions.

Understanding the \hat{u} Agent. Proposition 2 shows that the sophisticated IG agent's value function can be characterized using the time-consistent \hat{u} agent. This also implies that the IG agent's policy functions can be related to those of the \hat{u} agent. But, the \hat{u} agent is a reverse-engineered mathematical apparatus with a nonstandard

utility function. Accordingly, the key to understanding the \hat{u} agent's behavior is to recognize when the \hat{u} agent does, and does not, behave identically to the standard exponential agent.

Remark. *If the borrowing constraint never binds for any income state then the \hat{u} agent behaves identically to the standard exponential agent. In other words, $\hat{c}(x) = \check{c}(x)$ and $\hat{d}(x) = \check{d}(x)$. This is due to the fact that $\hat{u}_+(\hat{c})$ is a positive affine transformation of $u(c)$, and if the constraint never binds then $\hat{u}(\hat{c}, x) = \hat{u}_+(\hat{c})$ for all x .*

This affine-transformation property no longer holds when the borrowing constraint binds for some income state, because now the \hat{u} agent receives a nonstandard “utility boost” that represents the benefits of such constraints for present-biased agents.

5 The Effect of Present Bias on Policy Functions

5.1 Key Assumption: The Borrowing Limit Does Not Bind

For the remainder of this paper I assume that the hard borrowing constraint at \underline{b} does not bind in equilibrium.²² As emphasized in Section 4.1, this assumption allows for the IG agent to be characterized directly from the standard exponential agent.

Assumption 1. *The model is calibrated such that borrowing limit \underline{b} does not bind in equilibrium. Formally, if $b_0 > \underline{b}$ then $b_t > \underline{b}$ for all $t \geq 0$.*

While Assumption 1 eliminates binding hard borrowing constraints, recall that the flexible specification of borrowing rates in equation (2) allows for borrowing wedge $\mathcal{W}(b)$ to become arbitrarily large as the agent borrows more. Little realism is likely to be lost by eliminating binding borrowing limits and replacing them with arbitrarily large (but finite) borrowing rates. Indeed, the latter assumption is likely more realistic. If an agent is bound by a hard borrowing constraint, that means that the current self wants to consume more but it is *impossible* for them to do so (i.e., the cost of marginal consumption is infinite). I instead assume that the agent always has access to *some* technology that allows them to consume more if they choose to do so (e.g.,

²²A sufficient – but not necessary – condition is that \underline{b} is the natural borrowing limit.

a credit card, payday lender, pawn shop, or loan shark), but I place no restriction on the cost of that marginal borrowing except that it is finite.

The reason that the model’s predictions are sensitive to Assumption 1 is that present-biased agents’ equilibrium behavior interacts with binding borrowing limits. In particular, present-biased agents recognize that binding constraints provide a powerful “bright-line” restriction on future selves’ overconsumption (i.e., $c_t \leq y_t + r(\underline{b})\underline{b}$ at the borrowing limit).²³ Assumption 1, alternatively, provides no such upper bound — along the equilibrium path, the agent always *can* consume more if they choose to do so. The results in Sections 5.3 and 6 illustrate how powerful this optionality can be for unwinding commitment.

It is also worth briefly noting that even in a case where the researcher does not want Assumption 1 to hold, the IG specification is still methodologically useful. Specifically, it will still be the case that: (i) the equilibrium of the IG agent is unique; and (ii) the \hat{u} agent can be used to solve for the IG agent’s intrapersonal equilibrium.²⁴ The feature that is lost, however, is that the \hat{u} agent no longer behaves identically to the standard exponential agent, and hence the IG agent can no longer be directly related to the standard exponential agent.

5.2 Present Bias and Overconsumption

I begin by providing closed-form expressions for the effect of present bias on consumption.

Proposition 3. *Assumption 1 holds. Let β denote the agent’s true short-run discount factor, let $\beta^E \in [\beta, 1]$ denote the agent’s perceived short-run discount factor, and let $\psi^E = \frac{\gamma - (1 - \beta^E)}{\gamma}$. Relative to the standard exponential agent, the consumption of the*

²³As discussed in Section 4.3, the benefits of binding constraints are highlighted by the “utility boost” obtained by the \hat{u} agent at \underline{b} . For further analysis, the Supplementary Material Appendix presents a workhorse Aiyagari-Bewley-Huggett model with a binding hard borrowing constraint. That Appendix also provides additional theoretical results on the interaction of present bias with the binding constraint.

²⁴I.e., the proofs of Propositions 1 and 2 do not rely on Assumption 1.

IG agent is given by:

$$c(x) = \left(\frac{\beta^E}{\beta}\right)^{\frac{1}{\gamma}} \frac{1}{\psi^E} \times \check{c}(x). \quad (18)$$

Equation (18) simplifies in two cases of particular interest. Under sophistication, the IG agent consumes $\frac{1}{\psi}$ times the standard exponential agent, where $\psi = \frac{\gamma-(1-\beta)}{\gamma}$ (defined in (15)). Under full naivete, the IG agent consumes $\beta^{-\frac{1}{\gamma}}$ times the standard exponential agent.

Proposition 3 can also be used to compare sophisticates versus naifs:

Corollary 4. *Assumption 1 holds. If a sophisticate with $\beta^E = \beta$ consumes $c^S(x)$, a naif with $\beta^E \in (\beta, 1]$ will consume $\left(\frac{\beta^E}{\beta}\right)^{\frac{1}{\gamma}} \frac{\gamma-(1-\beta)}{\gamma-(1-\beta^E)} \times c^S(x)$. Consumption is increasing in naivete when $\gamma > 1$, and decreasing in naivete when $\gamma < 1$.*

Proof. The standard exponential agent's consumption function $\check{c}(x)$ is independent of β and β^E . This corollary then follows from equation (18). \square

In stylized environments featuring linear consumption functions, it is known that sophisticates and naifs adopt the same equilibrium consumption function when $\gamma = 1$ (e.g., Tobacman, 2007). Corollary 4 shows that this result continues to hold in much more general environments. Intuitively, 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 ($\gamma < 1$), and vice versa. These two effects exactly offset when $\gamma = 1$.

Taking this comparison of sophisticates and naifs even further, Proposition 3 also implies that there is an observational equivalence between sophisticates and naifs.²⁵

Corollary 5. *Assumption 1 holds. A sophisticated agent with short-run discount factor β will consume identically to a naive agent with perceived short-run discount factor $\beta^E \in [\beta, 1]$ and true short-run discount factor $\beta' = \beta^E \left[\frac{\gamma-(1-\beta)}{\gamma-(1-\beta^E)} \right]^\gamma$.*

²⁵See also Blow et al. (2021) for a related result.

Proof. The standard exponential agent’s consumption function $\check{c}(x)$ is independent of β and β^E . This corollary then follows from equation (18). \square

A key takeaway from Corollary 5 is that sophistication versus naivete cannot be easily identified from field data on consumption-saving decisions.²⁶ Instead, it is more likely that identification can be found by evaluating data on procrastination (O’Donoghue and Rabin, 1999), contract choices (DellaVigna and Malmendier, 2004, 2006; Gabaix and Laibson, 2006; Heidhues and Köszegi, 2010), or the accuracy of budgeting plans (Augenblick and Rabin, 2019; Kuchler and Pagel, 2020; Allcott et al., 2022), as these sorts of decisions follow more directly from agents’ misperception of their future actions.²⁷

State-Dependent Discounting: An Euler Equation. To gain more intuition for how present bias affects the consumption decision, continuous-time methods can also be used to obtain an Euler equation for the IG agent. For notational simplicity I assume full sophistication in the Euler equation below. But, recall from Corollary 5 that there is an observational equivalence between sophisticates and naifs whenever Assumption 1 holds. Additionally, Appendix A provides a generalized Euler equation that allows for naivete.

Proposition 6. *Assume the IG agent is sophisticated ($\beta^E = \beta$). Let $\varsigma(b_t)$ denote the marginal interest rate that the agent earns on their liquid wealth of b_t .²⁸ Whenever $c(x_t)$ and $r(b_t)$ are locally differentiable in b , consumption satisfies the following Euler equation:*

$$\frac{\mathbb{E}_t (du'(c(x_t))) / dt}{u'(c(x_t))} = \left[\rho + (1 - \beta)c_b(x_t) \right] - \varsigma(b_t). \quad (19)$$

²⁶Note that agents choose both consumption and illiquid deposits, but Corollary 5 only shows that there is an observational equivalence between the consumption of sophisticates and naifs. However, the observational equivalence continues to hold for illiquid deposits $d(x)$, as will be shown in Proposition 8 below.

²⁷Relatedly, see Lian (2021) for an analysis of how misperception can affect MPCs.

²⁸Fully, the marginal interest rate is

$$\varsigma(b) = \begin{cases} r & \text{if } b \geq 0 \\ r + \omega(b) & \text{if } b < 0 \end{cases}.$$

A generalized Euler equation that allows for naivete is given in the proof of this proposition (see Appendix A). Additionally, note that the Euler equation does not require Assumption 1.

Proposition 6 is the continuous-time analogue of the discrete-time Hyperbolic Euler Relation derived in Harris and Laibson (2001). A similar continuous-time Euler equation for sophisticates is presented in Harris and Laibson (2004).²⁹ The left side of equation (19) is the expected growth rate of marginal utility. When $\beta = 1$, we recover the standard Euler equation that the expected growth rate of marginal utility equals discount rate ρ minus interest rate $\varsigma(b)$ (Achdou et al., 2022). When $\beta < 1$, dynamic inconsistency means that the IG agent also cares about the extent to which future selves will consume out of a marginal dollar of savings, as captured by the instantaneous MPC of $c_b(x_t)$. This dynamic disagreement implies that the IG agent acts as if they have a state-dependent discount rate of $\rho + (1 - \beta)c_b(x_t)$.

As described in Harris and Laibson (2001), the intuition for why a present-biased agent's effective discount rate is increasing in $c_b(x_t)$ is as follows. A sophisticated agent knows that future selves will overconsume, which incentivizes the current self to set wealth aside in order to buffer against future overconsumption. However, the current self's ability to save for the future depends on the extent to which subsequent selves will overconsume out of any marginal savings. If $c_b(x)$ is large then marginal savings will be quickly consumed, reducing the current self's willingness to save.

Importantly, in models with stochastic income the consumption function is typically concave in liquid wealth, implying that MPCs will be higher for low-liquidity agents (Carroll and Kimball, 1996). When this is the case, equation (19) differentiates present bias from exponential discounting because (19) implies that present-biased consumers will act relatively more impatiently when their liquidity is low, and relatively more patiently as they build a buffer stock of liquid wealth.

Equation (19) can also be contextualized in relation to the literature on heterogeneous time preferences. When the consumption function is concave in liquid wealth, present bias will introduce similar effects as heterogeneous time preferences. An

²⁹However, their setup differs from the one here, and they do not generalize the Euler equation for naifs.

important difference, however, is that models with heterogeneous time preferences assume preference heterogeneity *across* individuals, while present bias endogenously generates effective time-preference heterogeneity *within* individuals that varies with liquid wealth. Relatedly, models with heterogeneous time preferences produce differences in wealth due to differences in patience, whereas models with present bias produce differences in patience due to differences in wealth.³⁰

Present Bias and High-Cost Borrowing. When $\beta = 1$, soft constraints as in equation (2) can generate a buildup of agents with exactly zero liquid wealth (see e.g. Achdou et al., 2022). When $\beta < 1$, soft constraints no longer prevent borrowing. To show this result in a simple environment, I assume here that there is only a single liquid asset b , and income is deterministic with $y_t = y > 0$ for all t . For notational simplicity I again present the following result for a sophisticated agent, but recall that there is an observational equivalence between sophisticates and naifs whenever Assumption 1 holds (Corollary 5).

Proposition 7. *Assumption 1 holds. Assume also that there is just a single liquid asset b , income is deterministic with $y > 0$, and the IG agent is sophisticated. Let $s(b) = y + r(b)b - c(b)$ denote the saving policy function. Set $r < \frac{\rho}{\beta}$ so that the IG agent dissaves for $b > 0$. Regardless of $\mathcal{W}(b)$, the IG agent chooses to accumulate debt at $b = 0$ by setting $s(0) < 0$.*

Proposition 7 provides the stark result that the IG agent will always choose to accumulate some amount of debt when $\beta < 1$, regardless of how large the initial interest rate on borrowing is. Note, however, that this is only a statement about the extensive margin of high-cost borrowing; just because the IG agent is willing to revolve some high-cost debt, this is not to say that they will necessarily borrow a lot.³¹

³⁰This result has some conceptual similarities to the temptation model of Banerjee and Mulainathan (2010). See also Aguiar et al. (2020) for empirical evidence that a positive correlation between liquidity and patience helps consumption-saving models fit the available data.

³¹Supplementary Material Appendix F presents a numerical example of the simplified model in Proposition 7. Appendix F also highlights an interesting implication of Proposition 2 in this deterministic model: when $r + \omega(b) > \rho$ for all $b < 0$, the IG agent's continuation-value function for $b \geq 0$ can be independent of $\omega(b)$ despite the IG agent eventually accumulating debt at that borrowing wedge. See also Lee and Maxted (2023) for further analysis of present bias and high-cost borrowing.

Note too that while Proposition 7 formally uses IG preferences, the intuition only depends on psychologically appropriate (i.e., short) – but not necessarily instantaneous – time-steps. To demonstrate this point, consider a discrete-time model in which a present-biased agent has zero liquid wealth ($b = 0$) and is deciding whether or not to borrow. Assume this agent has a β of 0.8, and the annualized borrowing rate is 25%. If periods are *one year* apart then the current self may decide to not borrow at a 25% rate. If periods are only *one day* apart, however, now that annualized borrowing rate of 25% corresponds to a daily rate of only 0.06%. Since the current self still discounts the next self (i.e., tomorrow) by $\beta = 0.8$, the current self may now be much more willing to borrow in order to increase consumption.

5.3 Present Bias May Not Engender a Demand for Illiquidity

I now turn to the present-biased agent’s demand for illiquid assets. Starting from the seminal papers of Strotz (1956) and Laibson (1997), much of the literature on present bias argues that present-biased agents seek out illiquid assets as a commitment against overconsumption.

In sharp contrast to this research, I present an irrelevance result showing that present bias does not necessarily affect the demand for illiquid assets:

Proposition 8. *Assumption 1 holds. The IG agent and the standard exponential agent choose the same asset allocation policy function: $d(x) = \check{d}(x)$. That is, asset allocation policy function $d(x)$ is independent of β and β^E .*

This irrelevance result arises in this class of two-asset models because the liquid asset eliminates any commitment benefits from the illiquid asset. The agent never needs the illiquid asset in order to finance current consumption — they can always adjust their holdings of the liquid asset instead. Indeed, Proposition 3 shows that the IG agent always consumes $\left(\frac{\beta^E}{\beta}\right)^{\frac{1}{\gamma}} \frac{1}{\psi^E}$ times the standard exponential agent’s consumption, meaning that asset illiquidity does not affect the relative overconsumption caused by present bias. Since the illiquid asset does not limit overconsumption, β has no effect on policy function $d(x)$.

Discussion. In the general class of two-(or more)-asset models considered here, Proposition 8 provides a sharp result highlighting the ineffectiveness of policies utilizing asset illiquidity to promote the saving of present-biased consumers.³² Despite the worldwide usage of illiquid retirement savings accounts (Beshears et al., 2015), Proposition 8 is consistent with the emerging empirical evidence that retirement savings plans are not particularly effective savings devices. For example, Argento et al. (2015) find that 30-40% of deposits into retirement accounts subsequently leak out before retirement.³³

More broadly, Proposition 8 speaks to the long-standing puzzle of why present-biased agents do not use commitment devices (e.g., Laibson, 2015; Bernheim and Taubinsky, 2018). This literature has concluded that commitment is often hampered by an important tradeoff between commitment and flexibility (e.g., Amador et al., 2006; Laibson, 2015). Proposition 8 suggests a complementary explanation: designing commitment devices is like playing a game of Whack-a-Mole. In my model, illiquid assets have no commitment benefits because the agent always has another margin (the liquid asset) that they can adjust in order to overconsume. Extending this intuition more generally, there are many margins that can potentially be adjusted to bring utility into the current period, ranging from the consumption of unhealthy food to decreasing exercise to staying up too late. Unless a commitment device can block all of these sources of temptation, there is no reason for the agent to choose an ineffective “commitment device” that actually serves only to limit flexibility.

One caveat to Proposition 8 is necessary. Though Proposition 8 allows for a flexible soft borrowing constraint it does still depend on Assumption 1 that there are no binding hard constraints. This contrasts with many of the earlier insights on illiquidity that have thus far been developed in the literature, which often rely on binding constraints. For example, in Laibson (1997) the agent is *always* endogenously liquidity constrained in equilibrium. At a minimum, this contrast clarifies that commitment

³²Proposition 8 says nothing about whether or not present-biased agents will save more when assets are illiquid. It only says that β does not have an independent effect on the asset allocation decision.

³³See Coyne et al. (2022) for additional evidence of leakage, and also Choukhmane (2019) who finds that workers who are auto-enrolled into their employer’s retirement plan adopt lower saving rates later in life.

does not come from the illiquid asset per se, but rather from the binding constraint on the liquid asset.³⁴ Interpreted more broadly, while binding constraints are a powerful modeling tool for simplifying models that can otherwise be extremely unruly, it’s less clear that our real-world economic decisions are truly characterized by such bright-line restrictions on consumption. So, the role of Assumption 1 is to instead allow for *the possibility* that the benefits of commitment can be undermined by “slips twixt the cup and the lip.” Proposition 8 highlights just how much such slips can matter.

This discussion also suggests one plausible way to break Proposition 8: it may fail to hold under less restrictive equilibrium refinements and/or when present bias is combined with boundedly rational optimization. For example, non-Markov or boundedly rational equilibria may feature “pseudo hard constraints” such as mental accounts (Thaler, 1985) and personal rules (Ainslie, 1992; Bernheim et al., 2015), which can reintroduce the types of notches in the consumer’s budget constraint that interact with present bias to produce non-irrelevance.

6 Welfare and the Present-Bias Dilemma

Welfare analyses can be difficult in models with time-inconsistent preferences because such preferences do not typically feature a single welfare criterion. In models with present bias, the common approach is to adopt a long-run view in which policymakers seek to maximize the continuation-value function v_t . However, this approach ignores the preferences of each individual self. This may be unsatisfactory if, for example, the current self is better able to evaluate their immediate preferences than a distant self (Bernheim and Rangel, 2009).

Unlike discrete-time models, continuous-time IG preferences feature a single welfare criterion (Harris and Laibson, 2013). This important property arises because each self lives for just an instant, and therefore composes only an infinitesimal part of the overall value function. More formally, the current self wants to adopt policies that

³⁴Agents against a hard constraint are prevented from freely financing consumption with the liquid asset (see equation (11)). If borrowing constraints bind, present-biased agents may seek to hold wealth in illiquid assets in order to put future selves at the constraint, thereby restricting future selves’ consumption.

maximize current-value function w_t . But since $w_t = \beta v_t$, any policy that maximizes v_t will also maximize w_t .³⁵ This single welfare criterion property, combined with the tractability of IG preferences, means that IG preferences are well-suited for studying policy and welfare in rich economic environments.

