Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy*

David Laibson
Harvard University

Peter Maxted
Berkeley Haas

Benjamin Moll
London School of Economics

November 21, 2023

Abstract

We study the effect of monetary and fiscal policy in a heterogeneous-agent model where households have present-biased time preferences and naive beliefs. The model features a liquid asset and illiquid home equity, which households can use as collateral for borrowing. Because present bias substantially increases households' marginal propensity to consume (MPC), present bias increases the impact of fiscal policy. Present bias also amplifies the effect of monetary policy but, at the same time, slows down the speed of monetary transmission. Interest rate cuts incentivize households to conduct cash-out refinances, which become targeted liquidity-injections to high-MPC households. But present bias also introduces a motive for households to procrastinate refinancing their mortgages, which slows down the speed with which this monetary channel operates.

*We are grateful to Adrien Auclert, John Campbell, Xavier Gabaix, Sun Yong Kim, Agnes Kovacs, Lu Liu, Patrick Moran, Dmitry Mukhin, Matt Rognlie, and seminar participants at various institutions for useful comments. Adrien Couturier, Seung Hyeong Lee, Jessica Schultz, and Rui Sousa provided excellent research assistance. The authors acknowledge support from the Pershing Square Fund for Research on the Foundations of Human Behavior.

1 Introduction

The idea that dynamically inconsistent preferences may alter individuals' dynamic choices has a long tradition going back to seminal work by Strotz (1956). A particular form of dynamic inconsistency, present bias, has received empirical support in both laboratory and field studies (e.g., Ashraf et al., 2006; Augenblick et al., 2015; Laibson et al., 2023; and the review by Cohen et al., 2020). Present bias implies that the current self draws a sharp distinction between a util that is experienced now versus a util experienced one time unit in the future, but draws relatively little distinction between a util consumed at any other two successive future dates (Phelps and Pollak, 1968; Laibson, 1997).[1]

The two most commonly used tools of macroeconomic stabilization policy – monetary and fiscal policy – operate in large part by affecting household consumption and investment decisions, two leading examples of the types of dynamic choices that are affected by present bias. It is therefore natural to ask whether and to what extent present bias alters the potency of these policy tools. To answer these questions, we develop and calibrate a heterogeneous-agent consumption model in order to evaluate the impact of present bias on policy outcomes.

Our modeling framework is motivated in part by Campbell's (2006) concept of *positive household finance*: households face a complex financial planning problem, and household behavior is influenced by a range of psychological factors. Our model aims to capture the complexities of household balance sheets that are important for the transmission of monetary and fiscal policy, as well as the channels through which present bias interacts with these balance sheet features.

We set our model in partial equilibrium in order to focus on the details of the household problem. Time is continuous, and we compare the exponential-discounting benchmark to a tractable, and empirically realistic, continuous-time limit of present-biased discounting. In addition to present bias, we assume that households have naive beliefs, meaning that households do not foresee their own future present bias (Strotz, 1956; Akerlof, 1991; O'Donoghue and Rabin, 1999). The modeling of present bias in continuous time builds on the foundational work of Barro (1999) and Luttmer and Mariotti (2003), and our specific approach follows Harris and Laibson (2013).

The household budget constraint includes stochastic auto-correlated labor income and interest rates. On the asset side of the household balance sheet, the model features a liquid savings account and an illiquid home. Households can build a buffer stock of liquid wealth to insure against income fluctuations, and can accumulate home equity by paying down their mortgage. On the liabilities side of the balance sheet, households have access to two forms of debt: credit cards and mortgages. Households can borrow on credit cards up to a

[1] We use "present bias" to refer to quasi-hyperbolic discounting. Other common terms are: "β-δ preferences" in discrete-time models; and "Instantaneous Gratification preferences" in continuous-time models.

calibrated limit. If they have enough home equity, they can also borrow against their home by refinancing their mortgage. We calibrate our heterogeneous-agent model to reproduce two empirical regularities on household balance sheets: the average quantity of credit card debt and the average loan-to-value ratio in the housing market.[2]

In order to focus on the effects of mortgage refinancing, we study homeowners, a large fraction of the population: two-thirds of U.S. housing units are owner-occupied. Homeowners represent an even larger fraction of aggregate income and consumption, and homeowners are an important channel for fiscal (Cloyne and Surico, 2017) and monetary policy (Wong, 2019; Cloyne et al., 2020).

Our main result is that relative to exponential discounting, present bias amplifies the balance-sheet channels of both fiscal and monetary policy, but with some important added subtlety in the case of monetary policy due to refinancing procrastination.

Fiscal policy is powerfully enhanced by present bias, because present bias sharply raises households' average marginal propensity to consume (MPC) (Angeletos et al., 2001). In our Exponential Benchmark model the quarterly MPC is predicted to be 4% and the quarterly marginal propensity for expenditure (MPX), which includes spending on both nondurables and durables, is predicted to be 13%. In our Present-Bias Benchmark, the MPC rises from 4% to 14% and the MPX rises from 13% to 30%. These higher propensities to consume and spend are more consistent with the empirical literature: estimates of the quarterly response of nondurable expenditure are on the order of 15-25%, and estimates of the quarterly response of total expenditure are typically two- to three-times larger.[3]

Present bias also amplifies the overall effect of expansionary monetary policy, but slows down the speed of monetary transmission (an offsetting effect). Interest rate cuts incentivize households to conduct cash-out refinances, which serve as targeted liquidity-injections to households with especially high MPCs (because they are near their liquidity constraint). But present bias with naive beliefs also introduces a motivation for households to procrastinate on refinancing their mortgage, which substantially slows down the speed at which this channel operates. Naive present bias implies that households will delay completing immediate-cost, delayed-reward tasks such as mortgage refinancing, which tends to take weeks and requires the borrower to go through the effortful process of negotiating with lenders, gathering documents, and filling out paperwork. Naive households will keep delaying refinancing, all the

[2]The version of our model with exponential discounting is similar to Guerrieri et al. (2020) with two main differences: we assume that housing is fixed while they model a costly housing adjustment decision, and our model features credit card debt while theirs does not. Also see Mitman (2016), Berger et al. (2018), Berger et al. (2021), Wong (2019), Kaplan et al. (2020), Kinnerud (2021), and Eichenbaum et al. (2022) for related models of housing and mortgage refinancing decisions. Like us, Guerrieri et al. (2020) and McKay and Wieland (2021) use the continuous-time methods of Achdou et al. (2021) to solve their models.

[3]For nondurable spending estimates, see e.g. Johnson et al. (2006), Parker et al. (2013), and the discussions in Kaplan et al. (2018) and Kaplan and Violante (2014). For total spending estimates, see e.g. Parker et al. (2013), Di Maggio et al. (2020), Fagereng et al. (2021), and the discussion in Laibson et al. (2021).

while (counterfactually) believing that the task will get done in the near future.

A noteworthy feature of our model is that present bias amplifies the direct effect of monetary policy on household consumption while, at the same time, also delivering larger MPCs. This is in contrast to standard heterogeneous-agent models, where modeling choices that amplify MPCs typically deliver *smaller* consumption responses to interest rate changes (Werning, 2015; Olivi, 2017; Kaplan et al., 2018; Auclert, 2019; Slacalek et al., 2020).[4] Our model instead delivers a larger responsiveness to monetary policy precisely *because* of the higher MPCs: interest rate cuts incentivize households to conduct cash-out refinances, which become targeted liquidity-injections to high-MPC households.

Though our model is stylized, the steady state of the present-biased economy replicates a variety of empirical patterns from the household finance literature that have, collectively, proven difficult to replicate in models with exponential discounting. The present-biased economy generates empirically-plausible levels of high-cost credit card borrowing by homeowners (Zinman, 2015), cash-out behavior, and loan-to-value (LTV) ratios. It also features a buildup of liquidity-constrained households that is consistent with empirical estimates of households' propensity to spend out of increases in credit card limits (Gross and Souleles, 2002; Agarwal et al., 2018). Present-biased households struggle to smooth consumption, resulting in a consumption function with discontinuities at the borrowing constraint (Ganong and Noel, 2019). Present bias delivers larger MPCs and MPXs, as well as MPCs and MPXs that remain elevated for large shocks (Fagereng et al., 2021).[5] The time-profile of consumer spending is consistent with the intertemporal MPC evidence in Auclert et al. (2018) (using data from Fagereng et al. (2021)). The present-biased economy also generates differential MPCs out of liquid cash transfers versus illiquid home equity increases, a pattern shown empirically by Ganong and Noel (2020). Finally, there is a large literature documenting refinancing inertia: the proclivity for households to delay refinancing when it is financially optimal to do so (e.g., Keys et al., 2016; Johnson et al., 2019; Andersen et al., 2020). We show that present bias with naive beliefs provides a natural motivation for this behavior.

One may wonder: what is specific to present bias? Why not just calibrate a model with exponential discounting that generates empirically realistic MPCs and use that for our policy experiments? The answer is twofold. First, such a model would not generate the procrastination behavior just described. Second, in two-asset models like ours, high-MPC calibrations with exponential discounting often make assumptions about interest rates that

[4]This statement is made precise by Auclert (2019) and Olivi (2017) who show that, in a standard one-asset consumption-saving problem, a household's MPC is a "sufficient statistic" to determine both income and substitution effects of interest rate changes and, in particular, enters the substitution effect with a negative sign. See also the exposition in Slacalek et al. (2020).

[5]Besides present bias, other deviations from standard exponential discounting also have the potential to help match empirical MPCs. For example, Attanasio et al. (2020) show that temptation preferences can help in this regard, and Lian (2021) shows that the anticipation of future mistakes can increase MPCs.

are difficult to reconcile with the data.[6] In contrast, our calibration with present bias delivers high MPCs with interest rates "taken from the data." We view this reconciliation as a step forward for the heterogeneous-agent literature.

Section 2 lays out our model of the household balance sheet. Section 3 characterizes the effect of present bias on household consumption and refinancing decisions. Section 4 discusses our calibration and what it implies for the model's steady state. Section 5 presents our main results about macroeconomic stabilization policy. Section 6 concludes.

2 A Model of Household Finances with Present Bias

Our model is set in partial equilibrium. The goal of the model is to capture household balance-sheet channels through which present bias can impact fiscal and monetary policy. Abstracting from general equilibrium considerations simplifies the analysis and allows for a richer investigation of the institutional factors that affect the household problem. Our partial equilibrium results should be interpreted as just one part of the overall macroeconomic analysis, providing inputs for a full general equilibrium analysis.[7]

2.1 The Household Balance Sheet

Our model focuses on homeowners. There is a unit mass of households that are heterogeneous in their wealth and income. Here we outline the evolution of each household's balance sheet.

Budget Constraints. Each household faces idiosyncratic income risk. The household's income is denoted y_t, and y_t follows a finite state Poisson process. We normalize the average income flow to 1.

Households hold three types of assets: liquid wealth b_t, illiquid housing h, and a fixed-rate mortgage m_t. For simplicity we assume that each household is endowed with a home of fixed value h.[8] The remainder of the household balance sheet evolves as follows:

$$\dot{b}_t = y_t + r_t b_t + \omega^{cc} b_t^- - (r_t^m + \xi) m_t - c_t, \tag{1}$$

$$\dot{m}_t = -\xi m_t, \tag{2}$$

[6] Specifically, such models often calibrate relatively low interest rates on credit card debt and/or relatively high returns on illiquid assets in order to generate the low levels of liquid wealth accumulation and high levels of credit card borrowing that are observed empirically.

[7] We return to this theme in Appendix G, which briefly discusses how present bias may alter the transmission of macroeconomic policy in general equilibrium.

[8] We study only the short-run response to fiscal and monetary policy, and house prices are slow-moving over short horizons (Case and Shiller, 1989). Appendix D.2.1 presents an extension with house price shocks.

subject to the borrowing constraint $b_t \geq \underline{b}$ and the loan-to-value (LTV) constraint $m_t \in [0, \theta h]$. Equation (1) characterizes the evolution of liquid wealth b_t. Equation (2) describes the evolution of mortgage balances. Explaining equation (1), households earn income y_t and have a consumption outflow of c_t. The return to liquid wealth is given by $r_t b_t + \omega^{cc} b_t^-$, where $b_t^- = \min\{b_t, 0\}$. $\omega^{cc} > 0$ is a credit card borrowing wedge, which generates a "soft constraint" at $b = 0$ (Kaplan and Violante, 2014; Achdou et al., 2021). The household's total mortgage payment is captured by $(r_t^m + \xi)m_t$. This is composed of a mortgage interest payment $(r_t^m \times m_t)$ and a principal repayment $(\xi \times m_t)$. To economize on state variables, we make the slightly non-standard assumption that the household pays down its mortgage at a constant proportional rate ξ (Agarwal et al., 2013). The more realistic assumption of a constant flow payment would require an additional state variable.

The borrowing constraint is important for our results, so we emphasize its effects here. In continuous time, the consumption rate c_t is unconstrained for all $b_t > \underline{b}$: any finite rate of consumption can be adopted without violating the borrowing constraint, so long as that rate of consumption persists for a short enough period of time (Achdou et al., 2021). However, at the liquid-wealth constraint of $b_t = \underline{b}$ the household is restricted to consume at a rate

$$c_t \leq y_t + (r_t + \omega^{cc})\underline{b} - (r_t^m + \xi)m_t. \tag{3}$$

That is, consumption features an occasionally-binding constraint when $b_t = \underline{b}$.

Equations (1) and (2) characterize how the household's balance sheet evolves continuously. Households also have the option of paying a fixed cost to discretely adjust their mortgage. We provide details of this discrete adjustment process, which includes the option to refinance, further below after outlining the interest rate process.

Interest Rates and Stimulus Payments. The real interest rate is denoted r_t. We assume that r_t follows a finite state Poisson process. Since our goal is to study the effect of changes in mortgage rates we treat r_t as a long rate (e.g., 10-year TIPS). When discussing monetary policy in Section 5, we implicitly assume that the Federal Reserve is implementing the necessary short rate adjustments to generate the corresponding changes in long rate r_t. Because this paper studies the refinancing channel of monetary policy, it is important that households have reasonable expectations about how mortgage rates evolve over time.[9] Our calibration of households' interest-rate expectations is discussed in Section 4.1. Each household pays a mortgage interest rate r_t^m. To capture features of the U.S. mortgage market, this mortgage rate is fixed until the household decides to refinance. At the time of refinancing, the household switches to a new mortgage rate of $r_t^m = r_t + \omega^m$, where ω^m

[9] Importantly, this means that interest rate shocks in our model are not "MIT shocks." This feature differentiates our model from many other heterogeneous-agent models that study monetary policy.

represents the mortgage borrowing wedge over the current interest rate r_t.

When discussing fiscal policy in Section 5, we consider an unexpected one-time stimulus payment to each household. While our analysis is set in partial equilibrium, we do impose a government budget constraint by levying a flow income tax on all households to finance the initial stimulus payment.

Discrete Adjustments and Mortgage Refinancing. Equations (1) and (2) describe how liquid wealth and mortgage balances evolve continuously. Households also have two methods for discretely adjusting their balance sheet position between liquid wealth and illiquid home equity. First, households can pay a small fixed cost to *prepay* their mortgage. Second, households can pay a larger fixed cost to *refinance*. We detail these options below.

A household's first adjustment option is to prepay its mortgage. We introduce prepayment because mortgage contracts typically allow for households to pay down their mortgage faster than contractually required. Prepayment requires a small fixed cost of κ^{prepay},[10] and the household chooses a new liquid wealth value b' and a new mortgage value m' such that

$$b' - m' = b_t - m_t - \kappa^{prepay}, \text{ subject to } m' \in [0, m_t) \text{ and } b' \geq \underline{b}. \tag{4}$$

Prepayment does not affect the mortgage interest rate, which remains the same as before the adjustment decision. By using part of their liquid wealth to prepay their mortgage, households are effectively shifting their portfolio from a low-return liquid asset to a high-return illiquid asset.[11] Home equity is illiquid because it can only be accessed through a cash-out refinance. But, the benefit of accumulating home equity is that wedge ω^m makes mortgage debt costly. Households will therefore first build a buffer-stock of liquid wealth and then use additional liquidity to prepay their mortgage.

A household's second adjustment option is to refinance its mortgage (or take out a new one if $m_t = 0$).[12] Refinancing requires payment of a fixed cost of κ^{refi}. When refinancing, the household chooses a new liquid wealth value b' and a new mortgage value m' such that

$$b' - m' = b_t - m_t - \kappa^{refi}, \text{ subject to } m' \in [0, \theta h] \text{ and } b' \geq \underline{b}. \tag{5}$$

By refinancing, the household also resets the interest rate on its mortgage to $r_t^m = r_t + \omega^m$.

[10] We impose a small fixed cost for numerical stability. This small fixed cost can be thought of as capturing, for example, the time cost of the additional budgeting required to make a mortgage prepayment.

[11] Kaplan and Violante (2014) highlight how high-return illiquid assets can prevent households from building sizable liquid buffer stocks. Laibson et al. (2023) demonstrate that asset illiquidity combined with present bias allows lifecycle consumption-saving models to match the joint accumulation of credit card debt and illiquid savings that characterizes the balance sheets of many U.S. households.