6.1 The Welfare Cost of Present Bias

IG preferences allow for a closed-form characterization of the welfare cost of present bias. In order to present a welfare metric that applies in the general consumption-saving environment that I am studying, I consider the following experiment. Assume that there exists a perfect commitment device that forces all future selves to behave with full self control ($\beta = 1$), but this commitment device comes at the cost of a perpetual consumption tax of τ . The welfare cost of present bias can be expressed in terms of consumption tax τ .

Proposition 9. *Assumption 1 holds. Let $\alpha = \psi^E \left(\frac{\beta}{\beta^E} \right)^{\frac{1}{\gamma}}$. The welfare cost of present bias is equivalent to a perpetual consumption tax of:*

$$\tau = 1 - \left(\frac{\alpha^\gamma}{1 - \gamma + \gamma\alpha} \right)^{\frac{1}{1-\gamma}}. \quad (20)$$

τ is decreasing in β , and τ is increasing in β^E if and only if $\gamma > 1$.³⁶

Proposition 9 is powerful because it is very general, and holds in this rich consumption-saving environment that allows for stochastic income, costly borrowing, and multiple assets of varying return and liquidity properties. As long as Assumption 1 applies, the welfare cost of present bias can be represented as a consumption tax of size τ .

The fact that such a simple formula characterizes the welfare cost of present bias across a general class of models may seem surprising. IG preferences make this welfare characterization possible. The proof relies on the fact that the value function of the IG agent can always be recast as the value function of an exponential agent with

³⁵While this single welfare criterion property only holds exactly under IG preferences, it is robust to discrete time with short time-steps, where each self is just a small part of the total value function.

³⁶When $\gamma = 1$, the tax is given by $\tau = 1 - \frac{\exp\left(\frac{\beta-1}{\beta}\right)}{\beta}$.

a modified utility function. Then, the proof boils down to finding the tax τ that provides an equivalent welfare loss as the modified utility function.³⁷

Discussion. The first takeaway from Proposition 9 is that present bias can be very costly. For example, if $\beta = \beta^E = 0.75$ and $\gamma = 2$ then the welfare cost of present bias is equivalent to a perpetual 2% consumption tax. For $\beta = 0.5$ and $\beta^E = 1$, as estimated in Laibson et al. (2023), the welfare cost rises to 17.2%. These costs are at least an order of magnitude larger than benchmark estimates of the welfare cost of business cycles (Lucas, 1987).

One caveat to these large welfare costs is that τ gives the realized cost of present bias. If the agent is naive, they may not perceive such a large welfare cost. A naif perceives that the welfare cost of their present bias is equivalent to a tax of $\tau^E = 1 - \left(\frac{(\psi^E)^\gamma}{\beta^E} \right)^{\frac{1}{1-\gamma}}$.³⁸

The second takeaway from Proposition 9 is that the welfare cost of present bias depends only on β , β^E , and γ . The welfare cost of present bias is decreasing in β , which is straightforward. The welfare cost is larger under naivete when $\gamma > 1$, and larger under sophistication when $\gamma < 1$. This follows from Corollary 4: overconsumption is exacerbated by naivete when $\gamma > 1$, and is reduced by naivete when $\gamma < 1$.

An alternate way to look at this second takeaway is to consider what the welfare cost of present bias does not depend on. The welfare cost of present bias is independent of liquid wealth b , illiquid wealth a , income state y , and market price parameters r , \mathcal{W} , χ , r^a , and σ^a . Though changes to these variables will certainly affect the agent's welfare, they do not independently affect the relative welfare cost of present bias. As Proposition 3 shows, the IG agent always consumes $\left(\frac{\beta^E}{\beta} \right)^{\frac{1}{\gamma}} \frac{1}{\psi^E}$ times the standard exponential agent's consumption, meaning that the relative overconsumption caused

³⁷It is worth noting that τ is only defined when $1 - \gamma + \gamma\alpha > 0$. This is not necessarily the case under naivete when $\gamma > 1$ and α is low. In these cases, the naif behaves so poorly that their realized value function goes to $-\infty$. Intuitively, the naif always thinks that they are only overconsuming for a single instant. When this mistake is made repeatedly, it can lead to a realized value function that is undefined (though at each point in time, the naif perceives that their value function is finite). The condition that $1 - \gamma + \gamma\alpha > 0$ can be thought of as a bound on the level of naivete that is theoretically admissible.

³⁸This is the maximum tax that a sophisticate with short-run discount factor β^E would accept in order to eliminate their present bias.

by present bias is constant over the state space and depends only on β, β^E , and γ . This intuition yields the following Proposition, which is a key policy implication of this welfare analysis:

Proposition 10. *Assumption 1 holds. A policy intervention that alters the income process, interest rates, and/or transaction costs improves the welfare of the IG agent if and only if it improves the welfare of the standard exponential agent.*

Proof. The proof of Proposition 10 follows directly from the proof of Proposition 9, which shows that the IG agent’s value function is a positive affine transformation of the standard exponential agent’s value function. So, any policy that increases one value function will increase the other, and vice-versa. \square

From the perspective of an individual IG agent, Proposition 10 implies that commitment devices, such as penalty borrowing rates and asset illiquidity, will not alleviate the welfare costs of present bias. To the extent that such devices are undesirable for the time-consistent exponential agent, they will also make the IG agent worse off. This forms the basis of the present-bias dilemma, as detailed in Section 6.2 below.

From the perspective of a policymaker, Proposition 10 is another irrelevance result: the policymaker does not need to consider present bias when determining whether or not a given policy is welfare improving. Instead, where present bias matters is in determining whether or not a given policy is feasible (i.e., whether it obeys a budget constraint). For example, consider a large interest rate subsidy on savings that is financed by a small consumption tax. Though such a policy may be welfare improving regardless of β and β^E , the revenue collected by the consumption tax may only be sufficient to cover the cost of the interest rate subsidy in an economy populated by present-biased consumers. This is because present-biased agents will not take full advantage of the policy, underusing the interest rate subsidy and overpaying the consumption tax. Appendix C provides a toy model that formalizes this discussion.

6.2 The Present-Bias Dilemma

Together, Propositions 9 and 10 show that present-biased agents face a quandary. On the one hand, present bias can be enormously costly (Proposition 9). On the

other hand, agents cannot use self-imposed financial commitment devices such as asset illiquidity or penalty borrowing rates to reduce this welfare cost (Proposition 10). This is the *present-bias dilemma*.

Building on Proposition 8, Proposition 10 suggests that the fact that we see so little commitment in the economy (Laibson, 2015) may not be such a puzzle after all. As Proposition 10 shows, commitment in the form of financial penalties does not improve IG agents’ welfare. Such commitment devices would only benefit the IG agent if they benefit the standard exponential agent, and the standard exponential agent would never choose to self-impose financial costs. For intuition, consider a penalty borrowing rate. Though this penalty rate will reduce borrowing, the borrowing that still occurs becomes more costly.³⁹

Potential Resolutions. I discuss three classes of potential resolutions to the present-bias dilemma. First, the inability for agents to create effective commitment devices relates to Proposition 3, which shows that the relative overconsumption caused by present bias is constant over the state space and depends only on β , β^E , and γ . Anything that breaks this constant-relative-overconsumption property may allow agents to develop effective commitment devices. For example, to the extent that it is realistic to depart from Assumption 1, doing so will explicitly restrict consumption at \underline{b} and hence eliminate the “slips” that otherwise undermine commitment (the “pseudo hard constraints” discussed in Section 5.3 could work similarly).

Second, if $\gamma > 1$ and agents are at least partially naive, then the welfare cost of present bias can be reduced by educating agents about their present bias (i.e., by making naifs more sophisticated). In the above example of $\beta = 0.5$ and $\gamma = 2$, turning a naif ($\beta^E = 1$) into a sophisticate ($\beta^E = \beta$) reduces the welfare cost of present bias from 17.2% to 11.1%.

Third, the present-bias dilemma is about the inability for any individual agent to self-impose financial commitments that alleviate their present bias. Government

³⁹Supplementary Material Appendix F presents a numerical example of this tradeoff. More broadly, this general intuition is supported by the evidence in John (2020), who documents that even in situations where researchers are able to induce subjects to take up a commitment device, it frequently fails. Additionally, nothing in this paper says that a demand for financial commitment devices cannot exist, only that present bias alone may not be the reason for that demand.

interventions are fundamentally different from what a consumer can self-impose. Policymakers not only have the ability to implement financial disincentives in the form of taxes, but also to redistribute tax revenues back to consumers. Thus, the present-bias dilemma suggests a potentially important role for policy in alleviating present bias.⁴⁰ To formalize this point, Appendix C presents a toy model in which the combination of taxes plus redistribution in the form of a savings subsidy can improve the welfare of present-biased agents.

7 Model Extensions and Numerical Methods

7.1 Model Extensions

Many of the results in this paper will continue to hold in even richer environments than the baseline model in Section 3. I discuss some relevant extensions here. These extensions are not mutually exclusive, and can typically be stacked together.

Liquid and Illiquid Assets. The baseline model has only a single liquid asset, but can be extended to allow for portfolio choice within the liquid asset. This can capture, for example, stock trading on wealth held in a brokerage account. The baseline model also has only a single illiquid asset, but results will continue to apply if the agent has access to multiple assets of varying return and liquidity properties. The model can also allow for time-varying and stochastic expected returns.

Second, the baseline model assumes that the illiquid asset has convex adjustment costs, and assets are adjusted with flow deposits/withdrawals. An alternative assumption is that agents pay fixed costs to adjust their illiquid wealth. With non-convex adjustment costs, adjustment decisions become an optimal-stopping (option-value) problem, and the model must be written using HJB Variational Inequalities instead

⁴⁰See e.g. Moser and Olea de Souza e Silva (2019) and Beshears et al. (2022b) for insightful analyses. Note that while penalty-plus-redistribution policies could also be implemented by private institutions, as it is the pooling of consumers that is essential, Laibson (1997) discusses why such schemes may be difficult for the private sector to implement. Additionally, political economy considerations are beyond the scope of this paper, but are important for understanding how governments can respond to the biases of their constituents.

of standard HJB equations.⁴¹ However, results continue to apply in this case. See Appendix D.1 for details. Similar optimal-stopping models can also be used to incorporate consumer bankruptcy.

Third, the illiquid asset is a strictly financial asset from which the agent derives no utility. This is a reasonable assumption for some illiquid assets (e.g., retirement accounts), but not for others (e.g., houses). Appendix D.2 presents the case in which the illiquid asset is a durable, like a house, that enters the utility function via a Cobb-Douglas aggregator. This case is still tractable with IG preferences, but the closed-form solutions are slightly different because illiquid housing affects the agent’s risk aversion (Flavin and Nakagawa, 2008).

β Heterogeneity. In any population there is very likely to be heterogeneity in β and β^E . Previously, modeling such heterogeneity was computationally difficult because it meant that an already sensitive model (see footnote 1) had to then be solved repeatedly for different levels of β and β^E . Now, the results in this paper imply that preference heterogeneity in β and β^E comes “for free.” Specifically, whenever Assumption 1 holds then the policy and value functions of a present-biased agent can be expressed in closed form based on those of the standard exponential agent. Computationally, this means that the model needs to be solved only once for the standard exponential agent. Then, introducing preference heterogeneity in β and β^E takes just a few extra lines of code and almost no additional runtime.

Additional Behavioral Biases. The tractability of IG preferences means that present bias can be modeled jointly with other types of behavioral biases, such as non-rational expectations (e.g., Bordalo et al., 2018; Maxted, 2020) or bounded rationality (e.g., Gabaix, 2016; Ilut and Valchev, 2020). Here, the point of comparison for the IG agent is no longer the standard exponential agent, but rather an agent without present bias ($\beta = 1$) but with the other behavioral biases.

⁴¹For other papers that use HJB Variational Inequalities, see Guerrieri et al. (2020), McKay and Wieland (2021), and Laibson et al. (2021).

7.2 Numerical Methods

Continuous-Time Algorithm. Models of the type presented in Section 3 typically require numerical solutions, and algorithms based on finite difference methods are a common way to solve these sorts of HJB equations.⁴² Barles and Souganidis (1991) prove that a finite difference scheme converges to the unique viscosity solution of an HJB equation whenever three conditions are met: (i) monotonicity; (ii) stability; and (iii) consistency. When applied directly to the Bellman equation of an IG agent, however, the monotonicity property may not be satisfied. To fix this problem, I develop a numerical algorithm for reestablishing a convergent finite difference scheme that follows directly from Proposition 2. First, solve the HJB equation of the time-consistent \hat{u} agent. Second, use the \hat{u} agent to back out the IG agent’s intrapersonal equilibrium. Appendix B provides details.

Discrete-Time Approximation. For researchers using discrete-time models, the closed-form expressions that I provide in continuous time can be viewed as approximate results in discrete time. Indeed, Laibson and Maxted (2022) show that the continuous-time IG specification provides a close approximation to discrete-time models that are written with short period lengths. So, researchers wanting a simple back-of-the-envelope way to study the robustness of their models’ predictions to present bias can use these closed-form expressions to approximate the behavior of present-biased agents in discrete time.

8 Conclusion

I present a new set of continuous-time methods for tractably modeling consumers with present bias. I use this methodological innovation to analytically characterize the effect of present bias on consumption-saving decisions and welfare. Along the way I uncover a variety of new findings, many of which diverge from conventional wisdom on present bias. These include an observational equivalence between sophisticates and naifs, an irrelevance result that present bias may not affect the demand for illiquid

⁴²For details, readers are encouraged to examine the extensive resources in Achdou et al. (2022).

assets, and the present-bias dilemma.

Given both the large welfare cost of present bias and the difficulty that individuals face in alleviating this cost, one particularly important subject for future analysis is the extent to which policy interventions can mitigate the present-bias dilemma. More broadly, the IG methods presented in this paper open many pathways for future research by enabling present bias to be tractably incorporated into frontier consumption-saving models.

References

- Acharya, Subas, David Jimenez Gomez, Dmitrii Rachinskii, and Alejandro Rivera**, “Present-Bias and the Value of Sophistication,” *SSRN 3722608*, 2022.
- Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll**, “Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach,” *The Review of Economic Studies*, 2022, 89 (1), 45–86.
- Aguiar, Mark A., Mark Bilis, and Corina Boar**, “Who are the Hand-to-Mouth?,” *NBER w26643*, 2020.
- Ainslie, George**, *Picoeconomics: The Strategic Interaction of Successive Motivational States Within the Person*, Cambridge University Press, 1992.
- Allcott, Hunt, Joshua Kim, Dmitry Taubinsky, and Jonathan Zinman**, “Are High-Interest Loans Predatory? Theory and Evidence from Payday Lending,” *The Review of Economic Studies*, 2022, 89 (3), 1041–1084.
- Amador, Manuel, Iván Werning, and George-Marios Angeletos**, “Commitment vs. Flexibility,” *Econometrica*, 2006, 74 (2), 365–396.
- Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg**, “The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation,” *Journal of Economic Perspectives*, 2001, 15 (3), 47–68.
- Argento, Robert, Victoria L. Bryant, and John Sabelhaus**, “Early Withdrawals from Retirement Accounts During the Great Recession,” *Contemporary Economic Policy*, 2015, 33 (1), 1–16.
- Auclert, Adrien, Matthew Rognlie, and Ludwig Straub**, “The Intertemporal Keynesian Cross,” *NBER w25020*, 2018.

- Augenblick, Ned**, “Short-Term Time Discounting of Unpleasant Tasks,” *Mimeo*, 2018.
- **and Matthew Rabin**, “An Experiment on Time Preference and Misprediction in Unpleasant Tasks,” *The Review of Economic Studies*, 2019, 86 (3), 941–975.
- Banerjee, Abhijit and Sendhil Mullainathan**, “The Shape of Temptation: Implications for the Economic Lives of the Poor,” *NBER w15973*, 2010.
- Barles, Guy and Panagiotis E. Souganidis**, “Convergence of Approximation Schemes for Fully Nonlinear Second Order Equations,” *Asymptotic Analysis*, 1991, 4 (3), 271–283.
- Barro, Robert J.**, “Ramsey Meets Laibson in the Neoclassical Growth Model,” *The Quarterly Journal of Economics*, 1999, 114 (4), 1125–1152.
- Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra**, “House Prices and Consumer Spending,” *The Review of Economic Studies*, 2018, 85 (3), 1502–1542.
- Bernheim, B. Douglas and Antonio Rangel**, “Beyond Revealed Preference: Choice-Theoretic Foundations for Behavioral Welfare Economics,” *The Quarterly Journal of Economics*, 2009, 124 (1), 51–104.
- **and Dmitry Taubinsky**, “Behavioral Public Economics,” in “Handbook of Behavioral Economics: Applications and Foundations 1,” Vol. 1, Elsevier, 2018, pp. 381–516.
- **, Debraj Ray, and Şevin Yeltekin**, “Poverty and Self-Control,” *Econometrica*, 2015, 83 (5), 1877–1911.
- Beshears, John, James Choi, David Laibson, and Peter Maxted**, “Present Bias Causes and Then Dissipates Auto-Enrollment Savings Effects,” *AEA Papers and Proceedings*, 2022, 112, 136–41.
- **, James J. Choi, Christopher Clayton, Christopher Harris, David Laibson, and Brigitte C. Madrian**, “Optimal Illiquidity,” *NBER w27459*, 2022.
- **, – , Joshua Hurwitz, David Laibson, and Brigitte C. Madrian**, “Liquidity in Retirement Savings Systems: An International Comparison,” *American Economic Review*, 2015, 105 (5), 420–25.

- Blow, Laura, Martin Browning, and Ian Crawford**, “Non-Parametric Analysis of Time-Inconsistent Preferences,” *The Review of Economic Studies*, 2021, 88 (6), 2687–2734.
- Bond, Philip and Gustav Sigurdsson**, “Commitment Contracts,” *The Review of Economic Studies*, 2018, 85 (1), 194–222.
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer**, “Diagnostic Expectations and Credit Cycles,” *The Journal of Finance*, 2018, 73 (1), 199–227.
- Bryan, Gharad, Dean Karlan, and Scott Nelson**, “Commitment Devices,” *Annual Review Economics*, 2010, 2 (1), 671–698.
- Cao, Dan and Iván Werning**, “Dynamic Savings Choices with Disagreements,” *NBER w22007*, 2016.
- and —, “Saving and Dissaving with Hyperbolic Discounting,” *Econometrica*, 2018, 86 (3), 805–857.
- Carrera, Mariana, Heather Royer, Mark Stehr, Justin Sydnor, and Dmitry Taubinsky**, “Who Chooses Commitment? Evidence and Welfare Implications,” *Mimeo*, 2020.
- Carroll, Christopher D. and Miles S. Kimball**, “On the Concavity of the Consumption Function,” *Econometrica*, 1996, pp. 981–992.
- Chatterjee, Satyajit and Burcu Eyigungor**, “Continuous Markov Equilibria with Quasi-Geometric Discounting,” *Journal of Economic Theory*, 2016, 163, 467–494.
- Choukhmane, Taha**, “Default Options and Retirement Saving Dynamics,” *Mimeo*, 2019.
- Cohen, Jonathan, Keith Marzilli Ericson, David Laibson, and John Myles White**, “Measuring Time Preferences,” *Journal of Economic Literature*, 2020, 58 (2), 299–347.
- Coyne, David, Itzik Fadlon, and Tommaso Porzio**, “Measuring Valuation of Liquidity with Penalized Withdrawals,” *NBER w30007*, 2022.
- DellaVigna, Stefano**, “Psychology and Economics: Evidence from the Field,” *Journal of Economic Literature*, 2009, 47 (2), 315–72.
- , “Structural Behavioral Economics,” in “Handbook of Behavioral Economics: Applications and Foundations 1,” Vol. 1, Elsevier, 2018, pp. 613–723.