[12] In our model, some households will have fully paid off their mortgage ($m_t = 0$) prior to taking cash out. For simplicity we refer to adjustments starting from both $m_t > 0$ and $m_t = 0$ as refinancing.

Though refinancing requires up-front costs, there are two reasons why households may choose to refinance. First, if the market interest rate falls then refinancing "locks in" lower mortgage interest payments. Second, refinancing allows households to rebalance their asset allocation across liquid wealth and illiquid home equity. For example, a cash-out refinance lets households convert illiquid home equity into liquid wealth. Accessing home equity is useful during spells of low income (consumption smoothing),[13] and also as a means of converting costly credit card debt into cheaper mortgage debt. Consistent with the empirical evidence in Berger et al. (2021), we find that these two motives for refinancing are not mutually exclusive in our model. Indeed, one of our main results is that interest rate cuts can be highly stimulative precisely because they induce a wave of home-equity extractions.

We assume that both types of discrete adjustments require a small effort cost $\bar{\varepsilon}$ (in addition to the monetary costs κ^{refi} and κ^{prepay}). This cost $\bar{\varepsilon}$ is intended to capture the effort associated with filling out paperwork, negotiating with mortgage brokers, etc. We will show in Section 2.3 that this effort cost provides a natural mechanism for producing refinancing procrastination. In that section we also slightly generalize the setup by making the effort cost stochastic to capture the idea that households face occasional windows of time in which the marginal cost of making effortful budgeting adjustments is lower than normal, such as a free weekend or a cancelled afternoon meeting.

Finally, we note that our model does not allow for home equity lines of credit, second mortgages, or reverse mortgages. These alternate products are much more likely to be used when interest rates are rising, in order to extract home equity without resetting the entire mortgage balance to a higher interest rate (Bhutta and Keys, 2016). We abstract from these alternate products because our paper focuses on the stimulative effect of rate cuts.

Other Structural Assumptions. To capture exogenous mortgage adjustment dynamics such as moving for a new job, we introduce an exogenous hazard rate λ^F at which households are forced to adjust their mortgage (and pay the cost to either refinance or prepay). We assume that households adjust their mortgage optimally when they are forced to do so.

To capture lifecycle dynamics, we assume that households retire at rate λ^R and are replaced by "first-time homeowners." To avoid needing an additional state variable we model retirement using a "perpetual youth" framework (Blanchard, 1985). A household who retires at time t receives a constant consumption flow of $y^R + \bar{r}(h - m_t + b_t)$ in perpetuity, where y^R is a fixed retirement income flow and $\bar{r}$ is the average interest rate. We denote the exponentially discounted value of the retirement consumption flow by $v^R(b_t, m_t) = [u(y^R + \bar{r}(h - m_t + b_t))]/\rho$, where ρ is an exponential discount rate. This parameterization captures a retirement pension of size y^R plus the annuity value of a household's assets at retirement.

[13] See also Hurst and Stafford (2004) and Chen et al. (2020) for related insights.

Summary. The goal of our model is to provide a simple characterization of the household balance sheet features that are important for the conduct of macroeconomic stabilization policy. Our partial equilibrium model has five state variables: (b, m, y, r^m, r). Liquid wealth b and stochastic income y introduce uninsurable income risk and wealth heterogeneity. Mortgage m introduces a realistic role for housing, which is the primary illiquid asset held by most American households (Campbell, 2006). Time-varying interest rate r provides a role for monetary policy. Mortgage interest rate r^m introduces a refinancing motive, and allows us to study the refinancing channel of monetary policy.

To simplify notation, let $x = (b, m, y, r^m, r)$ denote the vector of state variables that characterize the household problem. Households can be heterogeneous in dimensions b, m, y and r^m. All households face the same time-varying market interest rate r_t.

2.2 Utility and Value

Utility. Households have CRRA utility over consumption:[14]

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

Time Preferences: Instantaneous Gratification. This paper's key departure from rationality is the household's discount function. Households have naive *Instantaneous Gratification* (IG) time preferences. IG time preferences were first derived by Harris and Laibson (2013), and are extended in Laibson and Maxted (2022) and Maxted (2023).

In discrete time, the quasi-hyperbolic discount function is given by $1, \beta\delta, \beta\delta^2, \beta\delta^3, \ldots$. IG preferences are the continuous-time limit of this discount function, where each self lives for a vanishingly small length of time.[15] For $t \geq 0$, the limiting IG discount function $D(t)$ is:

$$D(t) = \begin{cases} 1 & \text{if } t = 0 \\ \beta e^{-\rho t} & \text{if } t > 0 \end{cases}. \tag{6}$$

Since the current instantaneous self discounts all future selves by factor β, discount function $D(t)$ features a discontinuity at $t = 0$ whenever $\beta < 1$. The IG household values *instantaneous* utility flows, and all later utility is discounted by β. Note that $\beta = 1$ recovers the standard, time-consistent, exponential discount function.

We assume that households are *naive* about their present bias. This means that the

[14] We could also let households earn utility from housing h. We ignore this element since housing h is constant. This assumption is isomorphic to households having separable utility over consumption and housing, or CES utility with a unitary elasticity of substitution.

[15] See Appendix B for a heuristic derivation, and Harris and Laibson (2013) for a rigorous derivation.

current self is unaware of the self-control problems of future selves. Under naive present bias the current self discounts the utility flows of all future selves by $\beta < 1$, while expecting that all future selves will be exponential discounters ($\beta = 1$). We assume naiveté because we need (at least partial) naiveté to generate procrastination from small effort costs. We extend our analysis to partial and full sophistication in Appendix D.5, and briefly discuss the key takeaways from that extension at the end of Section 3.

As detailed in Laibson and Maxted (2022), IG preferences are a mathematically tractable limit case. They are not a psychologically realistic model of time preferences. The temporal division between "now" and "later" is certainly longer than a single instant dt.[16] Nonetheless, Laibson and Maxted (2022) show that discrete-time models with psychologically appropriate time-steps (e.g., each period lasts for one day or one week) are closely approximated by the continuous-time IG model. The current paper leverages the tractability of the IG approximation to study the effect of present bias on macroeconomic stabilization policy.

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

Value Function ($\beta = 1$). We start by presenting the value function for an exponential ($\beta = 1$) household. We present the value function in steps, suppressing notation at first in order to clarify the structure of the household's decision-making problem. As a first step, the value function of a $\beta = 1$ household can be defined as the solution to the sequence problem:

$$\begin{aligned} v(x_0) &= \max_{\{c_t\},\tau} \mathbb{E}_0 \left[\int_0^\tau e^{-\rho t} u(c_t) dt + e^{-\rho\tau}(v^*(x_\tau) - \bar{\varepsilon}) \right] \quad \text{s.t. (1) and (2) hold and with} \\ v^*(x) &= \max \left\{ v^{prepay}(x), v^{refi}(x) \right\} \\ v^{prepay}(x) &= \max_{b',m'} \; v(b', m', y, r^m, r) \quad \text{s.t. prepayment constraint (4) holds} \\ v^{refi}(x) &= \max_{b',m'} \; v(b', m', y, r + \omega^m, r) \quad \text{s.t. refinancing constraint (5) holds} \end{aligned} \tag{7}$$

Equation (7) subsumes all Poisson shocks inside the expectation operator.

The integral $\mathbb{E}_0 \left[\int_0^\tau e^{-\rho t} u(c_t) dt \right]$ captures utility from consumption, which the household chooses continuously. The term $\mathbb{E}_0 \left[e^{-\rho\tau}(v^*(x_\tau) - \bar{\varepsilon}) \right]$ captures discrete adjustment, which the household chooses at time τ (a stopping time). These discrete adjustments form an optimal stopping (option value) problem. Function v^* denotes the optimal value function conditional on adjusting, which also requires effort cost $\bar{\varepsilon}$. Adjustment takes the form of either mortgage prepayment or refinancing. Note that the mortgage interest rate remains

[16] Augenblick (2018) finds that the division between "now" and "later" is roughly 2 hours. Using fMRI data, McClure et al. (2007) find a one-hour discount rate of 50% for food rewards. More generally, Augenblick (2018) and Augenblick and Rabin (2019) show that almost all discounting occurs within one week.

constant when the household chooses a mortgage prepayment, while refinancing resets the mortgage interest rate to $r_t + \omega^m$.

Equation (7) highlights that the household faces a simultaneous optimal control problem plus an optimal stopping problem. The household continuously chooses consumption c_t, and also possesses the option to discretely rebalance its asset allocation across liquid wealth and illiquid home equity. To capture these dual decisions, the value function in equation (7) can also be expressed as a Hamilton-Jacobi-Bellman Quasi-Variational Inequality (HJBQVI).[17] Starting with compact notation to highlight the general structure of variational inequalities:

$$\rho v(x) = \max\left\{\max_c u(c) + (\mathcal{A}v)(x)\ ,\ \rho(v^*(x) - \bar{\varepsilon})\right\}. \tag{8}$$

Operator $\mathcal{A}$ is an infinitesimal generator, which we will define momentarily by writing out equation (8) in a less compact fashion. The left branch of equation (8) captures the optimal control problem while the right branch captures the optimal stopping problem. If it is not optimal to adjust, the left branch imposes that value function $v(x)$ satisfies a standard HJB equation $\rho v(x) = \max_c u(c) + (\mathcal{A}v)(x)$, and the right branch imposes that $v(x)$ is larger than the value of adjusting: $v(x) \geq v^*(x) - \bar{\varepsilon}$. If it is optimal to adjust, the right branch imposes that the value function equals the value of adjusting: $v(x) = v^*(x) - \bar{\varepsilon}$.

Expanding the operator $\mathcal{A}$ to explicitly show the Poisson risks faced by the household, the HJBQVI can be written out fully as follows:

$$\begin{aligned}\rho v(x) = \max\Bigg\{ \max_c \Bigg\{ & u(c) + \frac{\partial v(x)}{\partial b}(y + rb + \omega^{cc}b^- - (r^m + \xi)m - c)\Bigg\} \\ & - \frac{\partial v(x)}{\partial m}(\xi m) \\ & + \sum_{y' \neq y} \lambda^{y \to y'}\left[v(b, m, y', r^m, r) - v(b, m, y, r^m, r)\right] \\ & + \sum_{r' \neq r} \lambda^{r \to r'}\left[v(b, m, y, r^m, r') - v(b, m, y, r^m, r)\right] \\ & + \lambda^R\left[v^R(x) - v(x)\right] \\ & + \lambda^F\left[v^*(x) - (v(x) - \bar{\varepsilon})\right], \\ & \rho(v^*(x) - \bar{\varepsilon})\Bigg\}.\end{aligned} \tag{8'}$$

[17] See Bensoussan and Lions (1982, 1984) and Bardi and Capuzzo-Dolcetta (1997). For additional details and a discussion of more economic applications, see http://benjaminmoll.com/liquid_illiquid_numerical/, and Guerrieri et al. (2020) and McKay and Wieland (2021). Relative to our formulation in equation (8), the mathematics literature studying HJBQVIs typically uses somewhat different but equivalent notation, for example: $0 = \min\left\{\rho v(x) - [\max_c u(c) + (\mathcal{A}v)(x)]\ ,\ v(x) - (v^*(x) - \bar{\varepsilon})\right\}$. We use the formulation in (8) with the max operator because it is more economically intuitive.

Each of the seven rows of equation (8′) reflects the value function's dependence on, respectively, liquid wealth b (row one), mortgage level m (row two), income state y (row three), interest rate state r (row four), retirement (row five), forced adjustment (row six), and discrete adjustment (row seven). In rows three and four of the equation, which correspond respectively to the income process and the interest rate process, we use notation $\lambda^{s \to s'}$ to denote the transition rate from state s to state s'. λ^R is the transition rate into retirement, and λ^F is the rate at which households are forced to adjust their mortgage.

Value Functions ($\beta < 1$)**.** We now introduce naive present bias. Naifs incorrectly perceive that all future selves will discount exponentially ($\beta = 1$). Thus, the value function $v(x)$ that solves equation (7) – or equivalently (8) – characterizes the naive IG household's perceived value function starting in the next instant. For this reason, we will refer to $v(x)$ as the *continuation-value function*. The *current-value function* characterizes the household's perceived value of future utility flows in the current period. Since the current self discounts all future selves by β, the current-value function of the naive IG household is given by:

$$
\begin{aligned}
w(x) &= \max \left\{ \beta v(x),\ w^*(x) - \bar{\varepsilon} \right\} \quad \text{with} \\
w^*(x) &= \max \left\{ w^{prepay}(x), w^{refi}(x) \right\} \\
w^{prepay}(x) &= \max_{b', m'}\ w(b', m', y, r^m, r) \quad \text{ s.t. prepayment constraint (4) holds} \\
w^{refi}(x) &= \max_{b', m'}\ w(b', m', y, r + \omega^m, r) \quad \text{ s.t. refinancing constraint (5) holds}
\end{aligned}
\tag{9}
$$

In the first line of equation (9), the left branch captures the current-value function if the household does not adjust its mortgage, and the right branch captures the current-value function if the household chooses to adjust its mortgage. Importantly, the utility flow of the current self does not show up in equation (9). Because each self lives for a single instant of dt, no individual self's utility flow has a measurable effect on the overall value function. Further discussion is provided in Appendix B, which derives the current-value function of the naive IG household as the continuous-time limit of a discrete-time model.

It is worth emphasizing again that, with naive present bias, the perceived value function v on the right-hand side of (9) *is the value function of a* $\beta = 1$ *household*, i.e. the one that solves (8). This property of naiveté is the key reason for its tractability. In particular, it implies that both theoretical and computational approaches can use the following two-step procedure: first, solve the value function v of an exponential $\beta = 1$ household from (8); second, find the value function of a present-biased $\beta < 1$ household immediately from (9).

Policy Functions. Households make two choices in the model: they choose consumption continuously, and they have the option to adjust their mortgage discretely. Introducing notation for these policy functions, let $c : x \to [0, \infty)$ denote a household's consumption policy in state $x = (b, m, y, r^m, r)$. Mortgage adjustment consists of two nested decisions. First, conditional on adjustment the household chooses its new mortgage level m' and liquid wealth b'. Let $m' : x \to [0, \theta h]$ and $b' : x \to [\underline{b}, \infty)$ denote the household's optimal mortgage and liquid wealth choice, conditional on adjustment. Second, the household chooses whether or not to adjust. Let $\mathfrak{R} : x \to \{0, 1, 2\}$ denote whether a household finds it optimal to not adjust ($\mathfrak{R} = 0$), prepay ($\mathfrak{R} = 1$), or refinance ($\mathfrak{R} = 2$).

2.3 Procrastination

There is a large literature documenting that households are slow to refinance after interest rate declines (e.g., Keys et al., 2016; Johnson et al., 2019; Andersen et al., 2020). Refinancing involves a series of up-front effort costs, such as negotiating with mortgage brokers and filling out paperwork, in exchange for long-run financial benefits. Households with naive present bias will delay completing these sorts of immediate-benefit delayed-reward tasks, instead deferring them for future selves (O'Donoghue and Rabin, 1999, 2001; DellaVigna and Malmendier, 2004). In this way, naive present bias provides one natural motivation for refinancing inertia: *procrastination*. From a theoretical standpoint, our incorporation of present-bias-driven procrastination also differentiates the analysis here from other papers using IG preferences, such as Harris and Laibson (2013) and Maxted (2023).

Keys et al. (2016) provide direct evidence of procrastination as an important channel through which refinancing inertia arises. The financial calculations involved in refinancing are complex (Agarwal et al., 2013), and refinancing generates a range of non-pecuniary short-term costs in exchange for uncertain long-term benefits. This creates an environment where a variety of psychological factors – such as trust in the financial system, financial illiteracy, sticky information, attention costs, and bounded rationality – underlie the effort costs that drive procrastination.[18] Our goal here is to provide a simple and transparent model that captures the intuition for how such cognitive costs can interact with present bias to produce refinancing inertia.

The key model ingredient that generates procrastination for $\beta < 1$ households is the effort cost that households face to discretely adjust their mortgage. While the setup with a constant effort cost that we spelled out in Sections 2.1 and 2.2 already gives rise to procrastination, we here generalize this setup slightly by assuming that the effort cost is stochastic. This stochasticity captures the idea that households face occasional windows of time in which the

[18]For references to some of these factors, see respectively Johnson et al. (2019); Agarwal et al. (2017); Mankiw and Reis (2002); Sims (2003); Woodford (2003); Gabaix (2019).

marginal cost of making effortful financial adjustments is lower than normal. As we explain below, a stochastic effort cost will generate a particular form of procrastination, namely "Calvo-style" procrastination, that is not only tractable but also a useful approximation of household refinancing behavior (Andersen et al., 2020).

Assumption 1 *The effort cost ε_t takes two values $\underline{\varepsilon}$ and $\bar{\varepsilon}$ satisfying*

$$\bar{\varepsilon} > \frac{1}{\beta}\underline{\varepsilon} > 0.$$

Effort cost ε_t usually equals the high value of $\bar{\varepsilon}$. But, at Poisson arrival rate ϕ, ε_t falls to the low value of $\underline{\varepsilon}$ for a single instant before immediately reverting back to the high value of $\bar{\varepsilon}$.

Though stylized with a two-state process for tractability, Assumption 1 captures the sorts of stochastic life events, such as a free weekend or a cancelled afternoon meeting, in which the household becomes more willing to complete chores because the opportunity cost of time has temporarily fallen.

What is critical about these sorts of stochastic windows of availability is that they typically have explicit end dates. For example, a free weekend may represent a low-cost period for a household, but that window closes on Sunday night. These sorts of deadlines are forcing mechanisms that encourage present-biased households to complete effortful tasks, because even present-biased households will want to take advantage of relatively low-cost periods before they come to an end (see e.g. Carroll et al., 2009; Allcott et al., 2022).

Assumption 1 makes the simplification that these low-cost windows last for exactly one instant dt. This simplification maintains the stationarity of our continuous-time model (avoiding the need to include time as a state variable), but it is not critical to the results. The essential feature of these low-cost windows is that they have a defined expiration date.

This simple generalization of our model to a stochastic effort cost requires us to append our model equations in a few places. Appendix B.3 spells out the full set of equations. For example, in addition to the current-value of a $\beta < 1$ household in the high-effort-cost state defined in equation (9), there is now an analogous equation for a household in the low-effort-cost state:

$$\underline{w}(x) = \max\left\{\beta v(x),\ w^*(x) - \underline{\varepsilon}\right\}. \tag{10}$$

Intuitively, since the low-cost state only lasts for an instant (Assumption 1), the household either takes advantage of adjusting its mortgage at the lower effort cost $\underline{\varepsilon}$ in which case its value is $w^*(x) - \underline{\varepsilon}$, or else the household reverts to the high-cost state and its value is $\beta v(x)$. Equation (10) highlights the cost of not refinancing when $\varepsilon = \underline{\varepsilon}$ — the low-cost period is lost and the effort cost reverts back to $\bar{\varepsilon}$. For future reference, we will denote by $\underline{\mathfrak{R}}(x), \underline{m}'(x)$, and $\underline{b}'(x)$ the corresponding refinancing policy function in the low-effort-cost state (i.e., the prepayment or refinancing decisions corresponding to (10)), while $\mathfrak{R}(x), m'(x)$, and $b'(x)$ continue to denote the analogous policy function when the effort cost is high.

Next, we make an assumption so that the effort cost only matters for $\beta < 1$ households:

Assumption 2 *$\bar{\varepsilon}$ and $\underline{\varepsilon}$ are vanishingly small.*

Assumption 2 represents the idea that households typically consider refinancing to be a nuisance, but not costly in an economically meaningful sense. By making the effort cost arbitrarily small, the effort cost is inconsequential for the behavior of $\beta = 1$ households.

However, this trivial effort cost becomes important when *interacted* with present bias. When $\beta < 1$, the small effort cost is sufficient to generate procrastination.[19] This is because naive present-biased households will always choose to delay the task of refinancing (for one instant in expectation) whenever $\varepsilon_t = \bar{\varepsilon}$. The perceived benefit of procrastinating is that the effort cost of adjustment gets pushed into the future, where it is discounted by β. When $\varepsilon_t = \bar{\varepsilon}$, the perceived cost of delaying for one instant is infinitesimal. So, naive present-biased households will continually procrastinate when $\varepsilon_t = \bar{\varepsilon}$.

For $\beta < 1$ households, procrastination persists until the household stochastically enters a low-cost window and the effort cost ε_t momentarily drops to $\underline{\varepsilon}$. Now, there is an explicit cost to waiting: further procrastination causes the low-cost window to expire and the effort cost to revert to $\bar{\varepsilon} > \frac{1}{\beta}\underline{\varepsilon}$. Since the effort cost of completion now, $\underline{\varepsilon}$, is less than the discounted effort cost of completion next instant, $\beta\bar{\varepsilon}$, this fleeting opportunity incentivizes the present-biased household to stop procrastinating. This is formalized in Proposition 2 below.[20]

3 The Effect of Present Bias on Policy Functions

We can now describe the effect of present bias on the consumption and mortgage adjustment decisions. In order to characterize the policy functions of a present-biased household relative to those of a standard exponential household, we use hat-notation to denote the policy functions of an otherwise-identical household that has $\beta = 1$. Accordingly, the policy functions denoted by hats are what the naive household perceives all future selves will follow.

[19] Of course, when $\beta = 1$ it is possible to rationalize refinancing inertia by making other assumptions about effort cost ε_t, for example by making $\bar{\varepsilon}$ arbitrarily large.

[20] As this discussion makes clear (and we confirm in Proposition 2), our baseline model in Sections 2.1 and 2.2 with a constant effort cost $\bar{\varepsilon}$ would generate indefinite procrastination by $\beta < 1$ households. Perhaps surprisingly, this is true even though this effort cost is arbitrarily small (Assumption 2). Thus, an additional rationale for extending our model to feature a stochastic effort cost, besides increased realism (low-cost windows like weekends and cancelled meetings do exist), is that it generates $\beta < 1$ households who refinance.

Consumption. Present-biased households want to bring utility into the present, which implies that present bias has a direct effect on households' consumption decisions. Specifically, present-biased households overconsume by factor $\beta^{-\frac{1}{\gamma}}$:

Proposition 1 (Continuous Control)

1. *For all* $b > \underline{b}$*, the household sets* $c(x) = \beta^{-\frac{1}{\gamma}}\widehat{c}(x)$*.*
2. *For* $b = \underline{b}$*, the household sets* $c(x) = \min\left\{\beta^{-\frac{1}{\gamma}}\widehat{c}(x),\ y + (r + \omega^{cc})\underline{b} - (r^m + \xi)m\right\}$*.*

The proof of Proposition 1 makes use of an important intermediate step, which we state first and which is proved in Appendix B:

Lemma 1 *When the borrowing constraint does not bind,* $b > \underline{b}$*, consumption is defined implicitly by the first-order condition:*

$$u'(c(x)) = \beta \frac{\partial v(x)}{\partial b}, \tag{11}$$

where the continuation-value function v *is equal to the value function of an exponential* $\beta = 1$ *household and solves* (7) *or equivalently* (8)*.*

Equation (11) is a first-order condition: consume until the marginal utility of consumption equals the marginal value of liquid wealth. For $\beta = 1$, the standard continuous-time first-order condition of $u'(c(x)) = \frac{\partial v(x)}{\partial b}$ is recovered. The additional discount factor β appears in equation (11) because the present-biased household discounts future consumption (and hence current wealth) by β.[21] It is worth reemphasizing that naiveté implies that the continuation-value function v in (11) is the value function of an exponential $\beta = 1$ household. This means that, similarly to the value function, one can recover the consumption policy function with a two-step procedure: first, solve the value function v of an exponential $\beta = 1$ household from (8); second, find the consumption policy function of a present-biased $\beta < 1$ household from (11) (with an additional condition when $b = \underline{b}$).

Proof of Proposition 1. Let $\widehat{c}(x)$ denote the consumption function that the naive household expects all future selves to adopt ($\beta = 1$). Expanding (11) under CRRA utility:

$$c(x)^{-\gamma} = \beta \frac{\partial v(x)}{\partial b} \quad \text{and} \quad \widehat{c}(x)^{-\gamma} = \frac{\partial v(x)}{\partial b}.$$

[21] While (11) looks like a standard discrete-time first-order condition, it is important to note that the interpretation is different. In particular, β is *not* the standard exponential discrete-time discount factor; instead, β is the discount factor between now and one instant from now, see (6), with $\beta = 1$ corresponding to exponential discounting, i.e. no present bias.

Rearranging gives $c(x) = \beta^{-\frac{1}{\gamma}}\widehat{c}(x)$. This holds as long as $b > \underline{b}$. For $b = \underline{b}$, overconsumption will be restricted if the borrowing constraint binds (see equation (3)). ■

Proposition 1 provides a tractable formula that relates the IG household's consumption to that of a standard exponential household. This can be used to construct an Euler equation for the IG household:

Corollary 1 (Maxted (2023)) *Let $r_t(b_t)$ denote the household's effective borrowing cost: $r_t(b_t) = r_t$ if $b_t \geq 0$, and $r_t(b_t) = r_t + \omega^{cc}$ if $b_t < 0$. Whenever $c(x_t)$ is locally differentiable in b, consumption obeys the following Euler equation:*

$$\mathbb{E}_t \frac{du'(c(x_t))/dt}{u'(c(x_t))} = \left[\rho + \gamma\left(1 - \beta^{\frac{1}{\gamma}}\right)\frac{\partial c(x_t)}{\partial b}\right] - r_t(b_t). \tag{12}$$

Proof. Appendix C extends the proof of Maxted (2023) to our environment. ■

Equation (12) is the naive, continuous-time, analogue of the Hyperbolic Euler Relation for sophisticates of Harris and Laibson (2001). The growth rate of marginal utility is given by $\mathbb{E}_t \frac{du'(c(x_t))/dt}{u'(c(x_t))}$, and the term in brackets can be interpreted as the household's effective discount rate at time t. When $\beta = 1$, equation (12) reduces to a standard Euler equation: $\mathbb{E}_t \frac{du'(c(x_t))/dt}{u'(c(x_t))} = \rho - r_t(b_t)$. When $\beta < 1$, the household's effective discount rate is increasing in its instantaneous MPC, $\frac{\partial c(x_t)}{\partial b}$. Intuitively, overconsumption will have a larger effect on the growth rate of marginal utility when consumption itself is sensitive to b_t.

For $\beta < 1$, an important consequence of equation (12) is that the household's effective discount rate varies over the state space. In particular, since households with a higher instantaneous MPC will consume more impatiently, households near $b = \underline{b}$ and $b = 0$ will act more impatiently, while households with plentiful liquidity will act more patiently.

Mortgage Adjustment. Next we characterize the effect of present bias on the mortgage adjustment decision. To this end, recall that $\mathfrak{R}(x) \in \{0, 1, 2\}$ and $\underline{\mathfrak{R}}(x) \in \{0, 1, 2\}$ denote the household's decision to not adjust, prepay, or refinance in the high- and low-effort-cost states, and $m'(x), b'(x), \underline{m}'(x)$, and $\underline{b}'(x)$ denote its adjustment targets conditional on adjusting. Unlike consumption, present bias has a muted impact on these adjustment decisions. In particular, the only way that $\beta < 1$ affects the mortgage adjustment decision is through procrastination.[22] This is formalized in the following proposition:

Proposition 2 (Optimal Stopping)

1. *Adjustment targets m' and b' are independent of β. Thus, $m'(x) = \widehat{m}'(x)$, $b'(x) = \widehat{b}'(x)$, $\underline{m}'(x) = \widehat{\underline{m}}'(x)$, and $\underline{b}'(x) = \widehat{\underline{b}}'(x)$ for all x.*

[22] Accordingly, setting $\varepsilon_t \equiv 0$ would remove the effect of present bias on the mortgage adjustment decision.

2. (a) *For $\beta = 1$, the refinancing policy function $\mathfrak{R}(x)$ converges pointwise to $\underline{\mathfrak{R}}(x)$ as the effort cost vanishes. This effectively means that the $\beta = 1$ household's mortgage adjustment behavior does not depend on the state of the effort cost.*

   (b) *For $\beta < 1$ and $\varepsilon = \bar{\varepsilon}$, $\mathfrak{R}(x) = 0$ for all x. This means that the present-biased household procrastinates and will not adjust its mortgage when $\varepsilon = \bar{\varepsilon}$.*

   (c) *For $\beta < 1$ and $\varepsilon = \underline{\varepsilon}$, $\mathfrak{R}(x)$ converges pointwise to $\widehat{\underline{\mathfrak{R}}}(x)$ as the effort cost vanishes. This effectively means that the present-biased household does not procrastinate when $\varepsilon = \underline{\varepsilon}$.*

Proof. See Appendix C. The intuition for clause 1 is that the current self composes only an infinitesimal part of the household's overall value function. This implies that $w^{prepay}(x)$ and $w^{refi}(x)$ in (9) can be rewritten as:

$$w^{prepay}(x) = \max_{b',m'} \ \beta v(b', m', y, r^m, r) \quad \text{s.t. prepayment constraint (4) holds}$$
$$w^{refi}(x) = \max_{b',m'} \ \beta v(b', m', y, r + \omega^m, r) \quad \text{s.t. refinancing constraint (5) holds}$$

Since maximizing βv is equivalent to maximizing v, the same (b', m') will be optimal for the $\beta < 1$ household and the $\beta = 1$ household. The intuition for clause 2 was discussed above in Section 2.3. ■

Clauses 2b and 2c of Proposition 2 state that present-biased households refinance only when they are in a refinancing region of the state space *and* they experience a low-cost window (weekends, cancelled meetings, etc). Recall from Assumption 1 that these low-cost windows occur at Poisson rate ϕ. Therefore the model with naive present bias reproduces the state- plus time-dependent refinancing behavior documented in Andersen et al. (2020).

Cash-Out Refinances as Targeted Liquidity Injections. The combination of Propositions 1 and 2 yields one of our key results: when households are present-biased, interest rate cuts incentivize households to conduct cash-out refinances, which become targeted liquidity-injections to high-MPC households. The explanation proceeds in two steps. First, Proposition 2 states that present bias ($\beta < 1$) does not affect refinancing behavior except through procrastination: conditional on refinancing, the mortgage adjustment and hence the size of any potential cash-out is the same. Second, Proposition 1 implies that present-biased households overconsume: hence they spend down any given cash-out amount faster than exponential discounters. This is precisely what we find when we conduct our monetary policy experiments in Section 5.2 further below.

A Comparison: Consumption versus Mortgage Adjustment. The juxtaposition of Propositions 1 and 2 also highlights the subtleties of present bias' impact on household balance-sheet decisions. In our model, present bias directly affects the consumption decision, whereas present bias only affects the mortgage adjustment decision through procrastination.

Present bias has a differential impact on these decisions because consumption and procrastination are "small" decisions (flow decisions) while discrete adjustment is a "large" decision (a stock decision). The current self wants to overconsume in the moment, but this overconsumption has only an infinitesimal effect on the household's balance sheet. Similarly, procrastination is expected to last only for an instant. The same is not true of the mortgage adjustment decision. This decision discretely adjusts the household's asset allocation between liquid wealth and illiquid home equity. Since each self only lasts for an instant, any short-term benefits from myopia-driven refinancing are dominated by the accumulation of costs borne across all future selves. In short, naive present bias causes a persistent accumulation of small mistakes, not the intermittent occurrence of large mistakes.

Though the proofs of Propositions 1 and 2 utilize IG preferences, the assumption that each self lasts for a single instant is mathematically convenient but not quantitatively necessary (Laibson and Maxted, 2022). Propositions 1 and 2 will be robust so long as the temporal division between the present and the future is relatively short (e.g., one week), meaning that each self composes a negligible part of the overall value function.[23]

The Effect of Present Bias on the Propensity to Refinance. There is an informal intuition in the literature that present bias increases the propensity for households to extract home equity in order to finance near-term consumption (see e.g. Mian and Sufi, 2011; Campbell, 2013). Our model illustrates that the effect of present bias on refinancing is more complex. When effort cost $\varepsilon = \underline{\varepsilon}$, Proposition 2 says that present bias has no effect on the refinancing decision at any given point x. When $\varepsilon = \bar{\varepsilon}$, procrastination slows down the rate at which present-biased households refinance. Both of these effects counter the intuition that present bias increases cash-out refinancing. However, present bias causes a build-up of credit card debt due to overconsumption. This incentivizes home-equity extractions as a means of converting costly credit card debt to cheaper mortgage debt. In summary, present bias inhibits refinancing at any given point x. But, present bias changes the distribution of households over the state space in a way that encourages cash-out refinancing.

[23]For intuition, consider a discrete-time model where each self lives for one week. Let $\delta = \exp\left(\frac{-\rho}{52}\right)$. Assume that each self consumes a constant amount, $\bar{c}$, in each period. Each self has a current value of $u(\bar{c}) + \frac{\beta\delta}{1-\delta}u(\bar{c})$, meaning that the utility of each self composes a share $1/\left(1 + \frac{\beta\delta}{1-\delta}\right)$ of the total value function. Under our benchmark calibration with $\beta = 0.83$ and $\rho = 0.88\%$ (Section 4.2) so that $\delta = \exp\left(\frac{-\rho}{52}\right) = 0.9998$, this share equals 0.02% of the total value function.

Generalization to (Partial) Sophistication. We extend our analysis to partial and full sophistication in Appendix D.5. Letting β^E denote the short-run discount factor that the current self expects all future selves to have, full naiveté means that $\beta^E = 1$. Partial sophistication instead sets $\beta^E \in (\beta, 1)$, such that the current self is partially aware of future selves' self-control problems. Full sophistication is the limiting case of $\beta^E = \beta$.