- **and Ulrike Malmendier**, “Contract Design and Self-Control: Theory and Evidence,” *The Quarterly Journal of Economics*, 2004, *119* (2), 353–402.
- **and –**, “Paying Not to Go to the Gym,” *American Economic Review*, 2006, *96* (3), 694–719.
- Fagereng, Andreas, Martin Blomhoff Holm, Benjamin Moll, and Gisle Natvik**, “Saving Behavior Across the Wealth Distribution: The Importance of Capital Gains,” *NBER w26588*, 2019.
- Flavin, Marjorie and Shinobu Nakagawa**, “A Model of Housing in the Presence of Adjustment Costs: A Structural Interpretation of Habit Persistence,” *American Economic Review*, 2008, *98* (1), 474–95.
- Frederick, Shane, George Loewenstein, and Ted O’Donoghue**, “Time Discounting and Time Preference: A Critical Review,” *Journal of Economic Literature*, 2002, *40* (2), 351–401.
- Gabaix, Xavier**, “Behavioral Macroeconomics via Sparse Dynamic Programming,” Technical Report 2016.
- **and David Laibson**, “Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets,” *The Quarterly Journal of Economics*, 2006, *121* (2), 505–540.
- Galperti, Simone**, “Commitment, Flexibility, and Optimal Screening of Time Inconsistency,” *Econometrica*, 2015, *83* (4), 1425–1465.
- Ganong, Peter and Pascal Noel**, “Consumer Spending During Unemployment: Positive and Normative Implications,” *American Economic Review*, 2019, *109* (7), 2383–2424.
- Gathergood, John**, “Self-Control, Financial Literacy and Consumer Over-Indebtedness,” *Journal of Economic Psychology*, 2012, *33* (3), 590–602.
- Gottlieb, Daniel and Xingtang Zhang**, “Long-Term Contracting with Time-Inconsistent Agents,” *Econometrica*, 2021, *89* (2), 793–824.
- Grenadier, Steven R. and Neng Wang**, “Investment Under Uncertainty and Time-Inconsistent Preferences,” *Journal of Financial Economics*, 2007, *84* (1), 2–39.

- Guerrieri, Veronica, Guido Lorenzoni, and Marta Prato**, “Slow Household Deleveraging,” *Journal of the European Economic Association*, 2020, 18 (6), 2755–2775.
- Harris, Christopher and David Laibson**, “Dynamic Choices of Hyperbolic Consumers,” *Econometrica*, 2001, 69 (4), 935–957.
- **and** —, “Hyberbolic Discounting and Consumption,” *Econometric Society Monographs*, 2003, 35, 258–297.
- **and** —, “Instantaneous Gratification,” *Mimeo*, 2004.
- **and** —, “Instantaneous Gratification,” *The Quarterly Journal of Economics*, 2013, 128 (1), 205.
- Heidhues, Paul and Botond Köszegi**, “Exploiting Naivete About Self-Control in the Credit Market,” *American Economic Review*, 2010, 100 (5), 2279–2303.
- Ilut, Cosmin L. and Rosen Valchev**, “Economic Agents as Imperfect Problem Solvers,” *NBER w27820*, 2020.
- John, Anett**, “When Commitment Fails: Evidence From a Field Experiment,” *Management Science*, 2020, 66 (2), 503–529.
- Kaplan, Greg and Giovanni L. Violante**, “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica*, 2014, 82 (4), 1199–1239.
- , **Benjamin Moll, and Giovanni L. Violante**, “Monetary Policy According to HANK,” *American Economic Review*, 2018, 108 (3), 697–743.
- Krusell, Per and Anthony A. Smith**, “Consumption–Savings Decisions with Quasi-Geometric Discounting,” *Econometrica*, 2003, 71 (1), 365–375.
- , **Toshihiko Mukoyama, Ayşegül Şahin, and Anthony A. Smith Jr.**, “Revisiting the Welfare Effects of Eliminating Business Cycles,” *Review of Economic Dynamics*, 2009, 12 (3), 393–404.
- Kuchler, Theresa and Michaela Pagel**, “Sticking to Your Plan: The Role of Present Bias for Credit Card Paydown,” *Journal of Financial Economics*, 2020.
- Laibson, David**, “Self-Control and Saving,” *Mimeo*, 1994.
- , “Golden Eggs and Hyperbolic Discounting,” *The Quarterly Journal of Economics*, 1997, 112 (2), 443–478.

- , “Life-Cycle Consumption and Hyperbolic Discount Functions,” *European Economic Review*, 1998, 42 (3-5), 861–871.
- , “Why Don’t Present-Biased Agents Make Commitments?,” *American Economic Review*, 2015, 105 (5), 267–72.
- **and Peter Maxted**, “The Beta-Delta-DELTA Sweet Spot,” *NBER w30822*, 2022.
- , – , **and Benjamin Moll**, “Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy,” *NBER w29094*, 2021.
- , **Sean Chanwook Lee, Peter Maxted, Andrea Repetto, and Jeremy Tobacman**, “Estimating Discount Functions with Consumption Choices over the Life-cycle,” *Mimeo*, 2023.
- Lee, Sean Chanwook and Peter Maxted**, “Household Borrowing and MPCs in Heterogeneous-Agent Models,” *SSRN 4389878*, 2023.
- Lian, Chen**, “Mistakes in Future Consumption, High MPCs Now,” *Mimeo*, 2021.
- Lucas, Robert E.**, *Models of Business Cycles*, Vol. 26, Basil Blackwell Oxford, 1987.
- Luttmer, Erzo G.J. and Thomas Mariotti**, “Subjective Discounting in an Exchange Economy,” *Journal of Political Economy*, 2003, 111 (5), 959–989.
- Madrian, Brigitte C. and Dennis F. Shea**, “The Power of Suggestion: Inertia in 401 (k) Participation and Savings Behavior,” *The Quarterly Journal of Economics*, 2001, 116 (4), 1149–1187.
- Maskin, Eric and Jean Tirole**, “Markov Perfect Equilibrium,” *Journal of Economic Theory*, 2001, 100 (2), 191–219.
- Maxted, Peter**, “A Macro-Finance Model with Sentiment,” *Mimeo*, 2020.
- McClure, Samuel M., Keith M. Ericson, David I. Laibson, George Loewenstein, and Jonathan D. Cohen**, “Time Discounting for Primary Rewards,” *Journal of Neuroscience*, 2007, 27 (21), 5796–5804.
- McKay, Alisdair and Johannes F. Wieland**, “Lumpy Durable Consumption Demand and the Limited Ammunition of Monetary Policy,” *Econometrica*, 2021, 89 (6), 2717–2749.
- Meier, Stephan and Charles Sprenger**, “Present-Biased Preferences and Credit Card Borrowing,” *American Economic Journal: Applied Economics*, 2010, 2 (1), 193–210.

- Moser, Christian and Pedro Olea de Souza e Silva**, “Optimal Paternalistic Savings Policies,” *SSRN 2959844*, 2019.
- O’Donoghue, Ted and Matthew Rabin**, “Doing It Now or Later,” *American Economic Review*, 1999, 89 (1), 103–124.
- and —, “Choice and Procrastination,” *The Quarterly Journal of Economics*, 2001, 116 (1), 121–160.
- Rivera, Alejandro**, “Contracting with a Present-Biased Agent: Sannikov Meets Laibson,” *SSRN 4009513*, 2022.
- Shigeta, Yuki**, “Quasi-Hyperbolic Discounting under Recursive Utility and Consumption-Investment Decisions,” *SSRN 3672999*, 2020.
- Skiba, Paige Marta and Jeremy Tobacman**, “Payday Loans, Uncertainty, and Discounting: Explaining Patterns of Borrowing, Repayment, and Default,” *Mimeo*, 2018.
- Strotz, Robert Henry**, “Myopia and Inconsistency in Dynamic Utility Maximization,” *The Review of Economic Studies*, 1956, 23 (3), 165–180.
- Thaler, Richard**, “Mental Accounting and Consumer Choice,” *Marketing Science*, 1985, 4 (3), 199–214.
- Tobacman, Jeremy**, “The Partial Naivete Euler Equation,” *Mimeo*, 2007.
- Tourin, Agnes**, “An Introduction to Finite Difference Methods for PDEs in Finance,” *Optimal Stochastic Target Problems, and Backward SDE*, 2013.

****ONLINE APPENDIX****

A Proofs

Throughout this appendix I assume that the reader understands the construction of the \hat{u} agent. Details of the \hat{u} construction are given in Section 4, and in Harris and Laibson (2013).

A.1 Proof of Proposition 1

The proof of Proposition 1 relies on Proposition 2.

Lemma 11. *The value function of the \hat{u} agent, denoted $\hat{v}(x)$, is unique.*

Proof. See Harris and Laibson (2013) for full details. Intuitively, the \hat{u} agent is an exponential discounter who optimally choses \hat{c} and \hat{d} to maximize $\hat{v}(x)$. Since there is only one maximal value function, $\hat{v}(x)$ must be unique. \square

The proof of Proposition 1 follows immediately from Lemma 11 and Proposition 2. In particular, Lemma 11 says \hat{v} is unique, and Proposition 2 says $v = \hat{v}$. Hence, v is unique.

A.2 Proof of Proposition 2

I now prove value function equivalence between the IG agent and the \hat{u} agent. The proof presented below builds off of related proofs in Harris and Laibson (2013) and Laibson and Maxted (2022).

Before continuing, I emphasize that most of the complexity in this proof arises when \underline{b} binds in equilibrium. The proof simplifies considerably when \underline{b} does not bind in equilibrium, as assumed in most of the paper (see Assumption 1).

The \hat{u} Agent's Bellman Equation. To begin, recall from equation (16) that the \hat{u} agent faces a choice at the borrowing constraint \underline{b} . The \hat{u} agent can either set $\hat{c} \leq \psi(y + r(\underline{b})\underline{b})$ or $\hat{c} = y + r(\underline{b})\underline{b}$. The former choice only earns utility \hat{u}_+ , but allows

the agent to save away from the constraint. The latter choice earns the “utility boost,” but requires the agent to stay at \underline{b} . I will refer to the former choice as “continuing,” and the latter choice as “stopping.”

Lemma 12. *The \hat{u} agent will choose to continue at \underline{b} when $\hat{v}_b(x) > \frac{1}{\beta}(y + r(\underline{b})\underline{b})^{-\gamma}$. Otherwise, the \hat{u} agent will choose to stop.*

Proof. The value function of the \hat{u} agent at \underline{b} can be expressed as:

$$\begin{aligned} \rho \hat{v}(\underline{b}, a, y) = \max \left\{ u(y + r(\underline{b})\underline{b}), \max_{\hat{c} \leq \psi(y + r(\underline{b})\underline{b})} \hat{u}_+(\hat{c}) + \hat{v}_b(x)(y + r(\underline{b})\underline{b} - \hat{c}) \right\} \\ + \hat{v}_a(x)(r^a a) + \frac{1}{2} \hat{v}_{aa}(x)(a\sigma^a)^2 \\ + \sum_{y' \neq y} \lambda^{y \rightarrow y'} (\hat{v}(\underline{b}, a, y') - \hat{v}(\underline{b}, a, y)), \end{aligned} \quad (21)$$

where the above equation imposes $\hat{d} = 0$ (recall that I assume that the agent cannot deposit/withdraw from the illiquid asset at $b = \underline{b}$).

Intuitively, the first line of equation (21) captures the choice that the \hat{u} agent faces at \underline{b} . The left branch of the first line is the “stopping” option: the agent sets $\hat{c} = y + r(\underline{b})\underline{b}$ and earns utility $u(y + r(\underline{b})\underline{b})$. The right branch is the “continuing” option: the agent chooses $\hat{c} \leq \psi(y + r(\underline{b})\underline{b})$ and earns utility $\hat{u}_+(\hat{c})$, but also accumulates liquid wealth ($db_t \geq 0$), which yields the additional term $\hat{v}_b(x)(y + r(\underline{b})\underline{b} - \hat{c})$.

In the right branch the \hat{u} agent chooses \hat{c} such that $\hat{u}'_+(\hat{c}) = \hat{v}_b(x)$, which implies $\hat{c} = \psi(\beta \hat{v}_b(x))^{-\frac{1}{\gamma}}$. Using this property, one can show that the \hat{u} agent is indifferent between the two choices when $\hat{v}_b(x) = \frac{1}{\beta}(y + r(\underline{b})\underline{b})^{-\gamma}$ (which implies $\hat{c}(x) = \psi(y + r(\underline{b})\underline{b})$). The \hat{u} agent chooses to “continue” when $\hat{v}_b(x) > \frac{1}{\beta}(y + r(\underline{b})\underline{b})^{-\gamma}$, and chooses to “stop” when $\hat{v}_b(x) < \frac{1}{\beta}(y + r(\underline{b})\underline{b})^{-\gamma}$. At the point of indifference, I assume that the \hat{u} agent stops. □

Now, the \hat{u} agent’s value function in equation (17) can be expressed recursively

as follows:

$$\begin{aligned}\rho\hat{v}(x) &= \hat{u}(\hat{c}(x), x) + \hat{v}_b(x) \left(y + r(b)b - \hat{d}(x) - \chi(\hat{d}(x), a) - \hat{c}(x) \right) \\ &\quad + \hat{v}_a(x) \left(r^a a + \hat{d}(x) \right) + \frac{1}{2} \hat{v}_{aa}(x) (a\sigma^a)^2 \\ &\quad + \sum_{y' \neq y} \lambda^{y \rightarrow y'} (\hat{v}(b, a, y') - \hat{v}(b, a, y)),\end{aligned}\tag{22}$$

subject to the optimality conditions

$$\hat{c}(x) = \begin{cases} \psi(\beta\hat{v}_b(x))^{-\frac{1}{\gamma}} & \text{if } b > \underline{b} \\ \psi(\beta\hat{v}_b(x))^{-\frac{1}{\gamma}} & \text{if } b = \underline{b} \text{ and } \hat{v}_b(x) > \frac{1}{\beta}(y + r(\underline{b})\underline{b})^{-\gamma}, \text{ and} \\ y + r(\underline{b})\underline{b} & \text{if } b = \underline{b} \text{ and } \hat{v}_b(x) \leq \frac{1}{\beta}(y + r(\underline{b})\underline{b})^{-\gamma} \end{cases}\tag{23}$$

$$\chi_d(\hat{d}(x), a) = \begin{cases} \frac{\hat{v}_a(x)}{\hat{v}_b(x)} - 1 & \text{if } b > \underline{b} \\ 0 & \text{if } b = \underline{b} \end{cases},\tag{24}$$

where Lemma 12 allows for the \hat{u} agent's consumption to be defined as in equation (23). Note that the first two lines of equation (23) are equivalent to the \hat{u} agent choosing \hat{c} such that $\hat{u}'_+(\hat{c}) = \hat{v}_b(x)$, while the third line of equation (23) imposes that $\hat{c} = y + r(\underline{b})\underline{b}$ if the \hat{u} agent is at \underline{b} and chooses to stop.

Proof Intuition. The intuition for this proof is as follows. Assume that $v(x) = \hat{v}(x)$ and $b > \underline{b}$. Then, equations (10) and (22) can be combined to yield:

$$u(c(x)) - v_b(x)c(x) = \hat{u}_+(\hat{c}(x)) - \hat{v}_b(x)\hat{c}(x).$$

Utility function \hat{u} is reverse-engineered so that this condition holds.

The IG Agent: A Modified Bellman Equation. Following Theorem 2 of Harris and Laibson (2013), let $f_+(\alpha)$ be the unique value of c satisfying $u'(c) = \alpha$. Let $h_+(\alpha) = u(f_+(\beta\alpha)) - \alpha f_+(\beta\alpha)$. Since the IG agent sets $u'(c(x)) = \beta v_b(x)$ for $b > \underline{b}$, it is the case that $h_+(v_b(x)) = u(f_+(\beta v_b(x))) - v_b(x)f_+(\beta v_b(x)) = u(c(x)) - v_b(x)c(x)$.

Next, let $\underline{f}(\alpha, x)$ be the unique value of c satisfying $u'(c) = \max\{\alpha, u'(y + r(\underline{b})\underline{b})\}$.

Let $\underline{h}(\alpha, x) = u(\underline{f}(\beta\alpha, x)) - \alpha\underline{f}(\beta\alpha, x)$. Again, since the IG agent sets $u'(c(x)) = \max \left\{ \beta v_b(x), u'(y+r(\underline{b})\underline{b}) \right\}$ for $b = \underline{b}$, it is the case that $\underline{h}(v_b(x), x) = u(\underline{f}(\beta v_b(x), x)) - v_b(x)\underline{f}(\beta v_b(x), x) = u(c(x)) - v_b(x)c(x)$.

Define:

$$h(\alpha, x) = \begin{cases} h_+(\alpha) & \text{if } b > \underline{b} \\ \underline{h}(\alpha, x) & \text{if } b = \underline{b} \end{cases}.$$

Function h can be used to rewrite the Bellman equation of the IG agent (equations (10) – (12)) as follows:

$$\begin{aligned} \rho v(x) = & h(v_b(x), x) + v_b(x) (y + r(b)b - d(x) - \chi(d(x), a)) \\ & + v_a(x) (r^a a + d(x)) + \frac{1}{2} v_{aa}(x) (a\sigma^a)^2 \\ & + \sum_{y' \neq y} \lambda^{y \rightarrow y'} (v(b, a, y') - v(b, a, y)), \end{aligned} \quad (25)$$

subject to the optimality condition:

$$\chi_d(d(x), a) = \begin{cases} \frac{v_a(x)}{v_b(x)} - 1 & \text{if } b > \underline{b} \\ 0 & \text{if } b = \underline{b} \end{cases}. \quad (26)$$

The \hat{u} Agent: A Modified Bellman Equation. For the \hat{u} agent, let $\hat{f}_+(\alpha)$ be the unique value of \hat{c} satisfying $\hat{u}'_+(\hat{c}) = \alpha$. Let $\hat{h}_+(\alpha) = \hat{u}_+(\hat{f}_+(\alpha)) - \alpha\hat{f}_+(\alpha)$. Since the \hat{u} agent sets $\hat{u}'_+(\hat{c}(x)) = \hat{v}_b(x)$ for $b > \underline{b}$, it is the case that $\hat{h}_+(\hat{v}_b(x)) = \hat{u}_+(\hat{f}_+(\hat{v}_b(x))) - \hat{v}_b(x)\hat{f}_+(\hat{v}_b(x)) = \hat{u}_+(\hat{c}(x)) - \hat{v}_b(x)\hat{c}(x)$.

Next, let $\underline{\hat{f}}(\alpha, x)$ be the unique value of \hat{c} satisfying $\hat{u}'_+(\hat{c}) = \alpha$ if $\hat{c} < \psi(y + r(\underline{b})\underline{b})$, and let $\underline{\hat{f}}(\alpha, x) = y + r(\underline{b})\underline{b}$ otherwise. Let $\underline{\hat{h}}(\alpha, x) = \hat{u}(\underline{\hat{f}}(\alpha, x), x) - \alpha\underline{\hat{f}}(\alpha, x)$. Again, now using Lemma 12, it is the case that $\underline{\hat{h}}(v_b(x), x) = \hat{u}(\underline{\hat{f}}(v_b(x), x), x) - v_b(x)\underline{\hat{f}}(v_b(x), x) = \hat{u}(\hat{c}(x), x) - \hat{v}_b(x)\hat{c}(x)$.