There are two key takeaways from the extension in Appendix D.5. First, our main-text analysis (which assumes full naiveté) is robust to all but the limiting case of *full* sophistication. In particular, all that is needed to generate refinancing procrastination is for the current self to think that future selves will be *less present biased* than the current self is, meaning that any amount of naiveté is sufficient (i.e., $\beta^E > \beta$). This is because the current self then perceives that the next self will be *more willing* to refinance, and hence pushes refinancing into the future. Second, for the limiting case of *full* sophistication, households do not procrastinate. This second result follows from Assumption 2 that effort costs are vanishingly small. Intuitively, without at least some scope for incorrect expectations, we cannot generate non-vanishing bouts of procrastination from vanishingly small effort costs.[24]

4 Calibration and Steady State Household Behavior

4.1 Externally Calibrated Parameters

We begin by describing the model's externally calibrated parameters. These parameter choices are summarized in Appendix Table 4.

Home Owners in the 2016 SCF. For many of our calibration targets we use the 2016 Survey of Consumer Finances (SCF) wave. We often express these targets relative to a measure of households' permanent income for which – following Kennickell (1995), Kennickell and Lusardi (2004), and Fulford (2015) – we use the SCF's question about "normal income." To align the SCF with the households in our model, we restrict our SCF sample to households that own a home, possess a credit card, have a head aged 25-61, and have a head or spouse in the labor force. The average after-tax permanent income for households in our SCF sample is roughly \$100,000. See Appendix A.2 for details.

Income. We assume that the Poisson income process takes one of three values $y_t \in \{y_L, y_M, y_H\}$. We limit the model to three income states – low, middle, and high – so

[24]Even with full sophistication, procrastination can still arise with non-vanishing effort costs. See the discussion in Appendix D.5.4.

that we can more fully illustrate the resulting model solution (see e.g. Figures 1 and 2).[25] Our calibration of the income process follows Guerrieri and Lorenzoni (2017), who assume that the logarithm of income follows an AR(1) process at a quarterly frequency. We first convert this quarterly AR(1) process to a continuous-time Ornstein-Uhlenbeck (OU) process and then discretize the OU process with a three-state Poisson process using a finite difference method. See Appendix A.3 for details.

Interest Rates. We view r_t as a longer-term interest rate, and assume that the monetary authority adjusts short rates in the background to generate these movements in the long rate. Our model aims to capture larger movements in the federal funds rate, such as those in 2001, 2008, and 2019-20, which have economically significant effects on long rates and hence on mortgage refinancing. The focus of our paper is not on smaller, fine-tuning, movements in the federal funds rate that do not significantly affect mortgage rates.

Rather than treating all shocks to r_t as unexpected "MIT shocks," households in our model have calibrated expectations about the data-generating process for r_t. In order to calibrate household interest rate expectations we estimate an AR(1) process via maximum likelihood estimation using weekly data on the yield of 10-year TIPS from 2003 – 2019. We first convert this weekly AR(1) process to a continuous-time OU process (with rate of mean reversion 0.29 and volatility 0.63%), and then discretize it into a four-state Poisson process with states $r_t \in \{-1\%, 0\%, 1\%, 2\%\}$ using a finite difference method.

We set the credit card wedge ω^{cc} to 10.3% to capture the difference between the commercial bank interest rate charged on credit cards and the 10-year treasury yield from 2015 – 2017. We set $\omega^m = 1.7\%$ to capture the difference between the average 30-year fixed rate mortgage and the 10-year treasury yield from 2015 – 2017.

Assets and Liabilities. Using the sample of home owners in the 2016 SCF, we estimate an average home value to permanent income ratio of 3.29 (roughly $329,000) and therefore set $h = 3.29$. We set θ, the maximum LTV ratio, equal to 0.8. Although this is a tight restriction on the maximum LTV allowed for first-time homebuyers, it is consistent with the maximum LTV available to households conducting a cash-out refinance.[26]

Mortgages are paid down at rate $\xi = 3.5\%$, which generates a 20-year half-life for mortgages.[27] We set the fixed cost of refinancing to $\kappa^{refi} = 0.05$ (approximately $5,000), which

[25] For $\beta = 1$ in consumption-saving models such as this, three income states are the minimum number of states that are needed to generate a mass of households who borrow, a mass of households at the soft constraint, and a mass of households who save (Achdou et al., 2021).

[26] See Greenwald (2018) for data on realized LTVs for first-time homebuyers and for cash-out refinances.

[27] Recall that mortgage payments are not constant in our model. We choose a 20-year half-life for mortgage paydowns so that the mortgage payment required by large mortgages is not exceedingly onerous.

is 3% of the average outstanding mortgage principal.[28] For numerical tractability, we also impose a small cost to prepaying mortgages of $\kappa^{prepay} = 0.002$ (or approximately $200).

We set credit card borrowing limit $\underline{b}$ to one third of permanent income. This is consistent with reported credit limits in the SCF, and in line with typical choices in the literature.[29]

Preferences. Discount function parameters ρ and β are calibrated internally to match home equity and credit card debt levels in the 2016 SCF — see Section 4.2 below. We set the intertemporal elasticity of substitution $\frac{1}{\gamma}$ equal to $\frac{1}{2}$, which is a standard calibration in the consumption-saving literature. We choose a Poisson rate governing procrastination of $\phi = -\ln(0.5)$. This implies that there is a 50% probability that a household for whom it is optimal to refinance will do so within one year, consistent with Andersen et al. (2020).[30]

Other Structural Parameters. We set the rate of forced adjustment to $\lambda^F = \frac{1}{15}$, which approximates the moving rate of homeowners reported in the Current Population Survey's Annual Social and Economic Supplement for 2016. The retirement rate $\lambda^R = \frac{1}{30}$ so that households exist in our model for 30 years on average. We set the retirement income flow y^R equal to the minimum income level y_L, since $y_L = 0.75$ and a retirement replacement rate of 70-80% is a common benchmark. Retired households are dropped from the model, and are replaced by new households with mortgage $m_t = \theta h$ and liquid wealth b_t drawn from a uniform distribution with support $[0, \frac{y_L}{2}]$.

4.2 Internally Calibrated Parameters and Steady State

"Steady State." We start all policy counterfactuals from a "steady state" in which the interest rate has been permanently fixed at $r^* = 1\%$. This implies that all households in our steady state will have a mortgage interest rate of $r_t^m = r^* + \omega^m = 2.7\%$. The assumption that the interest rate has been fixed at 1% is formally just one possible sample path of the interest rate process. This assumption is helpful for pedagogy because it reduces the dimensionality of our steady state by eliminating heterogeneity in r_t^m.[31]

[28]The Federal Reserve's website on refinancing suggests that refinancing costs roughly 3% of outstanding principle. The average LTV in our model is 0.51, suggesting $\kappa^{refi} = 0.03 \times (0.51h)$. This yields roughly $\kappa^{refi} = 0.05$. For details, see https://www.federalreserve.gov/pubs/refinancings/. While we assume a uniform fixed cost for simplicity, refinancing costs vary from state to state and also over time (e.g., the 2017 Tax Cuts and Jobs Act lowered the mortgage interest deduction threshold from $1 million to $750,000 and thus altered cash-out refinancing incentives for some households). Though not explored in depth in the current paper, our results suggest that variation in refinancing costs can affect both the steady-state distribution of households and their subsequent responses to macroeconomic stabilization policy.

[29]For example, Kaplan et al. (2018) calibrate a borrowing limit of one times quarterly labor income.

[30]Andersen et al. (2020) estimate that 84% of households are "asleep" each quarter, and $0.84^4 \approx 0.5$.

[31]This fixed-r assumption implies that once rates fall, all households can lower their mortgage rate by refinancing. This broadly reflects the rate cuts that followed the Early 2000s, Great, and COVID-19 Recessions,

Calibration Cases. We present two benchmark calibration cases for the "steady state": an *Exponential Benchmark* with $\beta = 1$ and a *Present-Bias Benchmark* with $\beta < 1$. To reproduce the home equity that we observe in our 2016 SCF sample, we calibrate the long-run discount rate ρ in both cases to match an estimated average LTV ratio of 0.51. To capture the fragility of household balance sheets, we also calibrate β in the Present-Bias Benchmark to fit an estimated average credit card debt to permanent income ratio of 0.09 (the Exponential Benchmark restricts $\beta = 1$). Note that SCF-measured credit card debt is designed to capture *revolving* debt that accrues interest. Moreover, to exclude balances on promotional low-interest products such as balance-transfer cards, we follow Lee and Maxted (2023) and restrict our focus to *high-interest credit card debt* on which households report paying an interest rate of greater than 5%. Finally, because credit card debt appears to be underreported in the SCF (Zinman, 2009), we adjust reported credit card borrowing upward by a factor of 1.5 following the methodology of Beshears et al. (2018, Appendix C).

Table 1 presents the discount function calibration. In the Exponential Benchmark, $\rho = 1.25\%$ matches the average LTV moment of 0.51. However, the Exponential Benchmark fails to generate the level of credit card borrowing observed in the data. In the Present-Bias Benchmark, $\beta = 0.83$ matches the average credit card debt to income ratio of 0.09,[32] while lowering ρ to 0.88% ensures that households still accumulate sufficient home equity to match the LTV moment. Importantly, Table 1 highlights that in a calibrated model, introducing present bias does not simply mean that households are more impatient. Rather, present-biased households are more impatient in the short run since $\beta < 1$, but are simultaneously more patient in the long run due to their lower ρ. This differential patience allows the Present-Bias Benchmark to fit high-cost credit card borrowing jointly with illiquid home equity accumulation (relatedly, see also Laibson et al., 2023).

For readers who are uncomfortable with our 1.5-times adjustment to SCF credit card borrowing, we also estimate that 53% of households in our SCF sample report carrying a high-interest credit card balance from one month to the next. Table 1 shows that this feature is qualitatively matched by our Present-Bias Benchmark, but not by our Exponential Benchmark. Indeed, we would calibrate an even lower β value if we were to set β to reproduce the share of households with revolving (high-interest) credit card debt.

Compared to heterogeneous-agent macro models with exponential time preferences, one key difference with our calibration is that we take both the credit card wedge ($\omega^{cc} = 10.3\%$) and the effective illiquid return spread ($\omega^{m} = 1.7\%$) "from the data." Despite the large

but may not reflect periods where prevailing mortgage rates exceed what homeowners previously locked in (e.g., year 2023). Additionally, we pick $r^* = 1\%$ so that we can study rate cuts from 1% to 0%, while still maintaining the possibility that rates could fall further (to -1%). This preserves the option value to delaying refinancing, thus preventing households from rushing to lock in the lowest possible mortgage rate.

[32] Meier and Sprenger (2010) combine surveyed time preferences with administrative credit card data to provide direct evidence of the relationship between present bias and credit card debt.

	Data	Exponential Benchmark	Present-Bias Benchmark
Discount Function			
β	-	1	0.83
ρ	-	1.25%	0.88%
Calibration Targets			
LTV	0.51	**0.51**	**0.51**
Avg. C.C. Debt	0.09	0.03	**0.09**
Share C.C. Debt > 0	53%	27%	47%

Table 1: Internally Calibrated Parameters.

Notes: This table presents the calibration of the discount function for the two benchmark cases we study.

borrowing wedge and the low illiquidity wedge, our Present-Bias Benchmark is nonetheless able to match the credit card debt levels observed empirically. In contrast, models with exponential time preferences such as Kaplan and Violante (2014) and Kaplan et al. (2018) typically internally calibrate at least one of the model's interest rates in order to generate similar levels of low-liquidity households. Doing so typically results in a lower borrowing wedge and higher illiquidity spread than in the data.[33] This also implies that our own Exponential Benchmark – which sets $\beta = 1$ *without* recalibrating interest rates to generate constrained households – has too few low-liquidity households and hence low MPCs.[34]

Supplement: Intermediate Cases. As shown by Proposition 2, the Present-Bias Benchmark features procrastination while the Exponential Benchmark does not. Because our Present-Bias Benchmark therefore introduces both present bias and procrastination, for conceptual clarity we also discuss various "intermediate cases" in Appendix D.4. These intermediate cases allow us to bridge the gap between the Exponential Benchmark and the Present-Bias Benchmark by studying present bias *without* refinancing inertia, and exponential discounting *with* refinancing inertia. These cases also allow us to explore sophistication, which effectively shows up as present bias without refinancing inertia.

[33]For example, Kaplan et al. (2018) internally calibrate a borrowing wedge of 6% to fit the share of households with negative net liquidity (a borrowing rate of 8% minus a liquid return of 2%), which is lower than our empirical credit card wedge of 10.3%. Kaplan et al. (2018) relatedly calibrate an illiquidity wedge of 3.7% (5.7% minus 2%), which is higher than our number of 1.7%. Additionally, our approach of taking interest rate wedges "from the data" also differentiates our approach from models with unsecured borrowing rates determined endogenously based on default risk (e.g., Mitman, 2016), though such models have also had success in producing realistic levels of unsecured debt and home equity accumulation.

[34]Another difference is that our model includes mortgages, which allows households to borrow against their illiquid (housing) wealth. Alternatively, in models such as Kaplan et al. (2018) that only have net illiquid wealth, households must sell their illiquid wealth in order to access it. Since illiquid assets offer high returns and hence are costly to sell, we conjecture that allowing households to borrow against their illiquid wealth effectively makes that wealth more liquid. This will also lead to fewer constrained households and hence lower MPCs in our model. A similar effect is also reported in McKay and Wieland (2021).

4.3 Steady State Household Behavior

Households solve an optimal control problem augmented with an optimal stopping problem, and the steady state features cross-sectional heterogeneity in four variables: b_t, m_t, y_t, and ε_t. When characterizing the steady state we focus on the features that will be important for the macroeconomic policy results to follow. As we detail below, many of the equilibrium behaviors that differentiate the Present-Bias Benchmark from the Exponential Benchmark are consistent with an emerging set of empirical findings in the household finance literature that have, collectively, proven challenging for models with exponential discounting to replicate.

Phase Diagrams of the Household Balance Sheet. Figure 1 uses phase diagrams to describe the evolution of households over the state space. From left to right, each panel represents a different income state. The top row shows the Exponential Benchmark, and the bottom row shows the Present-Bias Benchmark. The horizontal axis of each panel is liquid wealth b and the vertical axis is the household's mortgage balance, expressed as the LTV ratio m/h. The red and blue shaded areas indicate discrete adjustment regions: red indicates cash-out refinancing, and blue indicates discrete mortgage prepayment.[35] In areas of non-adjustment, the arrows indicate how the household balance sheet evolves over time.

Looking first at the red regions, households choose to take a cash-out refinance when they have relatively low income (either y_L or y_M) and are near the credit limit of $\underline{b}$. During spells of lower income, the first margin on which households adjust is to decumulate liquid wealth. The second margin is credit card debt: even with wedge ω^{cc}, the fixed cost of refinancing implies that temporarily taking on credit card debt can be prudent. Cash-outs are the final margin that households turn to, but only after accumulating sizable credit card debt.

Looking next at the blue regions, households choose to prepay their mortgage once they have built a buffer stock of liquid wealth. Having some liquid wealth is useful because it allows households to avoid taking on costly credit card debt during low-income spells. However, it is suboptimal to hold too much liquidity because the mortgage wedge ω^m implies that mortgage debt is more costly than the household's return on liquid wealth. Thus, high-liquidity households will use some of their accumulated liquidity to pay down their mortgage.

As shown in Proposition 2, differences in adjustment regions across the two calibrations are driven by variation in ρ. The top row (Exponential Benchmark) features a higher calibration of $\rho = 1.25\%$, while the bottom row (Present-Bias Benchmark) calibrates $\rho = 0.88\%$. Households who perceive themselves to be less patient (higher ρ) will be more willing to take a cash-out refinance and less willing to prepay their mortgage. The variation in the red cash-out region and blue prepayment region across the two calibrations follows accordingly.

[35] In the steady state with constant interest rates, the only reason to refinance is to withdraw home equity.

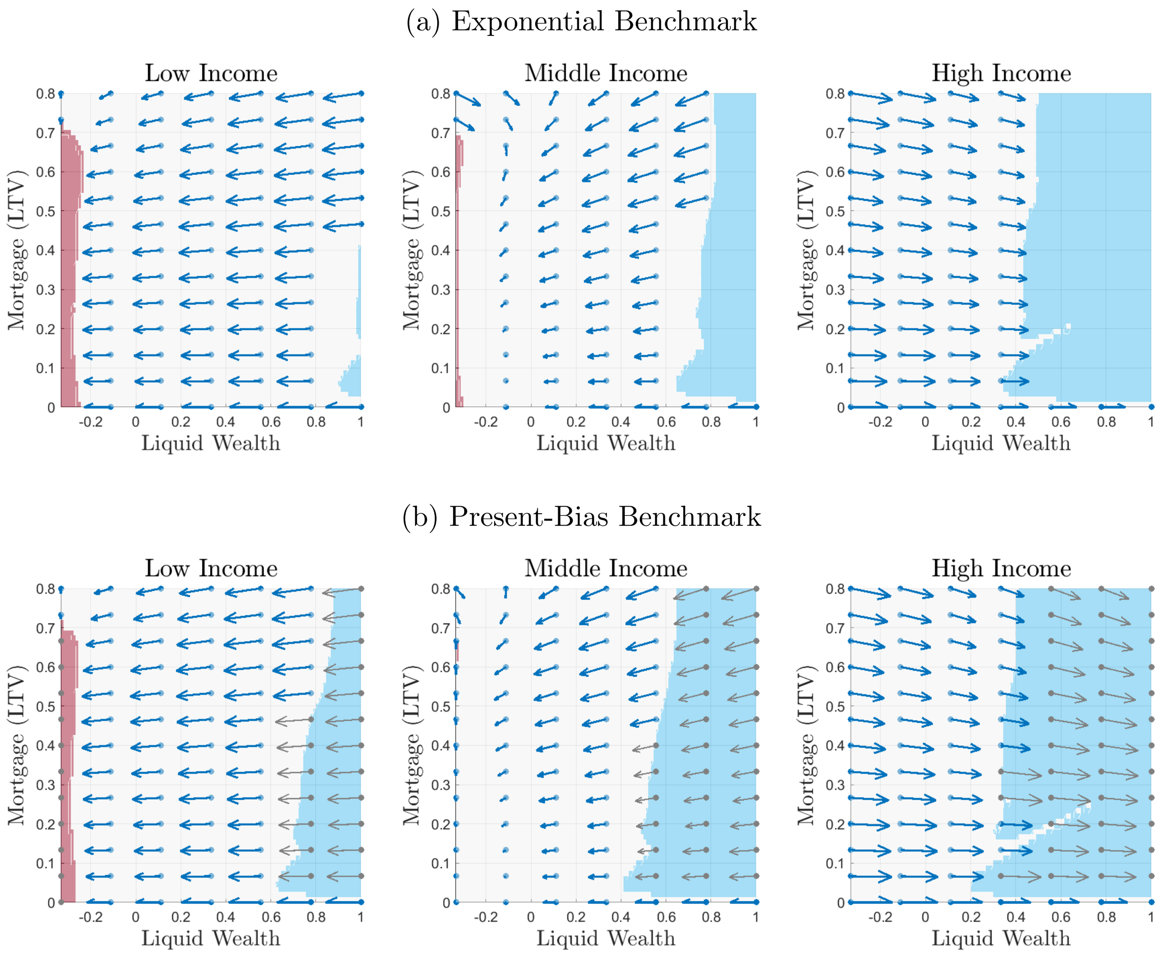

Figure 1: Steady State Phase Diagrams.

Notes: Arrows display the evolution of household balance sheets. Red regions mark where the household chooses to conduct a cash-out refinance, and blue regions mark mortgage prepayment. See text for details.

The blue arrows indicate how household balance sheets evolve over the non-adjustment regions. The arrows always point downward due to mortgage principal repayment, as specified by equation (2). The arrows point either right or left to indicate liquid saving or dissaving, and the length of the arrow corresponds to the speed of evolution. Arrows point strongly leftward when $y_t = y_L$, indicating liquid dissaving. Arrows point strongly rightward when $y_t = y_H$, indicating liquid saving. When $y_t = y_M$ the arrows typically point to the left but are small, indicating slight dissaving by households.

The bottom row features gray arrows in the adjustment regions, while the top row does not. In the top row there is no procrastination. Households will "jump" as soon as they move into an adjustment region, and therefore households will never find themselves in the shaded regions. The bottom row features procrastination. This means that households will find themselves in the adjustment regions, and the gray arrows indicate how household balance sheets evolve when households procrastinate.

Steady State Consumption Dynamics. Figure 2 plots the steady state consumption function. Liquid wealth is on the horizontal axis. Each panel plots the consumption function for an LTV $\in \{0, 0.4, 0.8\}$.

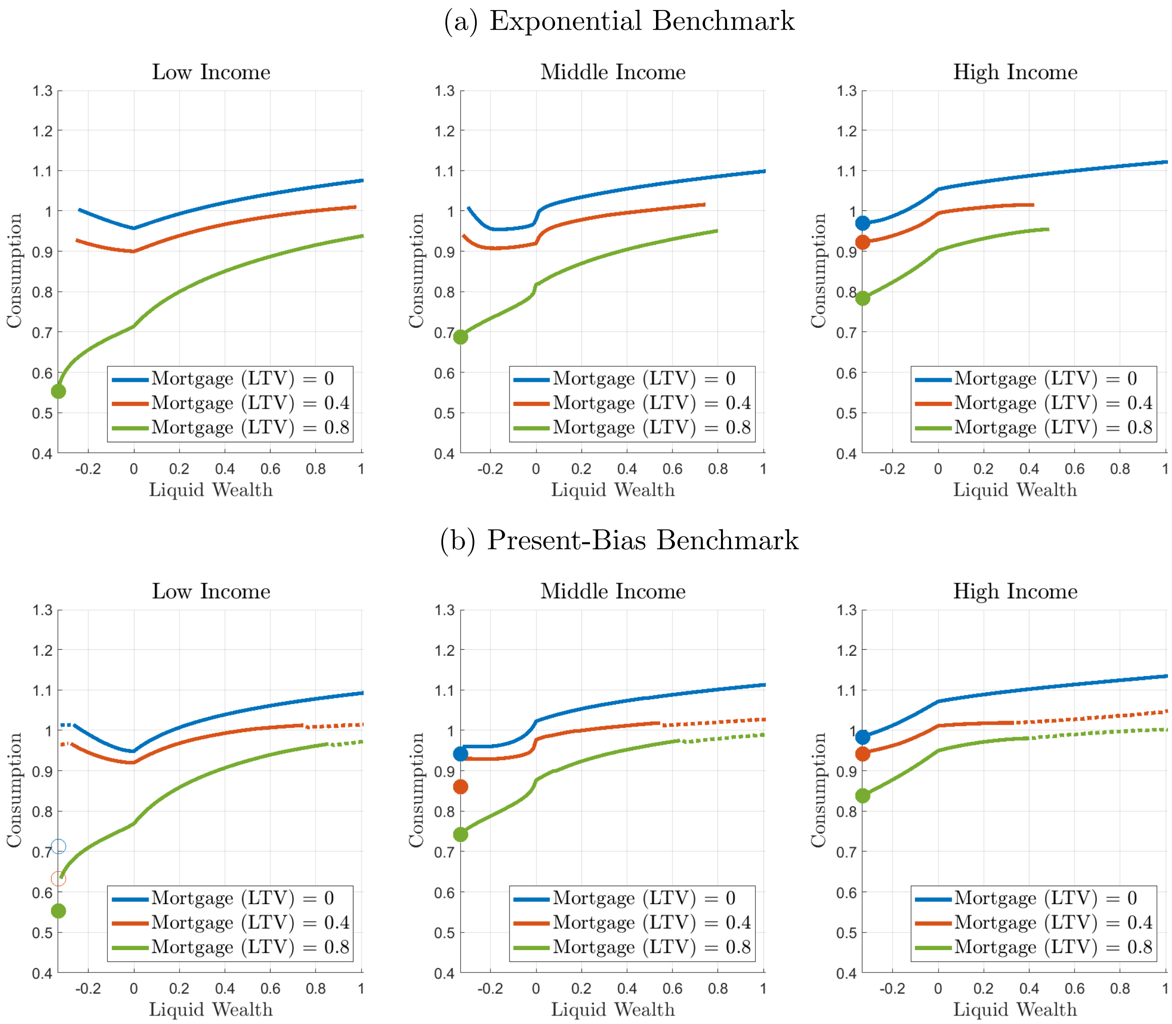

Figure 2: Steady State Consumption.

Notes: This figure shows the steady state consumption function over income and liquid wealth for LTV $\in \{0, 0.4, 0.8\}$. See text for details.

Starting with the Exponential Benchmark in the top row, note that the curves do not always span the entire liquid-wealth axis (e.g., the blue and orange curves do not extend to $\underline{b}$ in the Low Income panel). These missing areas reflect the adjustment regions. Households will never reach these parts of the state space.

The bottom row shows consumption in the Present-Bias Benchmark. There are two key differences under present bias. First, we now use dashed lines and open circles to mark consumption in the adjustment regions, since procrastination implies that households can enter these regions. Second, the consumption function occasionally features a discontinuity at $\underline{b}$. This discontinuity arises because the borrowing constraint restricts overconsumption (see constraint (3)). The consumption discontinuity is consistent with the evidence presented in Ganong and Noel (2019), who use high-frequency data to document that consumption drops sharply at the expiration of unemployment insurance.

This discontinuity can be particularly large when households procrastinate on refinancing. For example, households with $y_t = y_L$ and LTV $= 0.4$ experience a drop in consumption of approximately 30% if they fail to refinance before hitting $\underline{b}$. The reason for this large discontinuity is that households are naive about their procrastination. Once households are in an adjustment region they expect their next self to refinance. As a result, households choose consumption today in order to smooth consumption relative to where they expect to be following a cash-out refinance. Households do not anticipate hitting the borrowing constraint $\underline{b}$, creating a large downward discontinuity if the constraint does eventually bind.

Marginal Propensities to Consume (MPCs). Consumption behavior can also be investigated through MPCs. The marginal propensity to consume (MPC) over τ years is defined as follows (Achdou et al., 2021):

$$MPC_\tau(x) = \frac{\partial}{\partial b}\mathbb{E}\left[\int_0^\tau c(x_t)dt \mid x_0 = x\right]. \tag{13}$$

The left panel of Figure 3 plots the quarterly MPC out of \$1,000 as a function of liquid wealth b, averaging over the income and mortgage dimensions.[36] Consistent with the typical behavior in these sorts of models, the Exponential Benchmark features elevated MPCs at $b = 0$ where the interest rate jumps and at the borrowing constraint $\underline{b}$ (e.g., Kaplan and Violante, 2014; Kaplan et al., 2018). The key effect of present bias is that it drastically increases MPCs near the borrowing constraint. This follows directly from the consumption function discontinuities at $\underline{b}$ detailed above: present-biased households do not smooth consumption into the borrowing constraint, instead choosing a high consumption rate all the way into $\underline{b}$.[37]

Steady State Wealth Distribution. The right panel of Figure 3 presents the steady state distribution of liquid wealth for the Exponential Benchmark and the Present-Bias Benchmark. There are two key differences. First, present bias generates a leftward shift in the

[36] Equation (13) defines the MPC out of an infinitesimal increase in liquid wealth. However, tax rebates and fiscal stimulus payments increase liquid wealth discretely. The definition of the MPC can easily be extended to discrete liquidity shocks (see Appendix E.1 for details).

[37] The left panel of Figure 3 also illustrates that in both calibrations, MPCs are driven by borrowing-constrained households at either $b = 0$ or $b = \underline{b}$ and decline quickly with liquid wealth. Accordingly, present bias alone is not capable of matching the empirical evidence that some high-liquidity households also have high MPCs (see for example Figure 4 in Ganong et al., 2020, and the studies referenced therein).

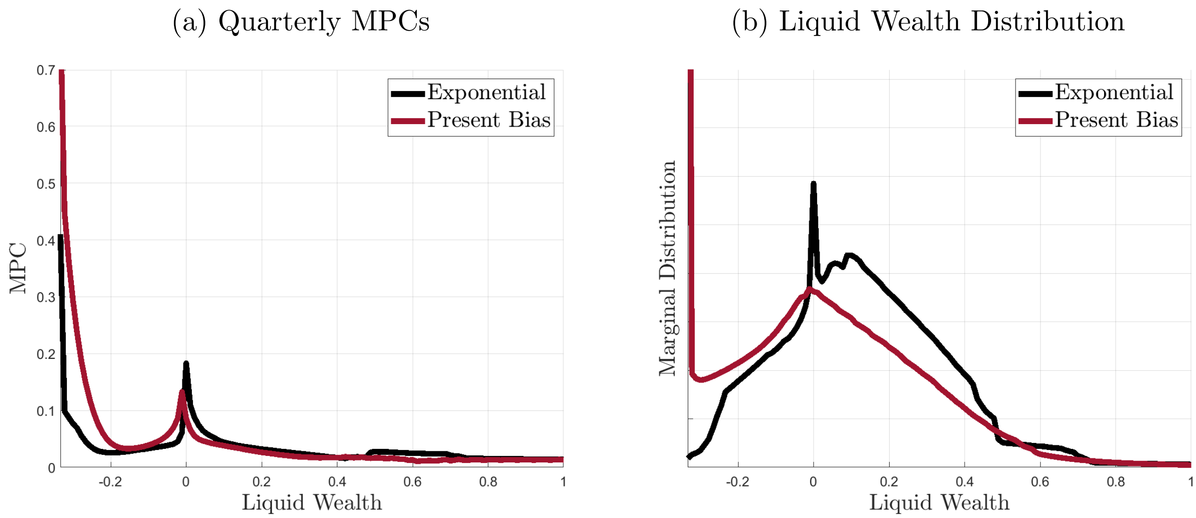

Figure 3: MPCs Over the Liquid Wealth Distribution.

Notes: The left panel presents the quarterly MPC as a function of liquid wealth, averaging over the mortgage and income dimensions. The right panel shows the household distribution over liquid wealth.

liquid wealth distribution because present-biased households overconsume out of liquidity.[38] Second, present bias produces a larger share of households at the borrowing constraint $\underline{b}$.

There are three factors that contribute to this large mass of constrained households at $\underline{b}$ when $\beta < 1$. First, households fail to maintain liquidity buffers. Second, the $\beta < 1$ calibration features a lower value of ρ, which reduces households' willingness to refinance at $\underline{b}$ (see the Middle Income panels in Figure 1). Third, procrastination means that households can be slow to refinance when they hit the constraint.

The buildup of households at $\underline{b}$ is consistent with the responsiveness of debt accumulation to borrowing limits documented empirically. Gross and Souleles (2002) study how borrowing responds to credit limit increases, and estimate that credit card debt increases by 10-14% of the increased limit after one year. In our model, the first-order effect of a small increase in the borrowing limit $\underline{b}$ is that constrained households accumulate additional debt by exactly the same amount, so that the overall marginal borrowing propensity simply equals the share of constrained households. Consistent with this evidence, the share of constrained households is 12% in the Present-Bias Benchmark, compared to only 0.1% in the Exponential Benchmark. Moreover, we find in our 2016 SCF sample of homeowners that roughly 14% of households have less than two-weeks' pay of available credit remaining on their credit card, again consistent with the buildup of households near $\underline{b}$ in the Present-Bias Benchmark.

[38] In particular, the overconsumption shown in Proposition 1 implies that the soft constraint does not typically bind for $\beta < 1$ households. See Maxted (2023, Prop. 7) and Lee and Maxted (2023) for details.

	Exponential	Present Bias
Avg. Quarterly MPC ($1,000)	4.2%	14.4%
(y_L, y_M, y_H)	(4.7, 5.5, 2.0)	(29.3, 12.4, 2.2)
Avg. Quarterly MPC ($10,000)	3.8%	8.3%
Avg. Quarterly MPX ($1,000)	12.7%	30.2%
Avg. Quarterly MPX ($10,000)	11.8%	24.8%
Share $b < 0$	26.7%	47.1%
Share $b = \underline{b}$	0.1%	12.0%

Table 2: Steady State Summary Statistics.

Notes: This table summarizes household consumption, expenditure, and saving behavior in the steady state.

Summary Statistics. Table 2 summarizes the model's steady state. The Present-Bias Benchmark features an average quarterly MPC of 14.4% out of a $1,000 transfer. For the Exponential Benchmark, this MPC is 4.2%. Only the Present-Bias Benchmark comes close to empirical MPC estimates, which range from 15-25% for fiscal transfers of $500 – $1,000.[39] The Present-Bias Benchmark also features much larger variation in MPCs based on transitory income (second row). Low- and middle-income households have high MPCs because they compose a larger share of households on or near the borrowing constraint. Though this MPC heterogeneity still exists in the Exponential Benchmark, it is much less drastic. Time-consistent households maintain liquidity buffers optimally, so the Exponential Benchmark features far fewer constrained households. We also find that the elevated MPC exhibited in the Present-Bias Benchmark is robust to the size of the wealth shock. The average quarterly MPC remains at 8.3% for a transfer of $10,000 (for more, Appendix Figure 12 reports quarterly MPCs out of transfers ranging from $1,000 to $50,000).[40]

When comparing the model to the data, it is important to delineate between the marginal propensity to consume (MPC) versus the marginal propensity for expenditure (MPX) (Laibson et al., 2021).[41] Durables drive a wedge between expenditure and model-based "notional" consumption, because spending on durables does not translate into immediate consumption. The difference can be substantial empirically. For example, Parker et al. (2013) document that over 75% of spending from the 2008 fiscal stimulus was on durable goods.

To bridge the gap between notional consumption and expenditure, we follow the method

[39]See footnote 3 for references reviewing the empirical literature on MPCs. While our model focuses on homeowners, the empirical evidence on large consumption responses to liquidity injections has also been found to extend to homeowners specifically (see e.g. Parker et al., 2013; Lewis et al., 2022).