Define:

$$\hat{h}(\alpha, x) = \begin{cases} \hat{h}_+(\alpha) & \text{if } b > \underline{b} \\ \underline{\hat{h}}(\alpha, x) & \text{if } b = \underline{b} \end{cases}.$$

Function \hat{h} can be used to rewrite the Bellman equation of the \hat{u} agent (equations

(22) – (24)) as follows:

$$\begin{aligned}\rho\hat{v}(x) &= \hat{h}(\hat{v}_b(x), x) + \hat{v}_b(x) \left(y + r(b)b - \hat{d}(x) - \chi(\hat{d}(x), a) \right) \\ &\quad + \hat{v}_a(x) \left(r^a a + \hat{d}(x) \right) + \frac{1}{2} \hat{v}_{aa}(x) (a\sigma^a)^2 \\ &\quad + \sum_{y' \neq y} \lambda^{y \rightarrow y'} (\hat{v}(b, a, y') - \hat{v}(b, a, y)),\end{aligned}\tag{27}$$

subject to the optimality condition:

$$\chi_d(\hat{d}(x), a) = \begin{cases} \frac{\hat{v}_a(x)}{\hat{v}_b(x)} - 1 & \text{if } b > \underline{b} \\ 0 & \text{if } b = \underline{b} \end{cases}.\tag{28}$$

From inspection, we can see that equations (25) and (26) are identical to equations (27) and (28) if and only if $h(\alpha, x)$ and $\hat{h}(\alpha, x)$ are the same. This can be confirmed directly.

A.3 Proof of Proposition 3

From equation (13), the (potentially naive) IG agent sets $u'(c(x)) = \beta v_b^E(x)$. By Proposition 2, one can construct a \hat{u} agent using β^E such that $v^E(x) = \hat{v}(x)$. This \hat{u} agent chooses consumption such that $\hat{u}'_+(\hat{c}(x)) = \hat{v}_b(x)$. This implies that $\hat{v}_b(x) = \frac{(\psi^E)^\gamma}{\beta^E} \hat{c}(x)^{-\gamma}$, where $\psi^E = \frac{\gamma - (1 - \beta^E)}{\gamma}$.

Using the value function equivalence property from Proposition 2 that $\hat{v}(x) = v^E(x)$:

$$u'(c(x)) = \beta \frac{(\psi^E)^\gamma}{\beta^E} \hat{c}(x)^{-\gamma}.$$

Rearranging gives

$$c(x) = \left(\frac{\beta^E}{\beta} \right)^{\frac{1}{\gamma}} \frac{1}{\psi^E} \times \hat{c}(x).$$

To complete the proof, note that the \hat{u} agent behaves identically to the standard exponential agent when Assumption 1 holds. This implies that the \hat{u} agent sets

$\hat{c}(x) = \check{c}(x)$ regardless of β and β^E . Therefore $c(x) = \left(\frac{\beta^E}{\beta}\right)^{\frac{1}{\gamma}} \frac{1}{\psi^E} \times \check{c}(x)$, as desired.

To see when consumption is increasing in naivete, consider:

$$\begin{aligned} \frac{\partial c(x)}{\partial \beta^E} &\propto \frac{1}{\gamma} \left(\frac{\beta^E}{\beta}\right)^{\frac{1-\gamma}{\gamma}} \frac{1}{\beta} \frac{1}{\psi^E} - \left(\frac{\beta^E}{\beta}\right)^{\frac{1}{\gamma}} \frac{1}{\psi^E} \frac{1}{\gamma - (1 - \beta^E)} \\ &\propto \frac{1}{\beta^E} - \frac{1}{\psi^E} \end{aligned}$$

For $\beta^E < 1$, one can show that $\psi^E > \beta^E$ when $\gamma > 1$, and $\psi^E < \beta^E$ when $\gamma < 1$. Thus, consumption is increasing in naivete when $\gamma > 1$, and decreasing in naivete when $\gamma < 1$.

A.4 Proof of Proposition 6

Generalizing the Proposition for Naivete. First, I extend Proposition 6 to the case of naivete. See Tobacman (2007) for a discrete-time analysis, and Laibson et al. (2021) for an application of this result.

Proposition. *Let $\varsigma(b_t)$ denote the marginal interest rate that the agent earns on their liquid wealth of b_t (see footnote 28 for details). Whenever $c(x_t)$ and $r(b_t)$ are locally differentiable in b , consumption satisfies the following Euler equation:*

$$\frac{\mathbb{E}_t(du'(c(x_t))) / dt}{u'(c(x_t))} = \left[\rho + \left((1 - \beta^E) \left(\frac{\beta}{\beta^E} \right)^{\frac{1}{\gamma}} + \gamma \left(1 - \left(\frac{\beta}{\beta^E} \right)^{\frac{1}{\gamma}} \right) \right) c_b(x_t) \right] - \varsigma(b_t). \quad (29)$$

Equation (29) simplifies in three special cases: complete sophistication ($\beta^E = \beta$), complete naivete ($\beta^E = 1$), and log utility ($\gamma = 1$). With complete sophistication or $\gamma = 1$, the Euler equation in (19) is recovered. With complete naivete, the Euler equation is:

$$\frac{\mathbb{E}_t(du'(c(x_t))) / dt}{u'(c(x_t))} = \left[\rho + \gamma \left(1 - \beta^{\frac{1}{\gamma}} \right) c_b(x_t) \right] - \varsigma(b_t).$$

As in the case with sophistication, the effective discount rate in brackets in equation (29) varies with the instantaneous MPC of the IG agent.

Proof of Sophisticated Case (Proposition 6). This proof extends the $\beta = 1$ case of Achdou et al. (2022). A similar result is given in Harris and Laibson (2004). Taking a derivative of (10) with respect to b gives

$$\begin{aligned}\rho v_b(x) &= u'(c(x))c_b(x) + v_{bb}(x)(y + r(b)b - d(x) - \chi(d(x), a) - c(x)) \\ &\quad + v_b(x)(r'(b)b + r(b) - d_b(x) - \chi_d(d(x), a)d_b(x) - c_b(x)) \\ &\quad + v_{ab}(x)(r^a a + d(x)) + v_a(x)d_b(x) + \frac{1}{2}v_{aab}(x)(a\sigma^a)^2 \\ &\quad + \sum_{y' \neq y} \lambda^{y \rightarrow y'}(v_b(b, a, y') - v_b(b, a, y)).\end{aligned}$$

Using the optimality condition (12) that $(\chi_d(d(x), a) + 1)v_b(x) = v_a(x)$ simplifies this to

$$\begin{aligned}\rho v_b(x) &= u'(c(x))c_b(x) + v_{bb}(x)(y + r(b)b - d(x) - \chi(d(x), a) - c(x)) \\ &\quad + v_b(x)(r'(b)b + r(b) - c_b(x)) \\ &\quad + v_{ab}(x)(r^a a + d(x)) + \frac{1}{2}v_{aab}(x)(a\sigma^a)^2 \\ &\quad + \sum_{y' \neq y} \lambda^{y \rightarrow y'}(v_b(b, a, y') - v_b(b, a, y)).\end{aligned}$$

Applying the first-order condition (11):

$$\begin{aligned}[\rho - (r'(b)b + r(b)) + (1 - \beta)c_b(x)]u'(c(x)) &= u''(c(x))c_b(x)(y + r(b)b - d(x) - \chi(d(x), a) - c(x)) \\ &\quad + u''(c(x))c_a(x)(r^a a + d(x)) \\ &\quad + \frac{1}{2}(u'''(c(x))(c_a(x))^2 + u''(c(x))c_{aa}(x))(a\sigma^a)^2 \\ &\quad + \sum_{y' \neq y} \lambda^{y \rightarrow y'}(u'(c(b, a, y')) - u'(c(b, a, y))).\end{aligned}$$

Applying Itô's Lemma to $u'(c(x_t))$ gives

$$\begin{aligned}
\mathbb{E}_t[du'(c(x_t))]/dt &= u''(c(x_t))c_b(x_t) (y_t + r(b_t)b_t - d(x_t) - \chi(d(x_t), a_t) - c(x_t)) \\
&\quad + u''(c(x_t))c_a(x_t) (r_t^a a_t + d(x_t)) \\
&\quad + u''(c(x_t)) \left(\frac{1}{2} c_{aa}(x_t) (a_t \sigma^a)^2 \right) \\
&\quad + \frac{1}{2} u'''(c(x_t)) (c_a(x_t) a_t \sigma^a)^2 \\
&\quad + \sum_{y' \neq y_t} \lambda^{y_t \rightarrow y'} (u'(c(b_t, a_t, y')) - u'(c(b_t, a_t, y_t))).
\end{aligned}$$

Plugging this in to the above equation results in

$$[\rho - (r'(b_t)b_t + r(b_t)) + (1 - \beta)c_b(x_t)] u'(c(x_t)) = \mathbb{E}_t[du'(c(x_t))]/dt.$$

Rearranging and using the property that $\varsigma(b_t) = r'(b_t)b_t + r(b_t)$ gives equation (19).

Proof of Generalized Proposition with Naivete. The proof begins the same as above, except now the agent's value function is based on their expected behavior. Taking a derivative of the expected value function with respect to b gives

$$\begin{aligned}
\rho v_b^E(x) &= u'(c^E(x))c_b^E(x) + v_{bb}^E(x) (y + r(b)b - d^E(x) - \chi(d^E(x), a) - c^E(x)) \\
&\quad + v_b^E(x) (r'(b)b + r(b) - d_b^E(x) - \chi_a(d^E(x), a)d_b^E(x) - c_b^E(x)) \\
&\quad + v_{ab}^E(x) (r^a a + d^E(x)) + v_a^E(x)d_b^E(x) + \frac{1}{2}v_{aab}^E(x)(a\sigma^a)^2 \\
&\quad + \sum_{y' \neq y} \lambda^{y \rightarrow y'} (v_b^E(b, a, y') - v_b^E(b, a, y)).
\end{aligned}$$

Using the perceived optimality condition (12) that $(\chi_d(d^E(x), a) + 1)v_b^E(x) = v_a^E(x)$ simplifies this to

$$\begin{aligned}\rho v_b^E(x) &= u'(c^E(x))c_b^E(x) + v_{bb}^E(x) (y + r(b)b - d^E(x) - \chi(d^E(x), a) - c^E(x)) \\ &\quad + v_b^E(x) (r'(b)b + r(b) - c_b^E(x)) \\ &\quad + v_{ab}^E(x) (r^a a + d^E(x)) + \frac{1}{2}v_{aab}^E(x)(a\sigma^a)^2 \\ &\quad + \sum_{y' \neq y} \lambda^{y \rightarrow y'} (v_b^E(b, a, y') - v_b^E(b, a, y)).\end{aligned}$$

Applying the first-order condition (13) for the naif's realized consumption:

$$\begin{aligned}[\rho - (r'(b)b + r(b))] u'(c(x)) + c_b^E(x) (u'(c(x)) - \beta u'(c^E(x))) \\ = u''(c(x))c_b(x) (y + r(b)b - d^E(x) - \chi(d^E(x), a) - c^E(x)) \\ + u''(c(x))c_a(x) (r^a a + d^E(x)) \\ + \frac{1}{2} (u'''(c(x))(c_a(x))^2 + u''(c(x))c_{aa}(x)) (a\sigma^a)^2 \\ + \sum_{y' \neq y} \lambda^{y \rightarrow y'} (u'(c(b, a, y')) - u'(c(b, a, y))).\end{aligned}$$

Next, since the naif perceives that $u'(c^E(x)) = \beta^E v_b^E(x)$, while equation (13) imposes that $u'(c(x)) = \beta v_b^E(x)$, we have $\beta u'(c^E(x)) = \beta^E u'(c(x))$. Using this:

$$\begin{aligned}[\rho - (r'(b)b + r(b)) + (1 - \beta^E)c_b^E(x)] u'(c(x)) &= u''(c(x))c_b(x) (y + r(b)b - d^E(x) - \chi(d^E(x), a) - c^E(x)) \\ &\quad + u''(c(x))c_a(x) (r^a a + d^E(x)) \\ &\quad + \frac{1}{2} (u'''(c(x))(c_a(x))^2 + u''(c(x))c_{aa}(x)) (a\sigma^a)^2 \\ &\quad + \sum_{y' \neq y} \lambda^{y \rightarrow y'} (u'(c(b, a, y')) - u'(c(b, a, y))).\end{aligned}$$

As above, applying Itô's Lemma to $u'(c(x_t))$ gives

$$\begin{aligned}\mathbb{E}_t[du'(c(x_t))]/dt &= u''(c(x_t))c_b(x_t) (y_t + r(b_t)b_t - d(x_t) - \chi(d(x_t), a_t) - c(x_t)) \\ &\quad + u''(c(x_t))c_a(x_t) (r_t^a a_t + d(x_t)) \\ &\quad + u''(c(x_t)) \left(\frac{1}{2} c_{aa}(x_t) (a_t \sigma^a)^2 \right) \\ &\quad + \frac{1}{2} u'''(c(x_t)) (c_a(x_t) a_t \sigma^a)^2 \\ &\quad + \sum_{y' \neq y_t} \lambda^{y_t \rightarrow y'} (u'(c(b_t, a_t, y')) - u'(c(b_t, a_t, y_t))).\end{aligned}$$

Plugging this in to the above equation results in

$$\begin{aligned}[\rho - (r'(b_t)b_t + r(b_t)) + (1 - \beta^E)c_b^E(x_t)]u'(c(x_t)) \\ = u''(c(x_t))c_b(x_t) (y + r(b_t)b_t - d^E(x_t) - \chi(d^E(x_t), a_t) - c^E(x_t)) \\ - u''(c(x_t))c_b(x_t) (y_t + r(b_t)b_t - d(x_t) - \chi(d(x_t), a_t) - c(x_t)) \\ + \mathbb{E}_t[du'(c(x_t))]/dt.\end{aligned}$$

Using the property that $d^E(x) = d(x)$, this simplifies immediately to

$$[\rho - (r'(b)b + r(b)) + (1 - \beta^E)c_b^E(x)]u'(c(x)) = u''(c(x))c_b(x) (c(x) - c^E(x)) + \mathbb{E}_t[du'(c(x_t))]/dt.$$

The naif perceives they set $u'(c^E(x)) = \beta^E v_b^E(x)$ while they actually set $u'(c(x)) = \beta v_b^E(x)$, which implies $c^E(x) = \left(\frac{\beta}{\beta^E}\right)^{\frac{1}{\gamma}} c(x)$. This gives

$$\begin{aligned}\left[\rho - (r'(b_t)b_t + r(b_t)) + (1 - \beta^E) \left(\frac{\beta}{\beta^E} \right)^{\frac{1}{\gamma}} c_b(x_t) \right] u'(c(x_t)) &= u''(c(x_t))c_b(x_t)c(x_t) \left(1 - \left(\frac{\beta}{\beta^E} \right)^{\frac{1}{\gamma}} \right) \\ &\quad + \mathbb{E}_t[du'(c(x_t))]/dt.\end{aligned}$$

Using the property that $-\gamma = \frac{u''(c(x))c(x)}{u'(c(x))}$ and $\varsigma(b_t) = r'(b_t)b_t + r(b_t)$, this can be rearranged to yield

$$\frac{\mathbb{E}_t[du'(c(x_t))]/dt}{u'(c(x_t))} = \left[\rho + \left((1 - \beta^E) \left(\frac{\beta}{\beta^E} \right)^{\frac{1}{\gamma}} + \gamma \left(1 - \left(\frac{\beta}{\beta^E} \right)^{\frac{1}{\gamma}} \right) \right) c_b(x_t) \right] - \varsigma(b_t),$$

which is equation (29) as desired.

A.5 Proof of Proposition 7

Notationally, since Proposition 7 simplifies the economic environment there is now just a single state variable, which is liquid wealth b . This will be reflected in my notation below (e.g., the consumption function will be written as $c(b)$).

Throughout this proof, since Assumption 1 holds the \hat{u} agent consumes identically to the standard exponential agent: $\hat{c}(b) = \check{c}(b)$.

First consider the case where $r \leq \rho$. The Euler equation of the standard exponential agent implies that $\check{c}(b) \geq y + r(b)b$ for all $b \geq 0$ (see Achdou et al. (2022)). Since the sophisticated IG agent sets $c(b) = \frac{1}{\psi}\hat{c}(b) = \frac{1}{\psi}\check{c}(b)$ (see Proposition 3), the IG agent strictly dissaves for all $b \geq 0$ when $r \leq \rho$.

Next consider the case where $r \in (\rho, \frac{\rho}{\beta})$. In this deterministic model the standard exponential agent consumes according to $\check{c}(b) = \frac{\rho - (1-\gamma)r}{\gamma}(b + \frac{y}{r})$ for $b \geq 0$ (see e.g. Fagereng et al. (2019) for details). The IG agent therefore sets $c(b) = \frac{1}{\psi}\check{c}(b) = \frac{\rho - (1-\gamma)r}{\gamma - (1-\beta)}(b + \frac{y}{r})$ for $b \geq 0$. One can show that $s(b) = y + r(b)b - c(b) < 0$ for $b \geq 0$ whenever $r < \frac{\rho}{\beta}$. Thus, the IG agent strictly dissaves for all $b \geq 0$ when $r \in (\rho, \frac{\rho}{\beta})$.

In both cases the IG agent strictly dissaves for all $b \geq 0$. This means that the IG agent dissaves at $b = 0$, completing the proof that $s(0) < 0$ whenever $r < \frac{\rho}{\beta}$. This holds regardless of how large $\mathcal{W}(0)$ is (indeed, \mathcal{W} doesn't even show up in the proof).

Note that this proof does not rely on some sort of consumption discontinuity at $b = 0$. The consumption function $c(b)$ is continuous at $b = 0$. To show this, recall that the IG agent's value function is given by

$$\rho v(b) = u(c(b)) + v'(b)(y + r(b)b - c(b)).$$

The IG agent sets $u'(c(b)) = \beta v'(b)$. Therefore

$$\rho v(b) = u(c(b)) + \frac{c(b)^{-\gamma}}{\beta}(y + r(b)b - c(b)).$$

Since $v(b)$ is continuous and $r(b)b$ is continuous, $c(b)$ is also continuous at $b = 0$.

A.6 Proof of Proposition 8

Equation (14) characterizes the deposit policy function $d(x)$ for the (possibly naive) IG agent, and equation (24) characterizes the deposit policy function $\hat{d}(x)$ for the \hat{u} agent. By Proposition 2, one can construct a \hat{u} agent using β^E such that $v^E(x) = \hat{v}(x)$. Given this value function equivalence, equations (14) and (24) imply that the IG agent chooses the same illiquid asset policy function as the \hat{u} agent: $d(x) = \hat{d}(x)$. When Assumption 1 holds the \hat{u} agent behaves identically to the standard exponential agent, and therefore $d(x) = \hat{d}(x) = \check{d}(x)$.

A.7 Proof of Proposition 9

Step 1: Value Function Equivalence for the Naive Agent ($\gamma \neq 1$). Recall that the \hat{u} utility function is constructed so that the value function of the sophisticated IG agent is equivalent to the value function of the \hat{u} agent. The first step of this proof generalizes the \hat{u} construction to allow for naivete. Specifically, I now construct a utility function, denoted $\hat{\hat{u}}$, such that the *realized* value function of the (potentially naive) IG agent is equivalent to the value function of the exponential agent with utility function $\hat{\hat{u}}$. I refer to this agent as the $\hat{\hat{u}}$ agent.

Note that when the IG agent is naive ($\beta^E \neq \beta$) their realized value function does not equal their expected value function. As given in the main text, the expected continuation-value function is $v_t^E = \mathbb{E}_t \left[\int_t^\infty e^{-\rho(s-t)} u(c_s^E) ds \right]$. Denote the realized value function by

$$v_t^R = \mathbb{E}_t \left[\int_t^\infty e^{-\rho(s-t)} u(c_s) ds \right].$$

v^R is based on the naif's realized consumption choices, while v^E is based on their

perceived consumption choices.

Let $\hat{u}(c) = \frac{\xi c^{1-\gamma}-1}{1-\gamma}$.⁴³ This is a positive affine transformation of CRRA utility function $u(c)$ whenever $\xi > 0$. When this is the case, the \hat{u} agent will behave identically to the standard exponential agent. Thus, I will directly use $\check{c}(x)$ to refer to the consumption of the \hat{u} agent.

In order to generate value function equivalence between the (possibly naive) IG agent and the \hat{u} agent, I construct \hat{u} so that the following condition holds for all $b > \underline{b}$:

$$u(c(x)) - v_b^R(x)c(x) = \hat{u}(\check{c}(x)) - \hat{v}_b(x)\check{c}(x). \quad (30)$$

Condition (30) ensures that $v^R(x) = \hat{v}(x)$ whenever Assumption 1 holds. See the proof of Proposition 2 for details.

I want to solve for ξ such that equation (30) holds. From Proposition 3, note that $\check{c}(x) = \alpha c(x)$, where $\alpha = \psi^E \left(\frac{\beta}{\beta^E} \right)^{\frac{1}{\gamma}}$. Additionally, the \hat{u} agent sets $\check{c}(x)$ such that $\xi \check{c}(x)^{-\gamma} = \hat{v}_b(x)$. Using these properties in equation (30) gives:

$$\frac{c(x)^{1-\gamma}}{1-\gamma} - \xi \alpha^{-\gamma} c(x)^{1-\gamma} = \frac{\xi (\alpha c(x))^{1-\gamma}}{1-\gamma} - \xi (\alpha c(x))^{1-\gamma}.$$

This can be rearranged to yield:

$$\xi = \frac{\alpha^\gamma}{1 - \gamma + \alpha\gamma}.$$

Note that $\xi = \frac{\psi^\gamma}{\beta}$ in the case of sophistication ($\beta^E = \beta$), in which case $\hat{u}(c) = u(c)$.

Step 2: The Effect of a Consumption Tax. I now introduce a constant perpetual consumption tax of $\tau \in [0, 1)$. Given consumption tax τ , let $\check{c}(x)$ denote the gross consumption expenditure rate of the standard exponential agent. In other words, the agent spends \check{c} to consume $(1 - \tau)\check{c}$, with the rest going to taxes. Here I show that a consumption tax of τ does not affect the standard exponential agent's gross consumption expenditure.

⁴³The construction of \hat{u} is simplified relative to the definition of \hat{u} in equation (16) because this proof assumes from the start that Assumption 1 holds.

With no tax, the standard exponential agent chooses consumption to maximize \check{v} :

$$\check{v}(x) = \max_{\{\check{c}_s\}} \mathbb{E}_t \left[\int_t^\infty e^{-\rho(s-t)} u(\check{c}_s) ds \right].$$

With a consumption tax, the standard exponential agent chooses consumption to maximize:

$$\check{v}(x; \tau) = \max_{\{\check{c}_s\}} \mathbb{E}_t \left[\int_t^\infty e^{-\rho(s-t)} u((1 - \tau)\check{c}_s) ds \right].$$

Note that $u((1 - \tau)c)$ is a positive affine transformation of $u(c)$. Thus, policy function $\check{c}(x)$ is unaffected by consumption tax τ . The only effect of the tax is that $\check{c}(x)$ now denotes gross consumption expenditure, whereas the agent only gets to consume $(1 - \tau)\check{c}(x)$ with the rest going to taxes.

Step 3: The Welfare Effect of Present Bias ($\gamma \neq 1$). Since \hat{u} is a positive affine transformation of u , the \hat{u} agent behaves identically to the standard exponential agent. Additionally, value function equivalence implies that the realized value function of the IG agent equals the value function of the \hat{u} agent whenever \underline{b} does not bind in equilibrium: $v^R(x) = \hat{v}(x)$. This was shown in Step 1 of this proof.

The final step is to derive the consumption tax τ that equates the realized value function of the IG agent ($v^R(x)$) with the value function of the standard exponential agent facing a consumption tax ($\check{v}(x; \tau)$). Using value function equivalence, the realized value function of the IG agent is:

$$v^R(x) = \hat{v}(x) = \mathbb{E}_t \left[\int_t^\infty e^{-\rho(s-t)} \hat{u}(\hat{c}_s) ds \right]. \quad (31)$$

The value function of the standard exponential agent facing a consumption tax is:

$$\check{v}(x; \tau) = \mathbb{E}_t \left[\int_t^\infty e^{-\rho(s-t)} u((1 - \tau)\check{c}_s) ds \right]. \quad (32)$$

The key to this proof is to note that $\hat{c}(x) = \check{c}(x)$. Therefore the consumption path in equation (31) is identical to the gross consumption expenditure path in equation

(32) (this property holds state by state, so it also holds in expectation). Thus, setting equation (31) equal to equation (32) is as simple as finding the value of τ such that:

$$\hat{u}(c) = u((1 - \tau)c).$$

This implies that $\xi = (1 - \tau)^{1-\gamma}$. Rearranging gives

$$\tau = 1 - \left(\frac{\alpha^\gamma}{1 - \gamma + \gamma\alpha} \right)^{\frac{1}{1-\gamma}}.$$

Special Case: $\gamma = 1$. In the special case of $\gamma = 1$ the naif and the sophisticate behave identically (Proposition 3). The realized value function $v^R(x)$ is therefore independent of β^E . So, I calculate the $\gamma = 1$ case under the assumption of sophistication using the \hat{u} agent.

I again derive the consumption tax τ that equates the realized value function of the IG agent ($v^R(x)$) with the value function of the standard exponential agent facing a consumption tax ($\check{v}(x; \tau)$). Using value function equivalence with the \hat{u} agent, the realized value function of the IG agent is:

$$v^R(x) = \hat{v}(x) = \mathbb{E}_t \left[\int_t^\infty e^{-\rho(s-t)} \hat{u}(\hat{c}_s) ds \right]. \quad (33)$$

The value function of the standard exponential agent facing a consumption tax is:

$$\check{v}(x; \tau) = \mathbb{E}_t \left[\int_t^\infty e^{-\rho(s-t)} u((1 - \tau)\check{c}_s) ds \right]. \quad (34)$$

Since $\hat{c}(x) = \check{c}(x)$, the consumption path in equation (33) is identical to the gross consumption expenditure path in equation (34). As above, I need to find the value of τ such that:

$$\hat{u}(c) = u((1 - \tau)c).$$

When $\gamma = 1$ this implies that $-\ln(\beta) + \frac{\beta-1}{\beta} = \ln(1 - \tau)$. Rearranging gives

$$\tau = 1 - \frac{\exp\left(\frac{\beta-1}{\beta}\right)}{\beta}.$$

The Effect of β and β^E . Assume that $1 - \gamma + \gamma\alpha > 0$ so that τ is defined. First, I show that τ is decreasing in α . The derivative

$$\frac{\partial \tau}{\partial \alpha} = \frac{-1}{1 - \gamma} \left(\frac{\alpha^\gamma}{1 - \gamma + \gamma\alpha} \right)^{\frac{\gamma}{1-\gamma}} \left(\frac{\gamma\alpha^{\gamma-1}}{1 - \gamma + \gamma\alpha} - \frac{\gamma\alpha^\gamma}{(1 - \gamma + \gamma\alpha)^2} \right)$$

implies that

$$\begin{aligned} \text{sgn}\left(\frac{\partial \tau}{\partial \alpha}\right) &= \text{sgn}(\gamma - 1) \times \text{sgn}\left(1 - \frac{\alpha}{1 - \gamma + \gamma\alpha}\right), \text{ or equivalently} \\ \text{sgn}\left(\frac{\partial \tau}{\partial \alpha}\right) &= \text{sgn}(\gamma - 1)\text{sgn}(1 - \gamma). \end{aligned}$$

Thus, τ is always decreasing in α .

The derivative of α with respect to β is:

$$\frac{\partial \alpha}{\partial \beta} = \frac{\psi^E}{\gamma\beta^E} \left(\frac{\beta}{\beta^E} \right)^{\frac{1-\gamma}{\gamma}} > 0.$$

As stated in the main text, this implies that $\frac{\partial \tau}{\partial \beta} < 0$.

The derivative of α with respect to β^E is:

$$\frac{\partial \alpha}{\partial \beta^E} = \frac{1}{\gamma} \left(\frac{\beta}{\beta^E} \right)^{\frac{1}{\gamma}} - \frac{1}{\gamma} \left(\frac{\beta}{\beta^E} \right)^{\frac{1}{\gamma}} \frac{\psi^E}{\beta^E}.$$

So, $\frac{\partial \alpha}{\partial \beta^E} > 0$ when $\beta^E > \psi^E$, and $\frac{\partial \alpha}{\partial \beta^E} < 0$ when $\beta^E < \psi^E$. Since $\beta^E > \psi^E$ when $\gamma < 1$ (and vice versa), this implies that α is increasing in β^E when $\gamma < 1$, and decreasing in β^E when $\gamma > 1$. This also implies that $\frac{\partial \tau}{\partial \beta^E} < 0$ when $\gamma < 1$, and $\frac{\partial \tau}{\partial \beta^E} > 0$ when $\gamma > 1$. As stated in the main text, naivete increases the welfare cost of present bias when $\gamma > 1$.

A.8 Proof of Proposition 10

This follows from the proof of Proposition 9, which shows that the realized continuation-value function of the IG agent is a positive affine transformation of the value function for the standard exponential agent. Accordingly, improving the (realized) welfare of the IG agent is equivalent to improving the welfare of the standard exponential agent.

B Numerical Methods Theory

Barles and Souganidis (1991) show that a finite difference scheme converges to the unique viscosity solution of an HJB equation as long as certain conditions hold. However, I show below that these conditions do not necessarily hold when $\beta < 1$. This failure means that one cannot directly solve the Bellman equation of the IG agent. Instead, the key algorithmic insight of this paper is that the following two-step approach can be used to solve for the IG agent's equilibrium. First, solve the HJB equation of the time-consistent \hat{u} agent. Second, compute the IG agent's equilibrium directly from the \hat{u} agent using Proposition 2 and equations (11) and (12) (or equations (13) and (14) if the agent is naive). I apply this algorithm to solve the Aiyagari-Bewley-Huggett model in Appendix E.

Failure of Monotonicity. Here I present a brief description of the problem: the Bellman equation of the IG agent fails to satisfy a monotonicity property. I follow Tourin (2013)'s treatment of Barles and Souganidis (1991). For simplicity, I assume here that income is deterministic with $y_t = y$, and there is just a single liquid asset b .

Let \mathcal{G} denote the discretized grid over liquid wealth b on which $v(b)$ is solved numerically. Assume this grid is uniformly spaced, and let Δ_b denote the size of the grid increment. At each gridpoint $g \in \mathcal{G}$, define:

$$S_g = \rho v_g - u(c_g) - \frac{v_{g+1} - v_g}{\Delta_b}(y + rb_g - c_g)^+ - \frac{v_g - v_{g-1}}{\Delta_b}(y + rb_g - c_g)^-,$$

where v_g, v_{g+1} , and v_{g-1} represent the value function at gridpoints $g, g+1$, and $g-1$, b_g is the wealth level at gridpoint g , and c_g is the consumption choice at gridpoint g .

For monotonicity to hold, S_g must be weakly decreasing in v_g , v_{g+1} , and v_{g-1} . To show that monotonicity fails when $\beta < 1$, assume that $y + rb_g - c_g < 0$. In this case, c_g is defined implicitly by $u'(c_g) = \beta \frac{v_g - v_{g-1}}{\Delta_b}$. Consider an increase in v_{g-1} :

$$\begin{aligned} \frac{\partial S_g}{\partial v_{g-1}} &= -u'(c_g) \frac{\partial c_g}{\partial v_{g-1}} + \frac{1}{\Delta_b} (y + rb_g - c_g)^- + \frac{v_g - v_{g-1}}{\Delta_b} \frac{\partial c_g}{\partial v_{g-1}} \\ &= (1 - \beta) \frac{v_g - v_{g-1}}{\Delta_b} \frac{\partial c_g}{\partial v_{g-1}} + \frac{1}{\Delta_b} (y + rb_g - c_g)^-, \end{aligned}$$

where the property that $u'(c_g) = \beta \frac{v_g - v_{g-1}}{\Delta_b}$ is used to go from the first to the second line.

If $\beta = 1$ then monotonicity holds: $\frac{\partial S_g}{\partial v_{g-1}} < 0$ since the first term drops out and $y + rb_g - c_g < 0$ by assumption.

If $\beta < 1$ then monotonicity may not hold. Since $\frac{\partial c_g}{\partial v_{g-1}} > 0$ the term $(1 - \beta) \frac{v_g - v_{g-1}}{\Delta_b} \frac{\partial c_g}{\partial v_{g-1}} > 0$ whenever $\beta < 1$. Now, it is possible for $\frac{\partial S_g}{\partial v_{g-1}} > 0$, in which case monotonicity does not hold.

The above example points to the difficulty of using finite difference methods to solve *directly* for the equilibrium of the IG agent. Since this difficulty only arises when $\beta < 1$, finite difference methods can still be used to solve for the equilibrium of the \hat{u} agent. Given a solution to the \hat{u} agent, the equilibrium of the IG agent can then be backed out: Proposition 2 implies that $v(x) = \hat{v}(x)$, and equations (11) and (12) define $c(x)$ and $d(x)$. In the case of the one-step extension to naivete, equations (13) and (14) can be used instead.

C Present Bias and Policy: A Simple Example

This Appendix provides a simple example to show how government interventions can improve the equilibrium of an economy with present-biased agents.⁴⁴ I study a simple “Cake-Eating” model of consumption-saving behavior. I assume that the model is deterministic, with income $y_t \equiv \bar{y}$. There is a single liquid asset b with a constant interest rate r , and there is a single representative agent with initial wealth b_0 . The

⁴⁴I thank David Laibson for suggesting this example. A similar result is presented in Laibson (1998).

borrowing limit is set to the natural borrowing constraint of $\underline{b} = \frac{-\bar{y}}{r}$. This agent has sophisticated IG preferences.

In this simple model, the IG agent consumes $c(b) = \frac{\rho-(1-\gamma)r}{\gamma-(1-\beta)}(b + \frac{\bar{y}}{r})$.⁴⁵ However, the first-best consumption level is $\check{c}(b) = \frac{\rho-(1-\gamma)r}{\gamma}(b + \frac{\bar{y}}{r})$.⁴⁶

For simplicity, I assume that $\rho = r$, $\gamma = 1$, and $b_0 > 0$. With these three assumptions, the first-best consumption level is $\check{c}(b_0) = rb_0 + \bar{y}$. In other words, it is optimal for the agent to consume the annuity value of their wealth plus the deterministic income flow.

I now introduce a social planner to improve the consumption-saving decisions of the representative IG agent. The social planner is allowed to use a combination of interest rate subsidies and consumption taxes, subject to a balanced-budget constraint. Interest rate subsidies encourage saving, while consumption taxes are a means of financing these subsidies.

Denote the consumption tax by ϕ_t , and the subsidized interest rate by r_t^s . The social planner runs a balanced budget for all t , so the interest rate subsidy of $(r_t^s - r)b_t$ must equal the total tax revenue collected at each point in time.

With the introduction of consumption taxes, I will now use c_t to denote gross consumption expenditures at time t . However, the agent only gets to consume share $1 - \phi_t$ of gross consumption expenditures, with the rest going to taxes.

I now show that the social planner can recover the first-best equilibrium using a constant consumption tax and interest rate subsidy. To implement the first-best equilibrium, the planner needs to choose r^s and ϕ such that:

$$(1 - \phi) \frac{rb_0 + \frac{r\bar{y}}{r^s}}{\beta} = rb_0 + \bar{y}, \text{ and} \quad (35)$$

$$\phi \frac{rb_0 + \frac{r\bar{y}}{r^s}}{\beta} = (r^s - r)b_0. \quad (36)$$

Under the simple calibration studied here, the IG agent will choose gross consump-

⁴⁵It is well known that a $\beta = 1$ agent would consume $\frac{\rho-(1-\gamma)r}{\gamma}(b + \frac{\bar{y}}{r})$ (see e.g. [Fagereng et al., 2019](#)). Then, Proposition 3 implies that the IG agent consumes $\frac{1}{\psi}$ times what the $\beta = 1$ agent consumes.

⁴⁶As discussed in Section 6, IG preferences feature a single welfare criterion even though they are time-inconsistent. Hence, the concept of “first best” is well defined under IG preferences.

tion expenditures of $c(b) = \frac{rb + \frac{r\bar{y}}{r^s}}{\beta}$.⁴⁷ However, actual consumption is only $(1 - \phi)c(b)$. Equation (35) imposes that realized consumption is at its first-best level: $(1 - \phi)c(b_0) = rb_0 + \bar{y}$. Equation (36) is the balanced-budget condition. It says that tax revenues of $\phi c(b_0)$ must equal the interest rate subsidy of $(r^s - r)b_0$.

One can show that the following set of policy tools produces the first-best equilibrium:

$$r^s = \frac{r}{\beta}, \quad (37)$$

$$\phi = \frac{rb_0(1 - \beta)}{rb_0 + \beta\bar{y}}. \quad (38)$$

For example, consider the calibration $\beta = 0.75$, $r = 3\%$, $\bar{y} = 1$, and $b_0 = 3$ (similar to Appendix E). The optimal consumption tax is $\phi = 2.68\%$, and the optimal subsidized interest rate is $r^s = 4\%$.

Welfare and Implementability when $\beta = 1$. Proposition 10 highlights the channels through which present bias matters for policymakers: present bias does not matter for determining whether a policy is welfare-improving, but does matter for determining whether a policy is feasible. This toy model can be used to formalize this discussion.

Proposition 10 implies that the interest rate subsidy plus consumption tax policy in equations (37) and (38) would also be welfare-improving for $\beta = 1$ agents. However, this policy is not possible in an economy populated by a representative $\beta = 1$ agent. At time 0, the $\beta = 1$ agent would consume only $r(b_0 + \frac{\bar{y}}{r^s})$, which is too little to generate the requisite taxes needed to support the interest rate subsidy of $(r^s - r)b_0$.

D Extensions to the Household Balance Sheet Model

This section generalizes the modeling of the illiquid asset along various dimensions, as discussed in Section 7.1. For notational simplicity I focus on the full sophistication

⁴⁷Here I use the property that a constant consumption tax does not change the gross consumption expenditure of the IG agent. See the proof of Proposition 9 for details.

case. Naivete can be captured with the “One-Step Extension” discussed in Section 3.