[40]See Kueng (2018) and Hamilton et al. (2023) for empirical evidence of high MPCs from large transfers. Hamilton et al. (2023) argue that a two-asset model with naive present bias is consistent with the evidence they document whereas the version with exponential discounting is not. Attanasio et al. (2020) show that temptation preferences can also generate high MPCs out of large transfers.

[41]The MPX is also likely to be the more relevant concept for general equilibrium analyses, since the MPX captures changes in overall consumption expenditure on both durables and nondurables which is what matters for aggregate demand.

of Laibson et al. (2021) for converting our model's predictions about MPCs into predictions about MPXs. Table 2 reports the model's MPX predictions. As with MPCs, the Present-Bias Benchmark features elevated quarterly MPXs that remain large even for large wealth transfers. Elevated MPXs out of large transfers is consistent with Fagereng et al. (2021), who use Norwegian administrative data to estimate that lottery winners of amounts ranging from \$8,300 – \$150,000 spend 50% of their prize within the year of winning.

The final two rows of Table 2 summarize credit card borrowing. In the Exponential Benchmark, only 0.1% of households are constrained, and only 27% of households hold credit card debt (compared to 53% in the SCF). The Present-Bias Benchmark is more in line with the data: 12% of households are constrained and 47% hold credit card debt.

5 Results: Macro Stabilization Policy with Present Bias

We now present our results for fiscal and monetary policy under present-biased time preferences. We start all policy counterfactuals from the pre-shock "steady state" with $r^* = 1\%$.

5.1 Fiscal Policy

We first study the efficacy of fiscal policy. Starting from the "steady state" at time $t = 0$, the government (unexpectedly) makes a one-time stimulus payment of \$1,000 to each household. This stimulus payment is financed by a flow income tax levied on all households in perpetuity, which is chosen to satisfy the government budget constraint in our environment with a stochastic interest rate on government debt. Specifically, the government follows a simple fiscal rule: at each point in time $t > 0$, levy a (stochastic) flow income tax on all households that is "just a little bit higher" than the additional (stochastic) interest payments resulting from the initial stimulus. As we explain in Appendix F.1, this simple fiscal rule ensures that the government budget constraint is satisfied and that government debt eventually reverts to its initial steady state level. This simple modeling trick could also prove useful in other environments with stochastic interest rates.

Given that a substantial fraction of households are financially constrained and that households have finite working lives (without transfers across generations), Ricardian equivalence does not hold in our model (Barro, 1974). In fact, the initial short-run consumption response under this tax scheme is similar to that without imposing a government budget constraint. Intuitively, the costs of the initial fiscal stimulus are spread over a long horizon (resulting in a small per-period tax) and borne largely by future generations (Blanchard, 1985).

Figure 4 plots the impulse response function (IRF) of aggregate consumption to this \$1,000 fiscal transfer. To make this IRF easier to connect to MPCs, we scale the consumption response at time t by the size of the initial fiscal transfer. We provide the cumulative (scaled) consumption responses over different time horizons in the corresponding table. The model's predicted consumption responses can also be interpreted using the "intertemporal MPC" framework of Auclert et al. (2018), who show that this dynamic spending response is important for characterizing the general equilibrium propagation of fiscal policy shocks.

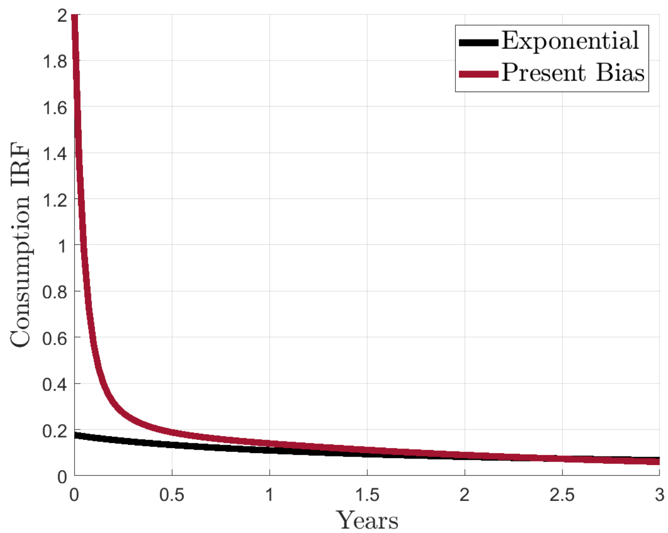

	Exponential	Present Bias
1 Year Consumption Response	14%	28%
2 Year Consumption Response	23%	39%
3 Year Consumption Response	30%	46%

Figure 4: Consumption Response to Fiscal Policy.
Notes: This figure presents the IRF of aggregate consumption to a fiscal stimulus of $1,000. The corresponding table provides the cumulative consumption response.

The results in the figure and table show clearly that the short-run potency of fiscal policy is amplified when $\beta < 1$. The intuition for this result is provided by Figure 3. The $\beta < 1$ calibration features larger MPCs at the borrowing constraint as well as a larger mass of constrained households. This dual result of large MPCs for constrained households combined with a large share of constrained households makes fiscal policy very powerful for $\beta < 1$.[42]

Fiscal Policy Implementation: Liquidity. In response to the 2007-08 Financial Crisis, policymakers utilized a mixture of liquid and illiquid fiscal transfers (e.g., stimulus checks versus mortgage principal reductions). To study these different transfer designs in our model, Appendix Figure 15 compares the consumption response following a liquid fiscal transfer to

[42]Table 2 and Appendix Figure 12 show that MPCs remain elevated even for large transfers, suggesting that present bias also allows fiscal policy to remain potent when implemented at larger scales.

an illiquid mortgage reduction of the same magnitude. In the Present-Bias Benchmark, we find that the large consumption response shown in Figure 4 relies critically on the transfer being liquid. This is because the liquidity-constrained households who drive that consumption response are unable to immediately consume out of illiquid home equity. Such sensitivity to transfer liquidity is consistent with Ganong and Noel (2020), who use variation in mortgage modification programs following the financial crisis to document that liquidity is the critical driver of short-term household consumption decisions. Alternatively, the consumption response in the Exponential Benchmark is far less sensitive to the liquidity of the transfer.

5.2 Monetary Policy

Next we study the impact of stimulative monetary policy. Starting from the "steady state" where $r^* = 1\%$ and all households have a mortgage interest rate of $r_t^m = r^* + \omega^m$, interest rates are cut from 1% to 0% and held at 0% for 3 years.[43] Figure 5 illustrates the consumption response to this 1% rate cut in our Exponential and Present-Bias Benchmarks compared to the "steady state." We also provide the present value of these consumption responses in the corresponding table (discounted at the market interest rate).

In both cases, approximately 70% of households find themselves in a refinancing region following the rate cut. By encouraging households to refinance, an important feature of this interest rate cut is that it also provides households with an opportunity to extract home equity. Indeed, roughly 75% of households who refinance engage in a cash-out refinance. Details of the refinancing decision are presented in Appendix F.3.

Starting with the Exponential Benchmark, Figure 5 shows that this case features a 3.7% increase in consumption on impact, which then decays slowly over longer horizons. Intuitively, the rate cut induces a wave of cash-out refinances on impact, and this extracted home equity is steadily spent down over time. Turning to the Present-Bias Benchmark, the solid red line shows a similar on-impact consumption response to the rate cut. But, in contrast to the Exponential Benchmark, the Present-Bias Benchmark also exhibits essentially no decay in potency over the three-year period that we study.

To provide intuition for the Present-Bias Benchmark, the transparent red line plots a counterfactual consumption response where we start from the Present-Bias Benchmark, but shut off refinancing procrastination for one quarter after the rate cut. Without procrastination, we see that monetary policy is roughly twice as effective on impact, but also burns out quickly over time. The intuition is the same as for fiscal policy: the rate cut incentivizes households – especially those who are constrained – to extract equity from their home. In

[43] Since we only study the three-year response, we do not need to specify the interest rate path beyond that horizon. While our monetary policy experiment examines just one sample path of interest rates for simplicity, households continue to have calibrated interest-rate expectations along that sample path.

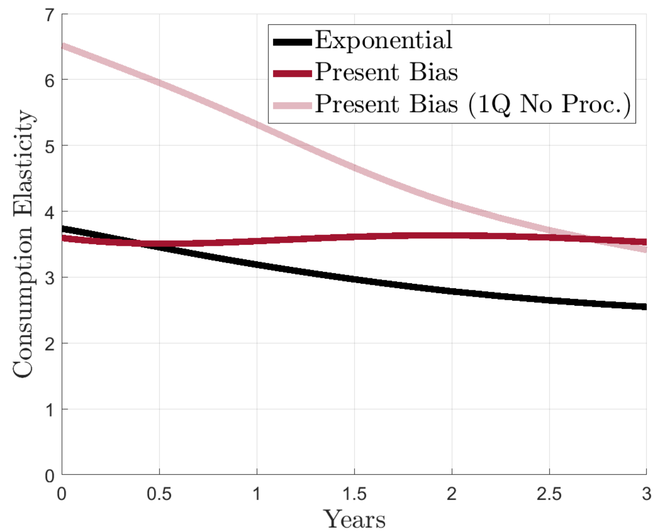

	Exponential	Present Bias
1 Year Present Value	3.45%	3.53%
2 Year Present Value	6.40%	7.14%
3 Year Present Value	9.03%	10.72%

Figure 5: Consumption Response to Monetary Policy.

Notes: This figure plots the percent change in aggregate consumption following an interest rate cut from 1% to 0%. The corresponding table provides the cumulative consumption response in present value terms. The transparent line labeled 1Q No Proc. plots a counterfactual consumption response that starts from the Present-Bias Benchmark, but shuts off refinancing procrastination for one quarter after the rate cut.

this way, *the refinancing channel of monetary policy imitates the liquidity-injection features of fiscal policy.* Since present-biased households overconsume out of liquid wealth, this wave of home-equity extractions produces a large consumption boom.

Connecting this counterfactual experiment to our Present-Bias Benchmark, the key difference is that households procrastinate on refinancing. Though procrastination lowers the consumption response to monetary policy on impact, it also helps to sustain that response over time as households slowly get around to refinancing.

A noteworthy feature of the Present-Bias Benchmark is that the consumption response to monetary policy in Figure 5 is mildly hump-shaped, reaching its peak after around 2 years. Appendix Figure 18 shows that this hump shape can also be more pronounced if the procrastination duration is shorter.[44] A hump-shaped response of aggregate consumption

[44]Intuitively, even short bouts of procrastination will limit the on-impact consumption response, since many households need to extract home equity before they can increase consumption. But when procrastination is only short-lived, a wave of home-equity extractions still follows quickly after the rate cut. Thus, short bouts of procrastination generate a muted on-impact, but large subsequent, response to monetary policy, which then burns out as households consume their cash-out — i.e., a hump-shaped consumption response.

to monetary policy shocks is a common finding in the literature estimating such IRFs using time-series data (e.g., Rotemberg and Woodford, 1997; Christiano et al., 2005). Present bias with procrastination thus has the potential to qualitatively generate this empirical finding.[45]

There is also an emerging literature showing that the refinancing channel of monetary policy is sensitive to the time-path of mortgage interest rates (Berger et al., 2021; Eichenbaum et al., 2022). We highlight a different form of sensitivity — sensitivity to procrastination. Procrastination means that households are slow to adjust on their refinancing margin, prolonging the pass-through of monetary policy to household consumption.

Consumption Response Decomposition: Three Channels. Monetary policy affects household consumption through three channels. First, there is the standard direct effect on liquid wealth — the change in interest rate r_t affects the household's return on b_t.[46] Second, the interest rate cut gives households the option to refinance into a lower-rate mortgage. Third, households can extract home equity when refinancing their mortgage.

We decompose the initial consumption response to monetary policy into its three components. First, we isolate the direct effect on liquid wealth by shutting down households' ability to adjust their mortgage. Second, we reintroduce the ability for households to conduct a rate refinance, but keep the cash-out channel shut down.[47] Third, we reintroduce the cash-out channel to get back to the benchmark results shown in Figure 5.

Table 3 presents this decomposition. Each cell gives the on-impact consumption elasticity in the modified model.[48] In both calibration cases, the liquid wealth channel and the rate refinancing channel each drive roughly one quarter of the total consumption response, with the bulk of the response coming from the cash-out channel of monetary policy.[49]

5.3 Summary

Present bias creates both high MPCs and a large share of constrained households. Because these households are quick to spend down liquidity, present bias makes fiscal policy a powerful and robust tool for generating a short-run consumption boom. Present bias also increases

[45]See also Auclert et al. (2020), who propose "sticky expectations" as a source of such hump-shaped IRFs.

[46]In particular, the direct effect on liquid wealth includes the usual income and substitution effects, which are the focus of most single-asset models of monetary policy.

[47]To do this, we modify the refinancing budget constraint in equation (5) so that it instead becomes: $b' - m' = b_t - m_t - \kappa^{refi}$, subject to $m' \in [0, m_t + \kappa^{refi}]$ and $b' \geq \underline{b}$. This modified budget constraint means that households cannot increase their liquid wealth by refinancing; i.e., $b' \leq b_t$.

[48]In all three steps we use the pre-shock distribution of households from the full model. This prevents the distribution of households from changing as we change households' access to refinancing technology.

[49]One important caveat here is that our model provides an incomplete picture of the spending response to cash-outs. In particular, while households often report using extracted home equity for residential investment (e.g., Greenspan and Kennedy, 2008), this channel is broadly missing from our model with *fixed* housing h.

	Exponential	Present Bias
Step 1. No Adjustment	0.81%	0.84%
Step 2. No Cash-Outs	1.78%	1.89%
Step 3. Full Response	3.74%	3.60%

Table 3: Consumption Response Decomposition.
Notes: This table decomposes the channels through which monetary policy produces a consumption response (on impact). The first row presents the consumption elasticity when households are not allowed to adjust their mortgage. The second row allows for rate refinances only. The third row presents the full response.

the potency of monetary policy for a similar reason: rate cuts produce cash-out refinances, which mirror the liquidity-injection features of fiscal policy. Though powerful, the refinancing channel of monetary policy is sensitive to refinancing inertia. Procrastination increases the lag between rate cuts and refinancing, generating a milder but longer-lived stimulus.

5.4 Extensions

Distributional Effects of Fiscal and Monetary Policy. Appendix D.1 leverages the heterogeneous-agent structure of the model to explore how present bias affects the distributional consequences of policy.[50] The key takeaway from this analysis is that present bias reverses the distributional consequences of fiscal versus monetary policy. In the Exponential Benchmark, monetary policy is an effective way to stimulate the consumption of low-consumption households: a cut to interest rates allows low-consumption households to refinance. In the Present-Bias Benchmark, procrastination implies that monetary policy no longer stimulates the short-run consumption of constrained households. Instead, fiscal policy is highly effective at increasing the short-run consumption of low-consumption households.

Shocks to House Prices and Income. Appendix D.2 examines the effect of house price shocks and recessionary income shocks on our results. House price shocks are of particular interest for our analysis, since house price shocks can quickly create or destroy the home equity that is central to the refinancing channel of monetary policy (e.g., Beraja et al., 2019). Though the magnitude of the consumption response to monetary policy is sensitive to house price shocks, our main result that present bias amplifies this response continues to hold.

A Call to ARMs? Thus far we have assumed that households have fixed-rate mortgages (FRMs) in order to reflect the typical features of the U.S. mortgage market. However, adjustable-rate mortgages (ARMs) are often the modal mortgage contract outside of the U.S. (Badarinza et al., 2016). Moreover, since the 2007-08 Financial Crisis, many economists

[50]See also Wolf (2021) for a comparison of the distributional effects of monetary and fiscal policy.

have argued that downwardly flexible mortgages may improve macroeconomic stability by creating a direct transmission of rate cuts to household mortgage payments (e.g., Eberly and Krishnamurthy, 2014; Andersen et al., 2020; Campbell et al., 2021; Guren et al., 2021).

In Appendix D.3 we explore how present bias interacts with ARMs, and highlight a novel tradeoff between FRMs and ARMs that policymakers should be aware of when considering different mortgage contract designs. On the one hand, the benefit of ARMs is that the direct pass-through of monetary policy offsets refinancing procrastination. On the other hand, ARMs also reduce the cash-out channel of monetary policy – which is particularly potent when households are present biased – because ARMs imply that households no longer need to refinance when rates fall. Though our model is too stylized to make rigorous quantitative claims, we find that these two effects roughly offset in our benchmark calibration.

Discussion of General Equilibrium. Appendix G discusses how present bias could affect the transmission of monetary and fiscal policy in a full general equilibrium analysis. We provide only a brief discussion through the lens of the literature on Heterogeneous Agent New Keynesian (HANK) models. Fully evaluating the impact of present bias in a general equilibrium model is an important task for future work.

6 Conclusion

This paper's main messages are twofold. First, present bias improves the model's ability to replicate a variety of empirical patterns exhibited in household consumption-saving behavior. Second, present bias amplifies the balance-sheet channels of both fiscal and monetary policy but, at the same time, slows down the transmission of monetary policy due to refinancing procrastination.

We conclude by repeating a number of limitations of our analysis. First, we do not model general equilibrium forces and touch upon this issue only briefly in Appendix G. Second, our model abstracts from many important macroeconomic dimensions. We focus on a subset of the population, homeowners. We do not model the endogenous responses of the financial sector nor do we model businesses, both of which are affected by macroeconomic stabilization policy. Third, even in partial equilibrium, the household side of our model is highly stylized. Fourth, our discussion on the timing of fiscal and monetary policy abstracts from policy lags which, in practice, are a critical difference between the speed of fiscal versus monetary policy. Fifth, we do not study the welfare consequences of fiscal and monetary policy. All of these considerations are likely fruitful areas for future research.

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*, 2021.

Agarwal, Sumit, Itzhak Ben-David, and Vincent Yao, "Systematic Mistakes in the Mortgage Market and Lack of Financial Sophistication," *Journal of Financial Economics*, 2017, *123* (1), 42–58.

_ , John C. Driscoll, and David I. Laibson, "Optimal Mortgage Refinancing: A Closed-Form Solution," *Journal of Money, Credit and Banking*, 2013, *45* (4), 591–622.

_ , Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel, "Do Banks Pass Through Credit Expansions to Consumers Who Want to Borrow?," *The Quarterly Journal of Economics*, 2018, *133* (1), 129–190.

Akerlof, George A., "Procrastination and Obedience," *American Economic Review*, 1991, *81* (2), 1–19.

Allcott, Hunt, Joshua J. 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.

Alves, Felipe, Greg Kaplan, Benjamin Moll, and Giovanni L. Violante, "A Further Look at the Propagation of Monetary Policy Shocks in HANK," *Journal of Money, Credit and Banking*, 2020, *52* (S2), 521–559.

Andersen, Steffen, John Y. Campbell, Kasper Meisner Nielsen, and Tarun Ramadorai, "Sources of Inaction in Household Finance: Evidence from the Danish Mortgage Market," *American Economic Review*, October 2020, *110* (10), 3184–3230.

Angeletos, George-Marios, David Laibson, Andrea Repetto, Jeremy Tobacman, and Stephen Weinberg, "The Hyberbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation," *Journal of Economic Perspectives*, 2001, *15* (3), 47–68.

Ashraf, Nava, Dean Karlan, and Wesley Yin, "Tying Odysseus to the Mast: Evidence from a Commitment Savings Product in the Philippines," *The Quarterly Journal of Economics*, 2006, *121* (2), 635–672.

Attanasio, Orazio, Agnes Kovacs, and Patrick Moran, "Temptation and Commitment: Understanding Hand-to-Mouth Behavior," *NBER w27944*, 2020.

Auclert, Adrien, "Monetary Policy and the Redistribution Channel," *American Economic Review*, 2019, *109* (6), 2333–67.

_ , Matthew Rognlie, and Ludwig Straub, "The Intertemporal Keynesian Cross," *NBER w25020*, 2018.

_ , _ , and _ , "Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model," *NBER w26647*, 2020.

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.

_ **, Muriel Niederle, and Charles Sprenger**, "Working Over Time: Dynamic Inconsistency in Real Effort Tasks," *The Quarterly Journal of Economics*, 2015, *130* (3), 1067–1115.

Badarinza, Cristian, John Y. Campbell, and Tarun Ramadorai, "International Comparative Household Finance," *Annual Review of Economics*, 2016, *8*, 111–144.

Bardi, Martino and Italo Capuzzo-Dolcetta, *Optimal Control and Viscosity Solutions of Hamilton-Jacobi-Bellman Equations*, Springer, 1997.

Barro, Robert J., "Are Government Bonds Net Wealth?," *Journal of Political Economy*, 1974, *82* (6), 1095–1117.

_ , "Ramsey Meets Laibson in the Neoclassical Growth Model," *The Quarterly Journal of Economics*, 1999, *114* (4), 1125–1152.

Bensoussan, Alain and Jacques-Louis Lions, *Applications of Variational Inequalities in Stochastic Control*, Vol. 12 of *Studies in Mathematics and its Applications*, North-Holland Publishing Co., Amsterdam-New York, 1982.

_ **and** _ , *Impulse Control and Quasivariational Inequalities* μ, Gauthier-Villars, Montrouge; Heyden & Son, Inc., Philadelphia, PA, 1984.

Beraja, Martin, Andreas Fuster, Erik Hurst, and Joseph Vavra, "Regional Heterogeneity and Monetary Policy," *The Quarterly Journal of Economics*, 2019, *134* (1), 109–183.

Berger, David, Veronica Guerrieri, Guido Lorenzoni, and Joseph Vavra, "House Prices and Consumer Spending," *The Review of Economic Studies*, 2018, *85* (3), 1502–1542.

Berger, David W., Konstantin Milbradt, Fabrice Tourre, and Joseph Vavra, "Mortgage Prepayment and Path-Dependent Effects of Monetary Policy," *American Economic Review*, 2021, *111* (9), 2829–78.

Bernanke, Ben S. and Kenneth N. Kuttner, "What Explains the Stock Market's Reaction to Federal Reserve Policy?," *The Journal of Finance*, 2005, *60* (3), 1221–1257.

Beshears, John, James J. Choi, David Laibson, and Brigitte C. Madrian, "Behavioral Household Finance," in "Handbook of Behavioral Economics: Applications and Foundations 1," Vol. 1, Elsevier, 2018, pp. 177–276.

Bhutta, Neil and Benjamin J. Keys, "Interest Rates and Equity Extraction During the Housing Boom," *American Economic Review*, 2016, *106* (7), 1742–74.

Blanchard, Olivier, *Fiscal Policy under Low Interest Rates*, MIT Press, 2023.

Blanchard, Olivier J., "Debt, Deficits, and Finite Horizons," *Journal of Political Economy*, 1985, *93* (2), 223–247.

Bohn, Henning, "The Behavior of U. S. Public Debt and Deficits," *The Quarterly Journal of Economics*, 1998, *113* (3), 949–963.

Campbell, John Y., "Household Finance," *The Journal of Finance*, 2006, *61* (4), 1553–1604.

_ , "Mortgage Market Design," *Review of Finance*, 2013, *17* (1), 1–33.

_ **, Nuno Clara, and Joao F. Cocco**, "Structuring Mortgages for Macroeconomic Stability," *The Journal of Finance*, 2021, *76* (5), 2525–2576.

Carroll, Gabriel D., James J. Choi, David Laibson, Brigitte C. Madrian, and Andrew Metrick, "Optimal Defaults and Active Decisions," *The Quarterly Journal of Economics*, 2009, *124* (4), 1639–1674.

Case, Karl and Robert Shiller, "The Efficiency of the Market for Single-Family Homes," *American Economic Review*, 1989, *79* (1), 125–37.

Chen, Hui, Michael Michaux, and Nikolai Roussanov, "Houses as ATMs: Mortgage Refinancing and Macroeconomic Uncertainty," *The Journal of Finance*, 2020, *75* (1), 323–375.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 2005, *113* (1), 1–45.

Cloyne, James, Clodomiro Ferreira, and Paolo Surico, "Monetary Policy when Households have Debt: New Evidence on the Transmission Mechanism," *The Review of Economic Studies*, 2020, *87* (1), 102–129.

Cloyne, James S. and Paolo Surico, "Household Debt and the Dynamic Effects of Income Tax Changes," *The Review of Economic Studies*, 2017, *84* (1), 45–81.

Cohen, Jonathan, Keith Marzilli Ericson, David Laibson, and John Myles White, "Measuring Time Preferences," *Journal of Economic Literature*, 2020, *58* (2), 299–347.

DellaVigna, Stefano and Ulrike Malmendier, "Contract Design and Self-Control: Theory and Evidence," *The Quarterly Journal of Economics*, 2004, *119* (2), 353–402.

Dempsey, Kyle and Felicia Ionescu, "Borrowing Premia in Unsecured Credit Markets," *Mimeo*, 2023.

Eberly, Janice and Arvind Krishnamurthy, "Efficient Credit Policies in a Housing Debt Crisis," *Brookings Papers on Economic Activity*, 2014, (2), 73–136.

Eichenbaum, Martin, Sergio Rebelo, and Arlene Wong, "State-Dependent Effects of Monetary Policy: the Refinancing Channel," *American Economic Review*, 2022, *112* (3), 721–761.

Fagereng, Andreas, Martin Blomhoff Holm, and Gisle James Natvik, "MPC Heterogeneity and Household Balance Sheets," *American Economic Journal: Macroeconomics*, 2021, *13* (4), 1–54.

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.

Fulford, Scott L., "The Surprisingly Low Importance of Income Uncertainty for Precaution," *European Economic Review*, 2015, *79*, 151–171.

Gabaix, Xavier, "Behavioral Inattention," in "Handbook of Behavioral Economics: Applications and Foundations 2," Vol. 2, Elsevier, 2019, pp. 261–343.

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

_ and _ , "Liquidity versus Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession," *American Economic Review*, 2020, *110* (10), 3100–3138.

_ , Damon Jones, Pascal J. Noel, Fiona E. Greig, Diana Farrell, and Chris Wheat, "Wealth, Race, and Consumption Smoothing of Typical Income Shocks," *NBER w27552*, 2020.

Gertler, Mark and Peter Karadi, "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, 2015, *7* (1), 44–76.

Gilchrist, Simon, David López-Salido, and Egon Zakrajšek, "Monetary Policy and Real Borrowing Costs at the Zero Lower Bound," *American Economic Journal: Macroeconomics*, 2015, *7* (1), 77–109.

Gornemann, Nils, Keith Kuester, and Makoto Nakajima, "Doves for the Rich, Hawks for the Poor? Distributional Consequences of Monetary Policy," *CEPR Discussion Papers 11233*, 2016.

Greenspan, Alan and James Kennedy, "Sources and Uses of Equity Extracted from Homes," *Oxford Review of Economic Policy*, 2008, *24* (1), 120–144.

Greenwald, Daniel, "The Mortgage Credit Channel of Macroeconomic Transmission," *SSRN 2735491*, 2018.

Gross, David B. and Nicholas S. Souleles, "Do Liquidity Constraints and Interest Rates Matter for Consumer Behavior? Evidence from Credit Card Data," *The Quarterly Journal of Economics*, 2002, *117* (1), 149–185.

Guerrieri, Veronica and Guido Lorenzoni, "Credit Crises, Precautionary Savings, and the Liquidity Trap," *The Quarterly Journal of Economics*, 2017, *132* (3), 1427–1467.

_ , _ , and Marta Prato, "Slow Household Deleveraging," *Journal of the European Economic Association*, 2020, *18* (6), 2755–2775.

Guren, Adam M., Arvind Krishnamurthy, and Timothy J. McQuade, "Mortgage Design in an Equilibrium Model of the Housing Market," *The Journal of Finance*, 2021, *76* (1), 113–168.

Gürkaynak, Refet S., Brian Sack, and Eric Swanson, "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements," *International Journal of Central Banking*, May 2005, *1* (1).

Hamilton, Steven, Geoffrey Liu, Jorge Miranda-Pinto, and Tristram Sainsbury, "Early Pension Withdrawal as Stimulus," *SSRN 4389699*, 2023.

Harris, Christopher and David Laibson, "Dynamic Choices of Hyperbolic Consumers," *Econometrica*, 2001, *69* (4), 935–957.

_ **and** _ , "Instantaneous Gratification," *The Quarterly Journal of Economics*, 2013, *128* (1), 205–248.

Hurst, Erik and Frank Stafford, "Home is Where the Equity Is: Mortgage Refinancing and Household Consumption," *Journal of Money, Credit and Banking*, 2004, *36* (6), 985–1014.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles, "Household Expenditure and the Income Tax Rebates of 2001," *American Economic Review*, 2006, *96* (5), 1589–1610.

Johnson, Eric J., Stephan Meier, and Olivier Toubia, "What's the Catch? Suspicion of Bank Motives and Sluggish Refinancing," *The Review of Financial Studies*, 2019, *32* (2), 467–495.

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor, "Betting the House," *Journal of International Economics*, 2015, *96* (S1), 2–18.

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.

_ **, Kurt Mitman, and Giovanni L. Violante**, "The Housing Boom and Bust: Model Meets Evidence," *Journal of Political Economy*, 2020, *128* (9), 3285–3345.

Kennickell, Arthur and Annamaria Lusardi, "Disentangling the Importance of the Precautionary Saving Motive," *NBER w10888*, 2004.

Kennickell, Arthur B., "Saving and Permanent Income: Evidence from the 1992 SCF," *Mimeo*, 1995.

Keys, Benjamin J., Devin G. Pope, and Jaren C. Pope, "Failure to Refinance," *Journal of Financial Economics*, 2016, *122* (3), 482–499.

Khandani, Amir E., Andrew W. Lo, and Robert C. Merton, "Systemic Risk and the Refinancing Ratchet Effect," *Journal of Financial Economics*, 2013, *108* (1), 29–45.

Kinnerud, Karin, "The Effects of Monetary Policy Through Housing and Mortgage Choices on Aggregate Demand," *Mimeo*, 2021.

Kueng, Lorenz, "Excess Sensitivity of High-Income Consumers," *The Quarterly Journal of Economics*, 2018, *133* (4), 1693–1751.

Laibson, David, "Golden Eggs and Hyperbolic Discounting," *The Quarterly Journal of Economics*, 1997, *112* (2), 443–478.

_ and Peter Maxted, "The Beta-Delta-DELTA Sweetspot," *Mimeo*, 2022.

_ , _ , and Benjamin Moll, "A Simple Mapping from MPCs to MPXs," *Mimeo*, 2021.

_ , Sean Chanwook Lee, Peter Maxted, Andrea Repetto, and Jeremy Tobacman, "Estimating Discount Functions with Consumption Choices Over the Lifecycle," *NBER w13314*, 2023.

Lee, Sean Chanwook and Peter Maxted, "Credit Card Borrowing in Heterogeneous-Agent Models: Reconciling Theory and Data," *Mimeo*, 2023.

Lewis, Daniel, Davide Melcangi, and Laura Pilossoph, "Latent Heterogeneity in the Marginal Propensity to Consume," *Mimeo*, 2022.

Lian, Chen, "Mistakes in Future Consumption, High MPCs Now," *Mimeo*, 2021.

Luttmer, Erzo G.J. and Thomas Mariotti, "Subjective Discounting in an Exchange Economy," *Journal of Political Economy*, 2003, *111* (5), 959–989.

Maggio, Marco Di, Amir Kermani, and Kaveh Majlesi, "Stock Market Returns and Consumption," *The Journal of Finance*, 2020, *75* (6), 3175–3219.

Mankiw, N. Gregory and Ricardo Reis, "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *The Quarterly Journal of Economics*, 2002, *117* (4), 1295–1328.

Maxted, Peter, "Present Bias Unconstrained: Consumption, Welfare, and the Present-Bias Dilemma," *Mimeo*, 2023.

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.

Mian, Atif and Amir Sufi, "House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis," *American Economic Review*, 2011, *101* (5), 2132–56.

Mitman, Kurt, "Macroeconomic Effects of Bankruptcy and Foreclosure Policies," *American Economic Review*, 2016, *106* (8), 2219–2255.

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.

Olivi, Alan, "Sufficient Statistics for Heterogeneous Agent Models," *Mimeo*, 2017.

Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland, "Consumer Spending and the Economic Stimulus Payments of 2008," *American Economic Review*, 2013, *103* (6), 2530–53.

Phelps, Edmund S. and Robert A. Pollak, "On Second-Best National Saving and Game-Equilibrium Growth," *The Review of Economic Studies*, 1968, *35* (2), 185–199.

Reis, Ricardo, "The Constraint on Public Debt when $r < g$ but $g < m$," *Mimeo*, 2021.

Rotemberg, Julio J. and Michael Woodford, "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy," *NBER Macroeconomics Annual*, 1997, *12*, 297–346.

Sims, Christopher A., "Implications of Rational Inattention," *Journal of Monetary Economics*, 2003, *50* (3), 665–690.

Slacalek, Jiri, Oreste Tristani, and Giovanni L. Violante, "Household Balance Sheet Channels of Monetary Policy: A Back of the Envelope Calculation for the Euro Area," *Journal of Economic Dynamics and Control*, 2020, *115*, 103879.

Strotz, Robert Henry, "Myopia and Inconsistency in Dynamic Utility Maximization," *The Review of Economic Studies*, 1956, *23* (3), 165–180.

Werning, Iván, "Incomplete Markets and Aggregate Demand," *NBER w21448*, 2015.

Wolf, Christian, "The Missing Intercept: A Demand Equivalence Approach," *Mimeo*, 2019.

_ , "Interest Rate Cuts vs. Stimulus Payments: An Equivalence Result," *Mimeo*, 2021.

Wong, Arlene, "Refinancing and the Transmission of Monetary Policy to Consumption," *Mimeo*, 2019.

Woodford, Michael, "Imperfect Common Knowledge and the Effects of Monetary Policy," *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, 2003, pp. 25–58.