D.1 Illiquid Assets with Fixed Adjustment Costs

This section studies asset illiquidity due to non-convex adjustment costs. In particular, I now assume that an adjustment to the stock of illiquid assets requires a fixed cost of $\mathcal{F} > 0$.⁴⁸ Because adjustments require fixed costs, the agent will adjust their asset allocation infrequently.

When the agent does not adjust their illiquid assets, the dynamic budget constraint is:

$$\begin{aligned} db_t &= (y_t + r(b_t)b_t - c_t)dt, \\ \frac{da_t}{a_t} &= r^a dt + \sigma^a dZ_t. \end{aligned}$$

When the agent adjusts their illiquid assets, the budget constraint is:

$$b' + a' = b_t + a_t - \mathcal{F}, \tag{39}$$

where b' and a' denote liquid and illiquid wealth immediately after adjustment. At all times, assets remain subject to the constraints that $b_t \geq \underline{b}$ and $a_t \geq 0$.

Equilibrium Under Sophistication. To begin, denote by w^* the shadow current-value that the agent would earn if they adjusted their illiquid assets. That is, if an agent at point $x = (b, a, y)$ would jump to point $x' = (b', a', y)$ conditional on adjusting, then $w^*(x) = w(x')$. Similarly, let $v^*(x) = v(x')$.

A stationary Markov-perfect equilibrium to the sophisticated IG agent’s intrapersonal problem is characterized by the following Bellman equation, which consists of

⁴⁸Adjustment costs can be dependent on the state space, and can include variable costs in addition to fixed costs. For simplicity, these alternate modeling choices are not explicitly studied here.

a differential variational inequality defined on x :

$$\begin{aligned} \rho v(x) = \max \bigg\{ & \rho v^*(x), \quad u(c(x)) + v_b(x) (y + r(b)b - c(x)) \\ & + v_a(x) (r^a a) + \frac{1}{2} v_{aa}(x) (a\sigma^a)^2 \\ & + \sum_{y' \neq y} \lambda^{y \rightarrow y'} (v(b, a, y') - v(b, a, y)) \bigg\}, \end{aligned} \quad (40)$$

subject to the optimality conditions:

$$u'(c(x)) = \begin{cases} \beta v_b(x) & \text{if } b > \underline{b} \\ \max \left\{ \beta v_b(x), u'(y + r(\underline{b})\underline{b}) \right\} & \text{if } b = \underline{b} \end{cases}, \quad \text{and} \quad (41)$$

$$v^*(x) = \max_{b', a'} v(x') \quad \text{s.t. constraint (39) holds.} \quad (42)$$

Equation (40) is similar to equation (10), except that equation (40) is written as a variational inequality in order to capture the fact that the agent always has the option of adjusting their illiquid assets.⁴⁹ Intuitively, when the agent does not adjust their asset allocation then the value function is pinned down by the righthand branch of equation (40), which is the standard Bellman equation for the IG agent (see equation (10)). When the agent does pay the fixed cost to adjust their asset allocation, the lefthand branch is selected and the value function is $v^*(x)$.

The consumption decision in equation (41) is identical to equation (11) in the main text.

Finally, equation (42) defines the asset allocation decision conditional on adjustment. Whenever the agent decides to adjust their illiquid asset holdings they choose b' and a' to maximize $w^*(x)$. Since $w(x) = \beta v(x)$, maximizing current-value function w is equivalent to maximizing continuation-value function v . Accordingly, equation (42) works directly with v .

⁴⁹A similar approach is used in Laibson et al. (2021). See also the notes at <https://benjaminmoll.com/codes/> under the heading “Stopping Time Problem” for additional details and mathematical references.

Results. Under this alternate setup for asset illiquidity the results in the main text still hold. In particular, the \hat{u} construction remains the same. Thus, the IG agent’s value function $v(x)$ is unique and equivalent to the \hat{u} agent’s value function (Propositions 1 and 2). When Assumption 1 holds, the IG agent’s consumption function equals $\left(\frac{\beta^E}{\beta}\right)^{\frac{1}{\gamma}} \frac{1}{\psi^E}$ times that of the standard exponential agent (Proposition 3), and the IG agent has the same demand for illiquid assets as the standard exponential agent (similar to Proposition 8).⁵⁰ The welfare results in Section 6 (Propositions 9 and 10) also continue to hold.

D.2 Illiquid Housing

This section generalizes the modeling of the illiquid asset to the case of an illiquid durable. In particular, I will refer to the durable as housing, which is a natural application since housing is an important illiquid asset for many U.S. households.

Let a_t denote the agent’s holding of illiquid housing. Unlike strictly financial assets, housing provides the agent with a flow of housing services, fa_t . I continue to assume that the durable has a return of r^a and volatility σ^a . I continue to let d_t denote adjustments to the durable stock, which are still subject to transaction cost $\chi(d_t, a_t)$. The dynamic budget constraint continues to satisfy equations (3) and (4).

In this extension with housing, the agent’s utility is given by a Cobb-Douglas aggregator over consumption c_t and flow housing services.⁵¹ I assume that the agent’s total flow of housing services is composed of the flow from the agent’s illiquid housing stock, fa_t , plus a small fixed flow of \underline{h} . Technically, this minimum housing flow \underline{h} ensures that the agent always consumes some amount of housing. It is a reduced-form alternative to a richer model in which the agent can switch between being a

⁵⁰In this model with discrete adjustments, “demand for illiquid assets” refers to two nested decisions. First, the agent chooses where in the state space to adjust their illiquid assets. Second, the agent chooses b', a' conditional on adjusting. Both of these decisions are independent of β (Laibson et al., 2021).

⁵¹Cobb-Douglas preferences are particularly relevant for the case of housing, since empirically the housing expenditure share has been stable over time. The case of separable utility is also easily handled.

homeowner and a renter.⁵² Fully, the agent's utility function is now:

$$u(c, a) = \begin{cases} \frac{(c^\eta(fa + \underline{h})^{1-\eta})^{1-\gamma} - 1}{1-\gamma} & \text{if } \gamma \neq 1 \\ \ln(c^\eta(fa + \underline{h})^{1-\eta}) & \text{if } \gamma = 1 \end{cases}. \quad (43)$$

Remark. *The model in the main text is recovered when $\eta = 1$. Additionally, throughout this section I impose that Assumption 1 holds (and I ignore boundary conditions accordingly).*

In this extension with durables, the sophisticated IG agent's intrapersonal problem is characterized by the following Bellman equation:

$$\begin{aligned} \rho v(x) = & u(c(x), a) + v_b(x) (y + r(b)b - d(x) - \chi(d(x), a) - c(x)) \\ & + v_a(x) (r^a a + d(x)) + \frac{1}{2} v_{aa}(x) (a\sigma^a)^2 \\ & + \sum_{y' \neq y} \lambda^{y \rightarrow y'} (v(b, a, y') - v(b, a, y)), \end{aligned} \quad (44)$$

subject to the optimality conditions:

$$u_c(c(x), a) = \begin{cases} \beta v_b(x) & \text{if } b > \underline{b} \\ \max \left\{ \beta v_b(x), u_c(y + r(\underline{b})\underline{b}, a) \right\} & \text{if } b = \underline{b} \end{cases}, \quad \text{and} \quad (45)$$

$$\chi_d(d(x), a) = \begin{cases} \frac{v_a(x)}{v_b(x)} - 1 & \text{if } b > \underline{b} \\ 0 & \text{if } b = \underline{b} \end{cases}. \quad (46)$$

Equations (44) – (46) are similar to equations (10) – (12) in the main text.

Construction of \hat{u}^{CD} Agent. To study this model with housing, I again want to reverse-engineer an agent with standard exponential time preferences ($\beta = 1$) but a penalized utility function \hat{u}^{CD} such that the value function of the \hat{u}^{CD} agent, denoted \hat{v}^{CD} , equals the value function v of the IG agent. I use the “CD” superscript

⁵²In continuous time, this rental versus ownership choice could be modeled using the variational inequality structure introduced in Appendix D.1. I abstract from these details here in order to simplify the presentation.

to emphasize that the Cobb-Douglas aggregator used here changes the construction relative to the main text.

As in the proof of Proposition 2, I construct the \hat{u}^{CD} utility function so that the following condition holds for all $b > \underline{b}$:

$$u(c(x), a) - v_b(x)c(x) = \hat{u}^{CD}(\hat{c}^{CD}(x), a) - \hat{v}_b^{CD}(x)\hat{c}^{CD}(x). \quad (47)$$

Condition (47) ensures that $v(x) = \hat{v}^{CD}(x)$ whenever Assumption 1 holds.

One can show that the requisite utility function \hat{u}^{CD} is as follows:

$$\begin{aligned} \hat{u}^{CD}(\hat{c}^{CD}, a) &= \frac{\psi^{CD}}{\beta} u\left(\frac{1}{\psi^{CD}} \hat{c}^{CD}, a\right) + \frac{\psi^{CD} - 1}{\beta}, \text{ where} \\ \psi^{CD} &= \frac{\gamma^{CD} - (1 - \beta)}{\gamma^{CD}} \text{ and } \gamma^{CD} = 1 - \eta(1 - \gamma). \end{aligned} \quad (48)$$

This is just like the \hat{u}^+ utility function in equation (15), except that ψ is replaced by ψ^{CD} and risk aversion parameter γ is replaced by $\gamma^{CD} = 1 - \eta(1 - \gamma)$. I write the equations in this way to highlight the following intuition: in any specific instant, housing stock a_t is fixed (because housing is illiquid). Accordingly, for that instant the utility function $u(c, a) = \frac{(c^\eta(fa + \underline{h})^{1-\eta})^{1-\gamma} - 1}{1-\gamma}$ is a positive affine transformation of $\frac{(c^\eta)^{1-\gamma} - 1}{\eta(1-\gamma)}$, or equivalently $\frac{c^{1-(1-\eta(1-\gamma))} - 1}{1-(1-\eta(1-\gamma))}$, which is standard CRRA utility but with risk aversion $\gamma^{CD} = 1 - \eta(1 - \gamma)$. Thus, with Cobb-Douglas preferences over housing we can use the same construction of \hat{u} , except that γ is replaced by $\gamma^{CD} = 1 - \eta(1 - \gamma)$.

Results. In this model with illiquid housing, a perturbed version of the results in the main text now obtains. In particular, the results differ with housing because one needs to use effective risk aversion $\gamma^{CD} = 1 - \eta(1 - \gamma)$ instead of γ . I emphasize that the irrelevance result in Proposition 8 continues to hold in this model with housing, suggesting that this irrelevance result does not rely on illiquid assets being strictly financial assets.

****SUPPLEMENTARY MATERIAL****

E The Aiyagari-Bewley-Huggett Model

This appendix studies IG preferences in a workhorse “Aiyagari-Bewley-Huggett” heterogeneous-agent model, following the continuous-time specification of [Achdou et al. \(2022\)](#). This model serves as an important building-block for a wide range of quantitative applications.⁵³

I model an endowment economy in which a continuum of agents have heterogeneous income and wealth profiles. Consumers are able to self-insure against income fluctuations by accumulating a buffer-stock of savings. At the aggregate level there exists an exogenous supply B of bonds. Interest rate r is determined in general equilibrium to equate the supply of savings with B .

With present bias, the solution to this model takes the form of a nested equilibrium. There is a sequence of two equilibria that must be solved jointly: (i) the intrapersonal equilibrium of the IG agent, taking prices as given; and (ii) the general equilibrium, in which individual-level policy functions are aggregated and markets clear. As explained below, in continuous time the joint solution to these equilibria takes the form of two coupled PDEs.

E.1 Consumer Problem (Intrapersonal Equilibrium)

The consumer’s side of this model is a simplification of the more general setup in Section 3. I discuss specifics where necessary, and refer the reader to the main text for details.

⁵³Foundational work includes [Bewley \(1986\)](#), [Huggett \(1993\)](#), and [Aiyagari \(1994\)](#), as well as [Imrohoroglu \(1989\)](#), [Zeldes \(1989\)](#), [Deaton \(1991\)](#), [Carroll \(1997\)](#), and [Gourinchas and Parker \(2002\)](#). For more recent surveys, see e.g. [Heathcote et al. \(2009\)](#), [Krueger et al. \(2016\)](#), [Benhabib and Bisin \(2018\)](#), and [Kaplan and Violante \(2018\)](#).

The Household Balance Sheet. There is a single liquid asset. Let b_t denote an agent's wealth at time t . b_t evolves as follows:

$$db_t = (y_t + rb_t - c_t)dt. \quad (49)$$

y_t is a stochastic endowment income process and c_t is consumption. For simplicity I assume that y_t follows a two-state Poisson process $y_t \in \{y_1, y_2\}$, with $0 < y_1 < y_2$. The income process jumps from state y_1 to y_2 with intensity $\lambda^{1 \rightarrow 2}$, and from y_2 to y_1 with intensity $\lambda^{2 \rightarrow 1}$.

Wealth is subject to the borrowing limit $b_t \geq \underline{b}$. With present bias the equilibrium is particularly sensitive to whether or not this borrowing constraint binds in equilibrium, because binding constraints form a commitment device of sorts by limiting overconsumption at \underline{b} (see the discussion in Section 5.1 for more).

Utility and Value. As in the main text, agents have CRRA utility over consumption. I assume here that agents are fully sophisticated about their present bias. For comparison, the case of full naivete is presented in Appendix E.5.

Intrapersonal Equilibrium. The intrapersonal equilibrium to this model is a simplification of the model in the main text. Here, let $x = (b, y)$, where $b \in [\underline{b}, \infty)$ and $y \in \{y_1, y_2\}$. A stationary Markov-perfect equilibrium to the sophisticated IG agent's intrapersonal problem is characterized by the following Bellman equation:

$$\rho v(x) = u(c(x)) + v_b(x) (y + rb - c(x)) + \lambda^{y \rightarrow y'} (v(b, y') - v(b, y)), \quad (50)$$

subject to the optimality condition:

$$u'(c(x)) = \begin{cases} \beta v_b(x) & \text{if } b > \underline{b} \\ \max \left\{ \beta v_b(x), u'(y + r\underline{b}) \right\} & \text{if } b = \underline{b} \end{cases}. \quad (51)$$

E.2 General Equilibrium

To solve for a general equilibrium to this heterogeneous-agent model, the intrapersonal equilibrium of IG agents is aggregated and bond market clearing is imposed. To close the model as simply as possible, I assume that there is an exogenous supply of safe debt $B \in (\underline{b}, \infty)$ that agents can hold (Huggett, 1993). It is well known that the model can be closed in alternate ways (e.g., Aiyagari, 1994). This paper focuses on the demand side of the economy, where present-biased preferences interact with incomplete markets. Simplicity is preferred on the supply side for expositional clarity.

Let $g_t(b, y)$ denote the distribution of wealth and income at time t , such that $\int_{\underline{b}}^{\infty} g_t(b, y_1)db + \int_{\underline{b}}^{\infty} g_t(b, y_2)db = 1$. Since this is an endowment economy with exogenous income, the one price that must be pinned down in general equilibrium is the interest rate r . Given r , the consumer's intrapersonal equilibrium is described by equations (50) and (51). The resulting policy functions give rise to a Kolmogorov Forward (KF) equation that characterizes the evolution of the aggregate wealth distribution.

In a stationary equilibrium the distribution of wealth is constant:

$$0 = -\frac{\partial}{\partial b}[s(b, y)g(b, y)] - \lambda^{y \rightarrow y'}g(b, y) + \lambda^{y' \rightarrow y}g(b, y'), \quad (52)$$

where $s(b, y)$ is the saving policy function $s(b, y) = y + rb - c(b, y)$. Equations (50) and (51), plus KF equation (52), define a steady state aggregate savings function:

$$S(r) = \int_{\underline{b}}^{\infty} bg(b, y_1)db + \int_{\underline{b}}^{\infty} bg(b, y_2)db. \quad (53)$$

The bond market clears when $S(r) = B$.

The action in this model occurs on the demand side of the economy, where consumers are heterogeneous. Present bias adds an additional layer of complexity by making the individual's problem itself a dynamic game. The benefit of the continuous-time IG approach is that the intrapersonal equilibrium can be characterized by a partial differential equation (equation (50)). Following Achdou et al. (2022), a general equilibrium can then be found by coupling the KF equation in (52) with the

IG agent’s intrapersonal equilibrium. Using this approach, this paper is among the first to solve a general equilibrium incomplete markets model where consumers have present bias and policy functions are nonlinear.⁵⁴

E.3 Model Solution and Results

I now solve this workhorse model numerically, and compare the equilibrium of an economy with IG consumers ($\beta < 1$) to an economy with exponential consumers ($\beta = 1$). Afterwards, I provide an additional set of theoretical results that formalize equilibrium properties of this workhorse heterogeneous-agent model.

Stylized Calibration. I roughly calibrate the economy to reflect the problem of a typical American household. Average income is normalized to one. I set $y_1 = 0.74$ and $y_2 = 1.26$, with a job switching rate of $\lambda^{1 \rightarrow 2} = \lambda^{2 \rightarrow 1} = 0.19$. This calibration is a two-state discretization of the income process used in [Guerrieri and Lorenzoni \(2017\)](#).⁵⁵

Borrowing constraint $\underline{b} = -\frac{1}{3}$, which corresponds to the average credit limit reported in the 2016 Survey of Consumer Finances ([Laibson et al., 2021](#)). I set the coefficient of relative risk aversion $\gamma = 2$. The exogenous supply of bonds is calibrated to $B = 3$ in order to roughly capture the ratio of wealth to income in the United States ([Kaplan et al., 2018](#)).

I target a steady-state interest rate of 3%, and the discount function is calibrated internally to produce this interest rate in equilibrium. In the exponential model with $\beta = 1$, $r = 3\%$ is produced with $\rho = 3.5\%$. In the IG calibration I set $\beta = 0.75$. The calibration of $\beta = 0.75$ is a conservative choice in the consumption-saving literature,

⁵⁴[Maliar and Maliar \(2006\)](#) solve a similar model in discrete time but are forced to make smoothness assumptions, which are only valid for β near 1, to solve the model. Such assumptions are not needed in continuous time with IG preferences.

⁵⁵[Guerrieri and Lorenzoni \(2017\)](#) assume that log-income follows an AR(1) at a quarterly frequency: $\log(y_{t+1}) = \rho \left(\log(y_t) - \frac{\sigma^2}{2} \right) + \sigma \varepsilon_{t+1}$. Using the estimates of [Floden and Lindé \(2001\)](#), this process is calibrated with persistence $\rho = 0.967$ and variance $\sigma^2 = 0.017$. As in [Laibson et al. \(2021\)](#), I convert this quarterly AR(1) into an Ornstein-Uhlenbeck process, and then discretize the Ornstein-Uhlenbeck process into two states using finite-difference methods. I set the income states to ± 1 standard deviation.

and the results that follow become more stark as β decreases.⁵⁶ Given $\beta = 0.75$, $\rho = 2.5\%$ produces a 3% steady-state interest rate.

To solve this model numerically I build on the finite difference methods presented in Achdou et al. (2022). Details are given in Appendix B.

Consumption and Saving. The top panel of Figure 2 plots the consumption function for the $\beta = 0.75$ calibration and the $\beta = 1$ calibration. The $\beta = 1$ consumption function is standard (see Achdou et al. (2022) for details). For $\beta = 0.75$, consumption is well-behaved in that it does not suffer from pathologies on the interior of the wealth space (Laibson and Maxted, 2022). However, the consumption function features a discontinuity at $b = \underline{b}$ when $y = y_1$. This consumption discontinuity is produced by the corresponding discontinuity in the IG agent’s discount function. Consider the self in control an instant before the constraint binds. This self does not want to smooth consumption with the next self (for whom the constraint will bind), since the self in control discounts the utility of the next self by β .

The bottom panel of Figure 2 plots the corresponding saving function $s(x) = y + rb - c(x)$ for both calibrations. Near \underline{b} the IG agent has a lower saving rate than the exponential agent. This pattern reverses as b increases. In short, the IG agent saves less when poor but saves more when wealthy. Relative to the $\beta = 1$ agent, the IG agent has both an incentive for lower saving ($\beta < 1$) and an incentive for higher saving (lower value of ρ). When the consumption function is nonlinear, the relative impact of β versus ρ varies over the state space. Near \underline{b} , $\beta < 1$ dominates and the IG agent saves less than the exponential agent. Away from \underline{b} , the low value of ρ dominates and the IG agent saves more than the exponential agent.

There are two reasons why the relative effect of $\beta < 1$ matters more for low levels of wealth. First, as presented above in Euler equation (19), present bias creates a disagreement between successive selves that is increasing in the slope of the consumption function (i.e., the instantaneous MPC). Since the consumption function is concave the slope of the consumption function is highest near \underline{b} , which decreases the

⁵⁶For example, Angeletos et al. (2001) set $\beta = 0.7$, arguing that this is consistent with laboratory experiments. Laibson et al. (2023) estimate $\beta = 0.5$ in a structural lifecycle model. Allcott et al. (2022) estimate $\beta = 0.75$ on a sample of payday loan users.

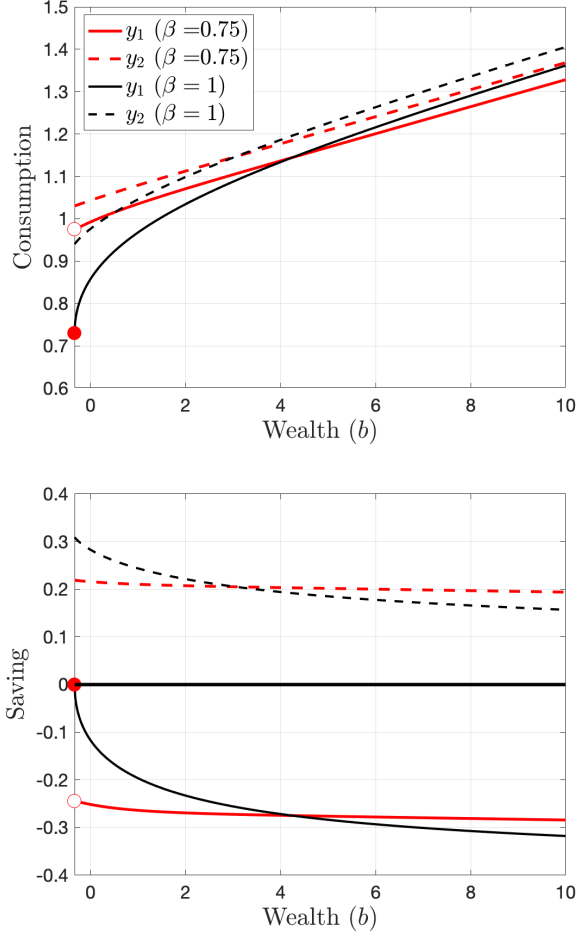


Figure 2: **Equilibrium Consumption-Saving Decisions.** This figure plots the consumption function (top) and the saving function (bottom) for the $\beta = 0.75$ calibration and the $\beta = 1$ calibration.

saving rate near \underline{b} . Second, in the case where the borrowing constraint binds, time inconsistency interacts with the consumer's effective planning horizon to lower the saving rate near \underline{b} .

The intuition for the second effect – which requires binding borrowing constraints as modeled here – is as follows. Sophisticated present bias means that the current self distrusts the consumption decisions of future selves. This has offsetting effects on the current self's incentive to save. On the one hand, the current self knows that wealth will be spent imprudently in the future. This decreases the saving rate of self t . On the other hand, because self t knows that future selves will not save enough, self t has an

incentive to set aside wealth today in order to buffer against future overconsumption. This increases the saving rate of self t . The relative strength of the second effect depends on the effective planning horizon. Binding borrowing constraints shorten the consumer's effective planning horizon and limit self t 's ability to pass wealth far into the future (because any marginal savings will be fully consumed in finite time), which reduces the current self's incentive to save. Proposition 13 below characterizes the interaction of present bias with the hard borrowing constraint at \underline{b} .

Distributions. Figure 3 plots the stationary distribution of wealth in the two calibrations. By bond market clearing (with $B = 3$), the average wealth level is constant across the two calibrations. However, the underlying distribution of wealth differs considerably. Relative to $\beta = 1$, the $\beta = 0.75$ calibration features both a larger share of agents near \underline{b} and a thicker right tail. This is consistent with the saving functions shown in Figure 2. Near \underline{b} the IG agent has difficulty generating savings. However, the lower long-run discount rate of the IG agent means that the IG agent also adopts a higher saving rate as b increases. This generates the thicker right tail observed in Figure 3.

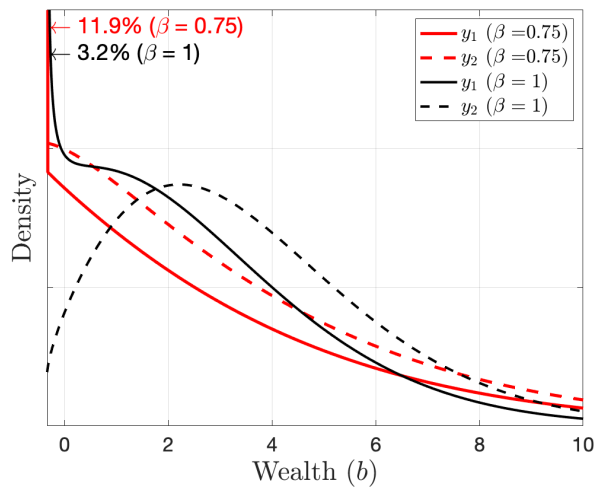


Figure 3: **The Distribution of Wealth.** This figure shows the stationary wealth distribution for the $\beta = 0.75$ calibration and the $\beta = 1$ calibration.

Quantifying these differences, Table 1 compares various measures of wealth inequality across the two calibrations. Wealth inequality is much higher for the $\beta = 0.75$

economy. The maximum wealth level attained in the $\beta = 0.75$ economy is twice as large as the maximum wealth level attained in the $\beta = 1$ economy. The $\beta = 0.75$ economy also features more wealth in the top 0.1%, 1%, 5%, and 10%. However, the $\beta = 0.75$ economy produces almost four times as many agents constrained at \underline{b} . As shown in Panel C of the table, all of these differences in wealth inequality arise even under the restriction that aggregate wealth is constant across the two calibrations.

	$\beta = 1$	$\beta = 0.75$
Panel A: Upper Wealth Moments		
Maximum Wealth (b_{max})	71.2	154.4
Average Wealth Top 0.1 %	14.7	21.9
Average Wealth Top 1 %	11.7	16.7
Average Wealth Top 5 %	9.3	12.5
Average Wealth Top 10 %	8.1	10.6
Panel B: Lower Wealth Moments		
Share $b < 0$	8.0%	18.0%
Share $b = \underline{b}$	3.2%	11.9%
Average Wealth Bottom 50 %	1.1	0.5
Panel C: Aggregate Wealth Moments		
Standard Deviation b	3.9	4.5
Mean b	3.0	3.0
Median b	2.6	2.0

Table 1: **Wealth Moments.** This table characterizes the wealth distribution for the $\beta = 0.75$ calibration and the $\beta = 1$ calibration.

As discussed in Section 5.2, one way to contextualize the results in Table 1 is to note that present bias generates similar effects on the wealth distribution as heterogeneous time preferences.⁵⁷ However, rather than assuming preference heterogeneity across individuals, present bias endogenously generates effective time-preference heterogeneity within individuals that varies with wealth (recall the Euler equation in Proposition 6). Near \underline{b} the IG agent acts impatiently. This produces the large mass of constrained agents. IG preferences simultaneously generate a longer right tail in the wealth distribution, since IG agents act more patiently as their wealth increases. It is well known that heterogeneous-agent models of the type solved here produce

⁵⁷Heterogenous time preferences are a tool that has frequently been employed in macroeconomic models to generate realistic wealth distributions (e.g., [Campbell and Mankiw, 1989](#); [Krusell and Smith, 1998](#)).

counterfactually low levels of wealth inequality when $\beta = 1$ (Carroll, 1997; Quadrini et al., 1997). While this numerical example is far too stylized to make any quantitative claims, present bias moves these models in the right direction.⁵⁸

The Marginal Propensity to Consume (MPC). Uninsurable income shocks generate state-dependent MPCs. To study consumers' MPCs in this model, I follow Achdou et al. (2022) who define the MPC as follows:

Definition (Achdou et al. (2022) Definition 1). *The Marginal Propensity to Consume over a period τ is given by*

$$\begin{aligned} MPC_\tau(x) &= C'_\tau(x), \text{ where} \\ C_\tau(x) &= \mathbb{E}_t \left[\int_t^{t+\tau} c(x_s) ds \mid x_t = x \right]. \end{aligned} \tag{54}$$

In equation (54) the MPC is defined over a discrete unit of time τ . While one could also study the instantaneous MPC of $c_b(x)$, the cumulative MPC is more empirically relevant because consumption is typically observed at a quarterly and/or annual horizon. The MPC can be computed numerically using the Feynman-Kac formula.

Figure 4 plots the quarterly MPC in the two calibrations. Near \underline{b} the MPC is larger for the IG agent because the IG agent lacks the self-control to smooth consumption into the borrowing limit. Away from \underline{b} the MPC is smaller for the IG agent.

The MPC analysis in Figure 4 does not account for the distribution of agents across the high- and low-MPC parts of the state space. In particular, there is more mass near \underline{b} in the $\beta = 0.75$ calibration than in the $\beta = 1$ calibration. In the $\beta = 0.75$ calibration, the average MPC is 24.6% for low-income agents, and 0.9% for high-income agents. In the $\beta = 1$ calibration, the average MPC is 4.5% for low-income agents, and 1.2% for high-income agents. Though this example is stylized, the stark difference in MPCs for low- versus high-income agents when $\beta = 0.75$ suggests that

⁵⁸Mian et al. (2020) document that wealthy households have higher saving rates than poor households, and Fagereng et al. (2019) find that (net) saving rates are approximately constant in wealth. Though the IG model cannot replicate these facts, the IG model performs better than the exponential model because the saving rate declines more slowly with wealth for the IG calibration.

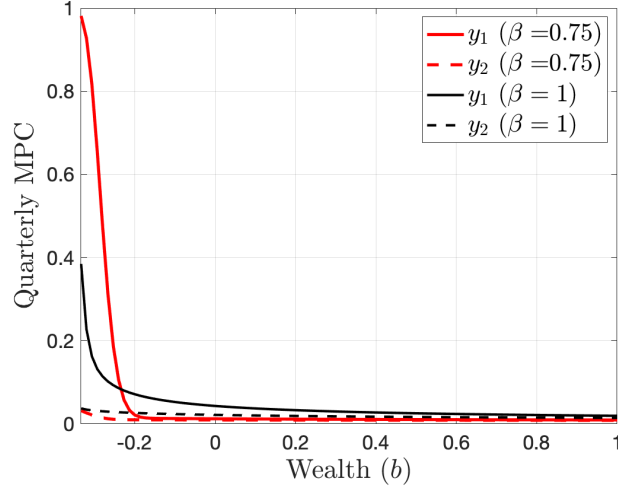


Figure 4: **MPCs.** This figure plots the quarterly MPC for the $\beta = 0.75$ calibration and the $\beta = 1$ calibration.

fiscal stimulus targeted to low-income households will be particularly effective when consumers are present biased. A more complete analysis is presented in [Laibson et al. \(2021\)](#). See also [Laibson et al. \(2022\)](#) for a discussion of how to map these model-based notional MPCs into MPXs.

E.4 Theoretical Properties

I now characterize some additional properties of the model's equilibrium. While the results in this section are specific to the economic environment modeled above, they can typically be generalized. Proofs are provided in [Appendix E.6](#).

Consumption Behavior at the Constraint. I begin by characterizing how present bias creates a consumption discontinuity at \underline{b} when the constraint binds.

Proposition 13. *Let $c(\underline{b}+, y) = \lim_{b \rightarrow \underline{b}+} c(b, y)$. If $\beta < 1$ and \underline{b} binds for income state y_j then there is a discontinuity in consumption at \underline{b} , such that $c(\underline{b}+, y_j) > c(\underline{b}, y_j)$. Specifically, $c(\underline{b}, y_j) = y_j + r\underline{b}$ while $c(\underline{b}+, y_j)$ is defined implicitly by*

$$u'(c(\underline{b}+, y_j)) = \beta \frac{u(c(\underline{b}+, y_j)) - u(y_j + r\underline{b})}{c(\underline{b}+, y_j) - (y_j + r\underline{b})}. \quad (55)$$

Proposition 13 formalizes the discontinuity in $c(b, y_1)$ seen in Figure 2. The IG agent does not smooth consumption into the constraint. Instead, the IG agent chooses a high level of consumption until the instant that the constraint binds, at which point consumption drops discretely. The following perturbation argument, first provided in Harris and Laibson (2004), provides the intuition for equation (55). Consider the self who lives one instant before the borrowing constraint binds. That self can cut consumption by db at a utility cost of $u'(c(\underline{b}+, y_j))db$. This allows future selves to consume at rate $c(\underline{b}+, y_j)$ rather than $y_j + r\underline{b}$ for a timespan of $dt = \frac{db}{c(\underline{b}+, y_j) - (y_j + r\underline{b})}$. The current self values this additional consumption at $\beta (u(c(\underline{b}+, y_j)) - u(y_j + r\underline{b})) dt$. The current self must be indifferent to this perturbation in equilibrium, implying $u'(c(\underline{b}+, y_j))db = \beta (u(c(\underline{b}+, y_j)) - u(y_j + r\underline{b})) dt$. Plugging in for dt yields equation (55).

Consumption Behavior of the Wealthy. The sole source of uncertainty in this model is income risk. Because income risk does not scale with wealth, income risk ceases to affect consumption-saving decisions as $b \rightarrow \infty$. Under exponential discounting, consumption and saving are asymptotically linear in b (Achdou et al., 2022). This linearity property also applies to IG agents.

Proposition 14. *Consumption and saving are asymptotically linear in b . Specifically,*

$$\lim_{b \rightarrow \infty} c(x) = \frac{\rho - (1 - \gamma)r}{\gamma - (1 - \beta)} b, \quad (56)$$

$$\lim_{b \rightarrow \infty} s(x) = \frac{r\beta - \rho}{\gamma - (1 - \beta)} b. \quad (57)$$

An interesting corollary is that when $\beta < 1$ the elasticity of intertemporal substitution (EIS) is no longer given by $\frac{1}{\gamma}$.

Corollary 15. *In the limit as $b \rightarrow \infty$, the EIS is given by:*

$$EIS = \lim_{b \rightarrow \infty} \frac{d \left[\frac{\mathbb{E}_t dc_t / dt}{c_t} \right]}{dr} = \frac{\beta}{\gamma - (1 - \beta)}.$$

A discrete-time version of this result is given in Laibson (1998). With IG pref-

erences, the EIS is less than $\frac{1}{\gamma}$ when $\gamma > 1$, and the EIS is greater than $\frac{1}{\gamma}$ when $\gamma < 1$. The intuition for this result is similar to the Euler equation in Proposition 6. The Euler equation shows that the IG agent chooses their consumption for strategic reasons in addition to standard consumption-smoothing considerations. The EIS determines the sensitivity of the IG agent's current consumption to these strategic motives. When $\gamma < 1$ the IG agent responds more than the standard exponential agent to interest rate changes. The reverse holds for $\gamma > 1$.

Proposition 14 also allows for an approximation of the maximum wealth level that can be attained in this model (see Table 1). Note that this result is an approximation because, for finite wealth, the extent to which consumption can be approximated by a linear function is calibration-dependent.

Remark 16. *The maximum level of wealth is approximately*

$$b_{max} \approx \frac{\kappa(\bar{y}/r) - y_2}{r - \kappa},$$

where $\kappa = \frac{\rho - (1-\gamma)r}{\gamma - (1-\beta)}$ is the consumption rate in equation (56) and $\bar{y} = \frac{\lambda^{2 \rightarrow 1}y_1 + \lambda^{1 \rightarrow 2}y_2}{\lambda^{1 \rightarrow 2} + \lambda^{2 \rightarrow 1}}$ denotes average income.

The intuition for Remark 16 is straightforward. Using equation (56), a wealthy, high-income, agent will save approximately $s(b, y_2) \approx y_2 + rb - \kappa(b + \frac{\bar{y}}{r})$. Setting this equal to 0 and rearranging yields the desired result.

E.5 Model Solution with Naivete

This section replicates the numerical example in Section E.3 under the assumption of complete naivete. To generate an equilibrium interest rate of 3%, I set $\beta = 0.75$, $\beta^E = 1$, and $\rho = 2.45\%$. The calibration is otherwise identical to Section E.3.

Overall the results are qualitatively similar. The main difference in behavior between naifs and sophisticates occurs near \underline{b} . Though naifs still overconsume near the borrowing constraint, Figure 5 illustrates that naifs overconsume by less than sophisticates. As described above, sophistication generates an interaction between present bias and the borrowing constraint, which increases consumption near \underline{b} . This

effect does not arise under naivete.

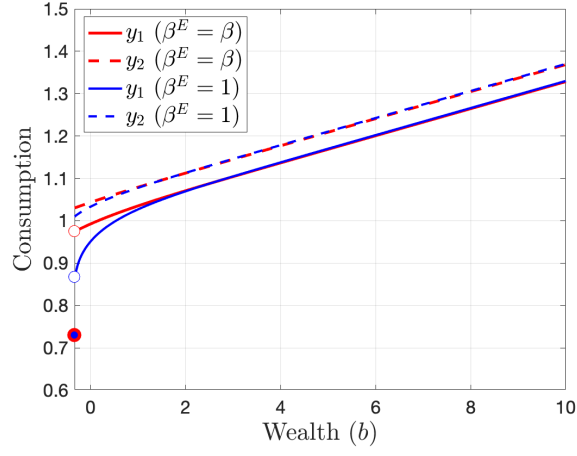


Figure 5: **Equilibrium Consumption-Saving Decisions.** This figure plots the equilibrium consumption function for the $\beta^E = \beta$ calibration (sophistication) and the $\beta^E = 1$ calibration (naivete).

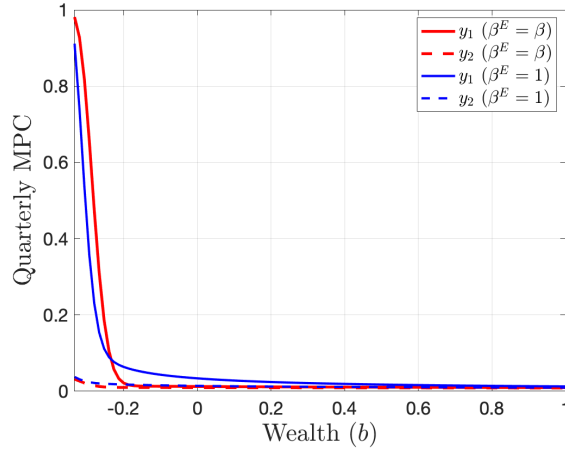


Figure 6: **MPCs.** This figure plots the quarterly MPC for the $\beta^E = \beta$ calibration and the $\beta^E = 1$ calibration.

E.6 Additional Proofs

Proof of Proposition 13. For full details, see Theorem 21 of [Harris and Laibson \(2004\)](#). The value function for the IG agent is given by (see equation (50)):

$$\rho v(x) = u(c(x)) + v_b(x)(y + rb - c(x)) + \lambda^{j \rightarrow i}(v(b, y_i) - v(b, y_j)).$$

If the constraint binds at \underline{b} for income state y_j , then $c(\underline{b}, y_j) = y_j + r\underline{b}$. Thus,

$$\rho v(\underline{b}, y_j) = u(y_j + r\underline{b}) + \lambda^{j \rightarrow i}(v(\underline{b}, y_i) - v(\underline{b}, y_j)).$$

Since the value function is continuous, $\rho v(\underline{b}, y_j) = \lim_{b \rightarrow +\underline{b}} \rho v(b, y_j)$. Therefore:

$$\begin{aligned} & u(y_j + r\underline{b}) + \lambda^{j \rightarrow i}(v(\underline{b}, y_i) - v(\underline{b}, y_j)) \\ &= \lim_{b \rightarrow +\underline{b}} [u(c(b, y_j)) + v_b(b, y_j)(y_j + rb - c(b, y_j)) + \lambda^{j \rightarrow i}(v(b, y_i) - v(b, y_j))], \end{aligned}$$

or simply

$$u(y_j + r\underline{b}) = \lim_{b \rightarrow +\underline{b}} [u(c(b, y_j)) + v_b(b, y_j)(y_j + rb - c(b, y_j))].$$

Using equation (51) gives

$$u(y_j + r\underline{b}) = \lim_{b \rightarrow +\underline{b}} \left[u(c(b, y_j)) + \frac{1}{\beta} u'(c(b, y_j))(y_j + rb - c(b, y_j)) \right].$$

This equation can be rearranged to yield:

$$\lim_{b \rightarrow +\underline{b}} u'(c(b, y_j)) = \beta \frac{\lim_{b \rightarrow +\underline{b}} u(c(b, y_j)) - u(y_j + r\underline{b})}{\lim_{b \rightarrow +\underline{b}} c(b, y_j) - (y_j + r\underline{b})},$$

which is equation (55).

Proof of Proposition 14. The proof of Achdou et al. (2022)’s Proposition 2 applies to the \hat{u} agent, giving

$$\lim_{b \rightarrow \infty} \hat{c}(x) = \frac{\rho - (1 - \gamma)r}{\gamma} b.$$

Since the (sophisticated) IG agent sets $c(x) = \frac{1}{\psi} \hat{c}(x)$ (see Proposition 3), this gives

$$\begin{aligned} \lim_{b \rightarrow \infty} c(x) &= \frac{1}{\psi} \frac{\rho - (1 - \gamma)r}{\gamma} b \\ &= \frac{\rho - (1 - \gamma)r}{\gamma - (1 - \beta)} b. \end{aligned}$$

The proof for $\lim_{b \rightarrow \infty} s(x)$ is similar.

Proof of Corollary 15. Using Itô’s Lemma, $\frac{\mathbb{E}_t dc(x_t)/dt}{c(x_t)} = \frac{c_b(x_t)s(x_t) + \lambda^{j \rightarrow i}(c(b_t, y_i) - c(b_t, y_j))}{c(x_t)}$. Equations (56) and (57) give $\lim_{b \rightarrow \infty} \frac{\mathbb{E}_t dc(x_t)/dt}{c(x_t)} = \frac{r\beta - \rho}{\gamma - (1 - \beta)}$. Taking a derivative with respect to r completes the proof.

F A Cake-Eating Model with a Credit Card

This appendix studies IG preferences in a deterministic cake-eating problem with credit card borrowing (as studied in Proposition 7). For more details see also Lee and Maxted (2023), who explore a similar model in order to illustrate that present bias is capable of producing the high extensive margin but modest intensive margin usage of high-cost debt that is observed empirically.

The Household Balance Sheet. As described in Proposition 7, I consider a simplified version of the general model of Section 3 in which there is just a single liquid asset b and income is deterministic with $y > 0$. Relative to a standard cake-eating model, the one slight complication in the model presented here is that I assume that there is a soft borrowing constraint at $b = 0$. Specifically, the agent earns a constant return of r when $b \geq 0$, but incurs a penalized “credit card” interest rate of $r^{cc} = r + \omega^{cc}$ whenever $b < 0$.

Illustrative Calibration. I assume that the IG agent is sophisticated with $\beta = 0.75$. I set $\gamma = 2$, $r = 3\%$, and $\rho = 5\%$. Deterministic income y is normalized to 1. I consider two calibrations of the credit card interest rate, $r^{cc} = 15\%$ and $r^{cc} = 30\%$. Assumption 1 holds.⁵⁹

Consumption and Saving. The top panel of Figure 7 plots the consumption function for $r^{cc} = 15\%$ (dashed gray) and $r^{cc} = 30\%$ (red). The bottom panel of Figure 7 plots the corresponding saving function $s(b) = y + r(b)b - c(b)$ for both calibrations. The dotted vertical lines mark where the IG agent sets $s(b) = 0$, which in this deterministic model also marks the absorbing point in both calibrations. The IG agent saves to the left of their absorbing point, and dissaves to the right of it. As stated in Proposition 7, note that the IG agent ends up taking on some credit card debt in both cases (i.e., the dotted vertical lines are to the left of $b = 0$).

An interesting feature of Figure 7 is that the consumption and saving functions are identical for $b \geq 0$, and only start to separate once $b < 0$. The separation when $b < 0$ is natural: the agent consumes less and saves more when r^{cc} increases. To understand why the consumption and saving functions are identical when $b \geq 0$, consider the behavior of the standard exponential agent (with $\beta = 1$ and $\rho = 5\%$) in this cake-eating model. Since $r^{cc} > \rho > r$ in both calibrations, the standard exponential agent will dissave whenever $b > 0$, but will always stop at the soft borrowing constraint of $b = 0$ instead of taking on credit card debt. Thus, the standard exponential agent's consumption function for $b \geq 0$ will be unaffected by whether $r^{cc} = 15\%$ or $r^{cc} = 30\%$ (since in either case, the standard exponential agent never gets to the $b < 0$ part of the state space). Turning to the IG agent, since Proposition 3 implies that the sophisticated IG agent consumes $\frac{1}{\psi}$ times the standard exponential agent, the IG agent's consumption function for $b \geq 0$ will also be unaffected by whether $r^{cc} = 15\%$ or $r^{cc} = 30\%$ (despite the IG agent eventually taking on credit card debt).

Value Functions. Next, the top panel of Figure 8 plots the IG agent's continuation-value function $v(b)$ for $r^{cc} = 15\%$ (dashed gray) and $r^{cc} = 30\%$ (red). The bottom

⁵⁹One sufficient condition for Assumption 1 is that \underline{b} is the natural borrowing limit.

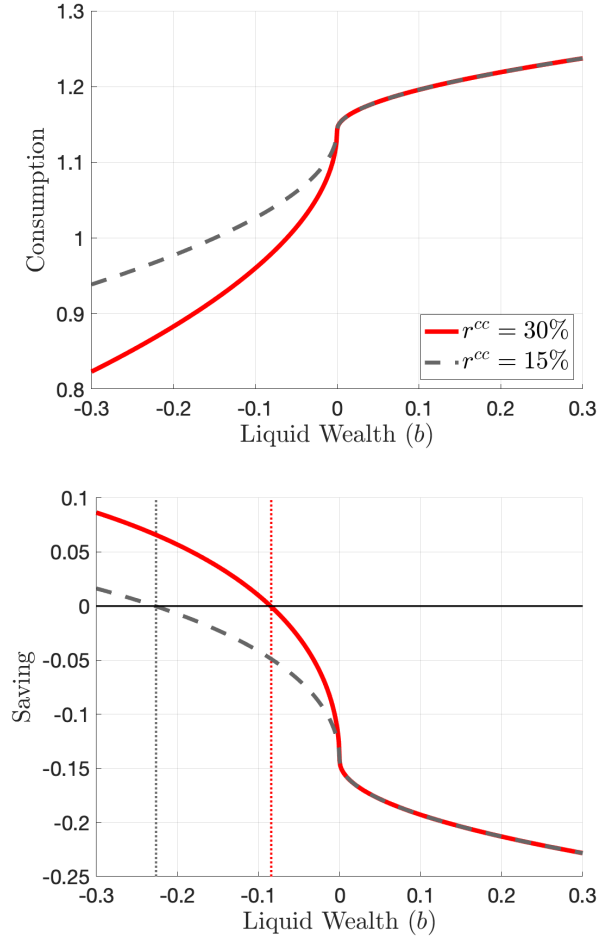


Figure 7: **Equilibrium Consumption-Saving Decisions.** This figure plots the consumption function (top) and the saving function (bottom) for $r^{cc} = 15\%$ (dashed gray) and $r^{cc} = 30\%$ (red).

panel plots the difference between the continuation-value function when $r^{cc} = 30\%$ and the continuation-value function when $r^{cc} = 15\%$.

Similar to Figure 7 above, we again see that the continuation-value functions are identical for $b \geq 0$, and only start to separate once $b < 0$. This result is particularly interesting here in that the IG agent's value function for $b \geq 0$ is unaffected by whether $r^{cc} = 15\%$ or $r^{cc} = 30\%$, despite the fact that the IG agent will eventually borrow at one of those two rates.⁶⁰

⁶⁰Note that here I do not mechanically calculate the IG agent's continuation-value function by applying an affine transformation to the standard exponential agent's value function. Instead, I first calculate the IG agent's consumption function, and then I iterate until convergence on the Bellman

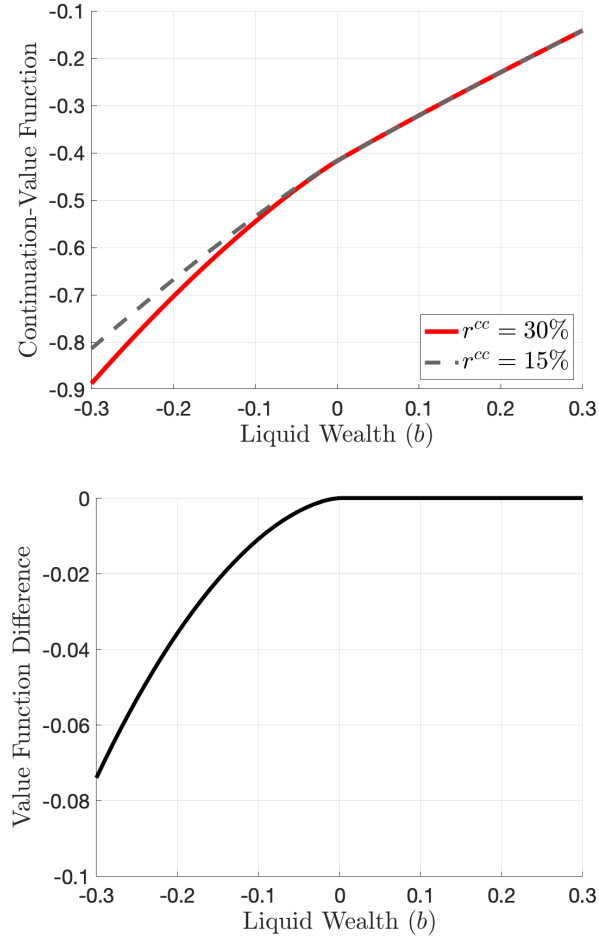


Figure 8: **Equilibrium Value Functions.** This figure plots the continuation-value function $v(b)$ (top) for $r^{cc} = 15\%$ (dashed gray) and $r^{cc} = 30\%$ (red), as well as the difference between the two value functions (bottom).

Intuitively, this result is due to the fact that increasing the borrowing rate from $r^{cc} = 15\%$ to $r^{cc} = 30\%$ introduces offsetting effects for the IG agent. On the one hand, this increased borrowing rate is good for the IG agent because it induces the IG agent to borrow less (see the bottom panel of Figure 7 above). On the other hand, this increased borrowing rate is bad for the IG agent because the borrowing that they still do becomes more costly. When $b \geq 0$ these two effects exactly offset. When $b < 0$, the latter effect dominates and the IG agent prefers the lower borrowing rate of $r^{cc} = 15\%$ (despite the fact that they borrow more in this case).

equation describing the IG agent's continuation-value function, $\rho v(b) = u(c(b)) + v_b(y + r(b)b - c(b))$.

Discussion. A benefit of this cake-eating example is that it provides a simplified special case to illustrate some of the more general welfare results presented in the main text.

First, this finding that the two continuation-value functions are identical for $b \geq 0$ follows from Proposition 2 that there is a value function equivalence (up to an affine transformation) between the IG agent and the standard exponential agent whenever Assumption 1 holds. Since the standard exponential agent never gets to the $b < 0$ part of the state space, their value function $\check{v}(b)$ will be unaffected for all $b \geq 0$ by whether $r^{cc} = 15\%$ or $r^{cc} = 30\%$. Figure 8 shows that the same is true of the sophisticated IG agent, despite the IG agent eventually taking on credit card debt.

Relatedly, this cake-eating example helps to illustrate the tradeoffs that underlie the present-bias dilemma. In particular, while raising the credit card borrowing rate from $r^{cc} = 15\%$ to $r^{cc} = 30\%$ does provide partial commitment in that it induces the IG agent to borrow less, it turns out that the IG agent is nonetheless made (weakly) worse off by this increased borrowing rate. Accordingly, the IG agent will never choose to self-impose a higher borrowing rate, exactly as implied by the present-bias dilemma.

Supplementary References

- Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll, “Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach,” *The Review of Economic Studies*, 2022, 89 (1), 45–86.
- Aiyagari, S. Rao, “Uninsured Idiosyncratic Risk and Aggregate Saving,” *The Quarterly Journal of Economics*, 1994, 109 (3), 659–684.
- Allcott, Hunt, Joshua Kim, Dmitry Taubinsky, and Jonathan Zinman, “Are High-Interest Loans Predatory? Theory and Evidence from Payday Lending,” *The Review of Economic Studies*, 2022, 89 (3), 1041–1084.
- Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg, “The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation,” *Journal of Economic Perspectives*, 2001, 15 (3), 47–68.

- Benhabib, Jess and Alberto Bisin**, “Skewed Wealth Distributions: Theory and Empirics,” *Journal of Economic Literature*, 2018, 56 (4), 1261–91.
- Bewley, Truman**, “Stationary Monetary Equilibrium with a Continuum of Independently Fluctuating Consumers,” *Contributions to Mathematical Economics in Honor of Gérard Debreu*, 1986, 79.
- Campbell, John Y. and N. Gregory Mankiw**, “Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence,” *NBER Macroeconomics Annual*, 1989, 4, 185–216.
- Carroll, Christopher D.**, “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis,” *The Quarterly Journal of Economics*, 1997, 112 (1), 1–55.
- Deaton, Angus**, “Saving and Liquidity Constraints,” *Econometrica*, 1991, 59 (5), 1221–1248.
- Fagereng, Andreas, Martin Blomhoff Holm, Benjamin Moll, and Gisle Natvik**, “Saving Behavior Across the Wealth Distribution: The Importance of Capital Gains,” *NBER w26588*, 2019.
- Floden, Martin and Jesper Lindé**, “Idiosyncratic Risk in the United States and Sweden: Is There a Role for Government Insurance?,” *Review of Economic Dynamics*, 2001, 4 (2), 406–437.
- Gourinchas, Pierre-Olivier and Jonathan A. Parker**, “Consumption over the Life Cycle,” *Econometrica*, 2002, 70 (1), 47–89.
- Guerrieri, Veronica and Guido Lorenzoni**, “Credit Crises, Precautionary Savings, and the Liquidity Trap,” *The Quarterly Journal of Economics*, 2017, 132 (3), 1427–1467.
- Harris, Christopher and David Laibson**, “Instantaneous Gratification,” *Mimeo*, 2004.
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante**, “Quantitative Macroeconomics with Heterogeneous Households,” *Annual Review of Economics*, 2009, 1 (1), 319–354.
- Huggett, Mark**, “The Risk-Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies,” *Journal of Economic Dynamics and Control*, 1993, 17 (5-6), 953–969.
- Imrohoroglu, Ayşe**, “Cost of Business Cycles with Indivisibilities and Liquidity Constraints,” *Journal of Political Economy*, 1989, 97 (6), 1364–1383.

- Kaplan, Greg and Giovanni L. Violante**, “Microeconomic Heterogeneity and Macroeconomic Shocks,” *Journal of Economic Perspectives*, 2018, 32 (3), 167–94.
- , **Benjamin Moll, and Giovanni L. Violante**, “Monetary Policy According to HANK,” *American Economic Review*, 2018, 108 (3), 697–743.
- Krueger, Dirk, Kurt Mitman, and Fabrizio Perri**, “Macroeconomics and Household Heterogeneity,” in “Handbook of Macroeconomics,” Vol. 2, Elsevier, 2016, pp. 843–921.
- Krusell, Per and Anthony A. Smith Jr.**, “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 1998, 106 (5), 867–896.
- Laibson, David**, “Life-Cycle Consumption and Hyperbolic Discount Functions,” *European Economic Review*, 1998, 42 (3-5), 861–871.
- **and Peter Maxted**, “The Beta-Delta-DELTA Sweet Spot,” *NBER w30822*, 2022.
- , – , **and Benjamin Moll**, “Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy,” *NBER w29094*, 2021.
- , – , **and –**, “A Simple Mapping from MPCs to MPXs,” *NBER w29664*, 2022.
- , **Sean Chanwook Lee, Peter Maxted, Andrea Repetto, and Jeremy Tobacman**, “Estimating Discount Functions with Consumption Choices over the Life-cycle,” *Mimeo*, 2023.
- Lee, Sean Chanwook and Peter Maxted**, “Household Borrowing and MPCs in Heterogeneous-Agent Models,” *SSRN 4389878*, 2023.
- Maliar, Lilia and Serguei Maliar**, “The Neoclassical Growth Model with Heterogeneous Quasi-Geometric Consumers,” *Journal of Money, Credit and Banking*, 2006, 38 (3), 635–654.
- Mian, Atif R., Ludwig Straub, and Amir Sufi**, “The Saving Glut of the Rich and the Rise in Household Debt,” *NBER w26941*, 2020.
- Quadrini, Vincenzo, José-Víctor Ríos-Rull et al.**, “Understanding the U.S. Distribution of Wealth,” *Federal Reserve Bank of Minneapolis Quarterly Review*, 1997, 21 (2), 22–36.
- Zeldes, Stephen P.**, “Optimal Consumption with Stochastic Income: Deviations from Certainty Equivalence,” *The Quarterly Journal of Economics*, 1989, 104 (2), 275–298.