Zinman, Jonathan, "Where is the Missing Credit Card Debt? Clues and Implications," *Review of Income and Wealth*, 2009, *55* (2), 249–265.

_ , "Household Debt: Facts, Puzzles, Theories, and Policies," *Annual Review of Economics*, 2015, *7* (1), 251–276.

Appendix

A Calibration Details

A.1 Summary of External Calibration

	Description	Value	Target / Source
Preferences			
$\frac{1}{\gamma}$	Intertemporal Elasticity of Substitution	$\frac{1}{2}$	Literature
ϕ	Procrastination Decay Rate	$-\log(0.5)$	Andersen et al. (2020)
Income			
y_t	Transitory Income	$\{0.75, 0.98, 1.28\}$	Guerrieri and Lorenzoni (2017)
$\mathbf{A}^y$	Income Transition Matrix	(see text)	Guerrieri and Lorenzoni (2017)
Interest Rates			
r_t	Interest Rate	$\{-1\%, 0\%, 1\%, 2\%\}$	10-Year TIPS
$\mathbf{A}^r$	Interest Rate Transition Matrix	(see text)	10-Year TIPS
ω^{cc}	Credit Card Wedge	10.3%	Credit Card - 10-Yr Treasury Spread
ω^m	Mortgage Wedge	1.7%	30-Yr FRM - 10-Yr Treasury Spread
Assets and Liabilities			
h	House Value	3.29	2016 SCF
θ	Max LTV	0.8	Greenwald (2018)
ξ	Mortgage Paydown	0.035	20 Year Half-Life
κ^{prepay}	Prepayment Fixed Cost	0.002	Numerical Stability
κ^{refi}	Refinancing Fixed Cost	0.05	FRB Documentation
$\underline{b}$	Credit Limit	$-\frac{1}{3}$	2016 SCF
Other Structural Assumptions			
λ^F	Rate of Forced Adjustment	$\frac{1}{15}$	2016 CPS Avg. Moving Rate
λ^R	Retirement Rate	$\frac{1}{30}$	Average Working Life
y^R	Retirement Fixed Income	y_L	Retirement Replacement Rate
-	New-Homeowner Distribution	$m_0 = \theta h, b_0 \sim U(0, \frac{y_L}{2})$	Lifecycle Dynamics

Table 4: Externally Calibrated Parameters.

Notes: This table presents the model's externally calibrated parameters. See Section 4.1 for details.

A.2 SCF Details

Many of our calibrated parameters rely on data from the 2016 SCF. To construct a sample of households that is consistent with our model we impose the following data filters. The household must own a home, possess a credit card, have a head or spouse in the labor force, and have a head aged 25-61. In order to limit both measurement error and extreme heterogeneity in home values and income, we also restrict our analysis to households with after-tax permanent income between the 1st and 99th percentile, and a home value to permanent income ratio that is below the 95th percentile. Our sample is broadly representative of working-age homeowners, and captures 74% of homeowners aged 25-61.[51]

All of our variables are scaled relative to permanent income. Following Kennickell (1995), Kennickell and Lusardi (2004), and Fulford (2015) we use the SCF's question about "normal income" to measure each household's permanent income.[52] Though this is an imperfect proxy for the household's permanent income, it has the benefit of being both straightforward and respecting the household's information set. We adjust each household's normal income for 2015 federal taxes, and deduct an additional 5% for state taxes.

We use the 2016 SCF to estimate six moments that are used in our calibration: (i) permanent income; (ii) average home value to permanent income; (iii) average LTV; (iv) average credit card debt to permanent income; (v) share of households with revolving credit card debt; and (vi) average credit limit to permanent income. Moments (ii) – (v) are reported in the main text. The average after-tax permanent income for our sample of homeowners is \$95,718. The average credit limit to permanent income is 0.35.

A.3 Calibration of Income and Interest-Rate Processes

To calibrate our income and interest rate processes, we assume that these processes are discretized versions of continuous-time Ornstein-Uhlenbeck (OU) processes.

Using Discrete-Time Estimates to Calibrate Continuous-Time Process. Consider a generic mean-zero OU process $u(t) = \int_0^t e^{-\kappa(t-s)}\sigma dZ_s$. Process $u(t)$ has the conditional distribution $u(t+\tau)|u(t) = \mathcal{N}\left(u(t)e^{-\kappa\tau}, \frac{\sigma^2}{2\kappa}(1-e^{-2\kappa\tau})\right)$.

Assume that $u(t)$ is only observed in snapshots every Δ years. Let $d_s = u(s\Delta)$ denote the s'th snapshot of process $u(t)$. The discrete process d_s can be modeled as an AR(1) process:

$$\begin{aligned} d_{s+1} &= \rho d_s + \sigma_d \varepsilon_{s+1}, \text{ where} \\ \rho &= e^{-\kappa\Delta} \\ \sigma_d^2 &= \frac{\sigma^2}{2\kappa}(1-e^{-2\kappa\Delta}). \end{aligned}$$

Given any discrete-time AR(1) estimate, we can use the above formulas to back out the parameters of the underlying OU process: κ and σ. We discretize the OU process using finite difference methods. For details, see the Numerical Appendix of Achdou et al. (2021).

[51]Working-age homeowners represent 61% of all households aged 25-61 in the 2016 SCF.

[52]SCF respondents are asked whether or not their 2015 income was normal. If not, they are asked to report what their total income would be if it had been normal.

Implementation for Income Process. As already stated in the main text our income process is calibrated following Guerrieri and Lorenzoni (2017) who in turn use data from Floden and Lindé (2001). Specifically, Guerrieri and Lorenzoni assume that the logarithm of income follows an AR(1) process at a quarterly frequency, and calibrate this process with persistence of $\rho = 0.967$ and variance of $\sigma_d^2 = 0.017$. From the formulas we just discussed, we get the parameters for our OU process of $\kappa = 0.134$ and $\sigma = 0.265$.

We set the three income states $y_t \in \{y_L, y_M, y_H\}$ as follows. We set the low and high income states equal to -1 and $+1$ annualized standard deviations of the log income process, more precisely $y_L = y_M e^{-\sigma}$ and $y_H = y_M e^{\sigma}$ where $\sigma = 0.265$. We then set the middle income state y_M to normalize mean income to 1 which yields $y_M = 0.98$ and hence $\{y_L, y_M, y_H\} = \{0.75, 0.98, 1.28\}$. In the stationary distribution of the income process, 31% of households are low income, 39% are middle income, and 31% are high income. The expected persistence of the low, middle, and high income states are 1.6 years, 1 year, and 1.6 years, respectively.

B Naive Present Bias: Passing to Continuous Time

Here we present a heuristic derivation of naive IG preferences as the continuous-time limit of a model where some of the decisions are made discretely. This heuristic approach is designed to capture the intuition of the more rigorous derivation in Harris and Laibson (2013). We begin by assuming a constant effort cost, as in Sections 2.1 and 2.2. The full setup with a stochastic effort cost, as introduced in Section 2.3, is presented in Appendix B.3.

B.1 Naive IG Current-Value Function

Assume that the current self lives for a discrete length of time, denoted Δ. After this time has elapsed, starting with the next self, time progresses continuously again.[53] Since the naive present-biased household incorrectly perceives that all future selves will discount exponentially, continuation-value function $v(x)$ characterizes the equilibrium starting with the next self at time Δ. The current self discounts all future selves by β, so the current-value function for the naive present-biased household is given by:

$$
\begin{aligned}
w(x) &= \max\left\{\max_{c}\ u(c)\Delta + \beta e^{-\rho\Delta}\mathbb{E}[v(x_\Delta)|x],\ w^*(x) - \bar{\varepsilon}\right\} \quad \text{with} \\
w^*(x) &= \max\left\{w^{prepay}(x), w^{refi}(x)\right\} \\
w^{prepay}(x) &= \max_{b',m'} w(b', m', y, r^m, r) \quad \text{s.t. prepayment constraint (4) holds} \\
w^{refi}(x) &= \max_{b',m'} w(b', m', y, r+\omega^m, r) \quad \text{s.t. refinancing constraint (5) holds}
\end{aligned}
\tag{14}
$$

and where x_Δ denotes the vector of household states after time interval Δ has elapsed, for example $b_\Delta = b + (y + rb + \omega^{cc}b^- - (r^m + \xi)m - c)\Delta$.

Equation (14) captures the consumption/adjustment decisions made by the current self. In the left branch of the first line the household does not adjust, and chooses consumption rate c over the next Δ units of time to maximize the current-value function. In the right branch of the first line the household pays effort cost $\bar{\varepsilon}$ and fixed monetary cost κ^i to discretely adjust its mortgage. Importantly, this discrete-time value function is written such that there is no delay to refinancing (i.e., the current self benefits from refinancing).[54] Though this is unrealistic – there are time delays in refinancing – we write the Bellman equation in this

[53] This mixed discrete- and continuous-time setup is of course slightly non-standard. Alternatively, we could have assumed that future selves also make decisions in discrete time, as done in Laibson and Maxted (2022). In this case the continuation-value function $v(x)$ would be the discrete-time analogue of the continuous-time $v(x)$ that we use below.

[54] To see how the value function is written in this way, note that refinancing gives the current self the current-value function of w^*. As the first line of equation (14) shows, this value function consists of an undiscounted utility flow earned for the current self, $u(c)\Delta$.

way to emphasize that our results do not rely on assumptions about temporal delays.

Discrete-time Bellman equation (14) can be used to derive the current-value function in continuous time. Taking the time-step Δ to its continuous-time limit, we see that the term $u(c)\Delta$ drops out of the current-value function, leaving:

$$w(x) = \max \left\{ \beta v(x),\ w^*(x) - \bar{\varepsilon} \right\}.$$

This recovers equation (9) in the main text.

B.2 Continuous Control: Consumption (Proof of Lemma 1)

We now derive the continuous-time first-order condition for consumption stated in Lemma 1. As shown in equation (14), the household makes a consumption choice in every period. For the consumption decision, equation (14) implies that consumption is given by the following first-order condition:[55]

$$u'(c(x)) = \beta e^{-\rho\Delta} \frac{\partial}{\partial b} \mathbb{E}[v(x_\Delta)|x].$$

Taking $\Delta \to 0$ yields

$$u'(c(x)) = \beta \frac{\partial v(x)}{\partial b},$$

which is equation (11) in Lemma 1. This derivation continues to hold in the full setup with a stochastic effort cost presented in Appendix B.3 below.

B.3 Full Setup with a Stochastic Effort Cost (Section 2.3)

Here we briefly spell out the full set of equations for the generalization with a stochastic effort cost that evolves according to the two-state process in Assumption 1. In what follows, we will denote value and policy functions in the normal high-cost state by the same functions as in the baseline model with a constant effort cost, e.g. $v(x)$ or $\mathfrak{R}(x)$. Alternatively, we will denote the corresponding value and policy functions in the temporary low-cost state with underlines, e.g. $\underline{v}(x)$ or $\underline{\mathfrak{R}}(x)$.

We first show how to generalize equation (8′), the HJBQVI equation for the value function

[55] We ignore difficulties such as the kink in the budget constraint when taking this first-order condition.

$v(x)$ of a $\beta = 1$ household:

$$\rho v(x) = \max\Bigg\{ \max_c \Big\{ u(c) + \frac{\partial v(x)}{\partial b}(y + rb + \omega^{cc}b^- - (r^m + \xi)m - c) \Big\} \qquad (8'')$$
$$- \frac{\partial v(x)}{\partial m}(\xi m)$$
$$+ \sum_{y' \neq y} \lambda^{y \to y'} \big[v(b, m, y', r^m, r) - v(b, m, y, r^m, r) \big]$$
$$+ \sum_{r' \neq r} \lambda^{r \to r'} \big[v(b, m, y, r^m, r') - v(b, m, y, r^m, r) \big]$$
$$+ \lambda^R \big[v^R(x) - v(x) \big]$$
$$+ \lambda^F \big[v^*(x) - (v(x) - \bar{\varepsilon}) \big]$$
$$+ \phi \big[\underline{v}(x) - v(x) \big],$$
$$\rho(v^*(x) - \bar{\varepsilon}) \Bigg\}.$$

Relative to ($8'$), there is a new entry $\phi\big[\underline{v}(x) - v(x)\big]$. Parameter ϕ is the arrival rate of the low-effort-cost state, and $\underline{v}(x)$ is the household's value in this state. This value is given by

$$\underline{v}(x) = \max\{v(x), v^*(x) - \underline{\varepsilon}\}. \qquad (15)$$

Intuitively, since the low-cost state only lasts for an instant (Assumption 1), the household either takes advantage of refinancing at the lower effort cost $\underline{\varepsilon}$ or it loses the opportunity in the next instant in which case its value reverts back to $v(x)$.

We next show how to generalize (9), the equation for the current-value function $w(x)$:

$$\begin{aligned}
w(x) &= \max\Big\{ \beta v(x),\ w^*(x) - \bar{\varepsilon} \Big\} \quad \text{and} \\
\underline{w}(x) &= \max\Big\{ \beta v(x),\ w^*(x) - \underline{\varepsilon} \Big\} \quad \text{with} \\
w^*(x) &= \max\big\{ w^{prepay}(x), w^{refi}(x) \big\} \\
w^{prepay}(x) &= \max_{b', m'}\ w(b', m', y, r^m, r) \quad \text{s.t. prepayment constraint (4) holds} \\
w^{refi}(x) &= \max_{b', m'}\ w(b', m', y, r + \omega^m, r) \quad \text{s.t. refinancing constraint (5) holds}
\end{aligned} \qquad (16)$$

Relative to (9), there is a new line $\underline{w}(x) = \max\Big\{ \beta v(x),\ w^*(x) - \underline{\varepsilon} \Big\}$ that captures the current-value of a household that has the opportunity to refinance at the low-effort-cost $\underline{\varepsilon}$.

Like in Appendix B.1, the current-value function in (16) can be derived from a setup in which the current self lives for a discrete length of time Δ:

$$
\begin{aligned}
w(x) &= \max\left\{ \max_c \; u(c)\Delta + \beta e^{-\rho\Delta}\left[e^{-\phi\Delta}\mathbb{E}[v(x_\Delta)|x] + \left(1-e^{-\phi\Delta}\right)\mathbb{E}[\underline{v}(x_\Delta)|x]\right],\; w^*(x) - \bar{\varepsilon}\right\},\\
\underline{w}(x) &= \max\left\{ \max_c \; u(c)\Delta + \beta e^{-\rho\Delta}\left[e^{-\underline{\phi}\Delta}\mathbb{E}[\underline{v}(x_\Delta)|x] + \left(1-e^{-\underline{\phi}\Delta}\right)\mathbb{E}[v(x_\Delta)|x]\right],\; \underline{w}^*(x) - \underline{\varepsilon}\right\},\\
w^*(x) &= \max\left\{w^{prepay}(x), w^{refi}(x)\right\}\\
w^{prepay}(x) &= \max_{b',m'} w(b',m',y,r^m,r) \quad \text{s.t. prepayment constraint (4) holds}\\
w^{refi}(x) &= \max_{b',m'} w(b',m',y,r+\omega^m,r) \quad \text{s.t. refinancing constraint (5) holds}\\
\underline{w}^*(x) &= \max\left\{\underline{w}^{prepay}(x), \underline{w}^{refi}(x)\right\}\\
\underline{w}^{prepay}(x) &= \max_{b',m'} \underline{w}(b',m',y,r^m,r) \quad \text{s.t. prepayment constraint (4) holds}\\
\underline{w}^{refi}(x) &= \max_{b',m'} \underline{w}(b',m',y,r+\omega^m,r) \quad \text{s.t. refinancing constraint (5) holds}
\end{aligned}
$$

where ϕ and $\underline{\phi}$ denote the Poisson switching rates between the two effort-cost states. As stated in Assumption 1 we assume that $\underline{\phi} \to \infty$. Therefore $e^{-\underline{\phi}\Delta} \to 0$ and

$$
\begin{aligned}
w(x) &= \max\left\{ \max_c \; u(c)\Delta + \beta e^{-\rho\Delta}\left[e^{-\phi\Delta}\mathbb{E}[v(x_\Delta)|x] + \left(1-e^{-\phi\Delta}\right)\mathbb{E}[\underline{v}(x_\Delta)|x]\right],\; w^*(x) - \bar{\varepsilon}\right\},\\
\underline{w}(x) &= \max\left\{ \max_c \; u(c)\Delta + \beta e^{-\rho\Delta}\mathbb{E}[v(x_\Delta)|x],\; \underline{w}^*(x) - \underline{\varepsilon}\right\},\\
w^*(x) &= \max\left\{w^{prepay}(x), w^{refi}(x)\right\}\\
w^{prepay}(x) &= \max_{b',m'} w(b',m',y,r^m,r) \quad \text{s.t. prepayment constraint (4) holds}\\
w^{refi}(x) &= \max_{b',m'} w(b',m',y,r+\omega^m,r) \quad \text{s.t. refinancing constraint (5) holds}\\
\underline{w}^*(x) &= \max\left\{\underline{w}^{prepay}(x), \underline{w}^{refi}(x)\right\}\\
\underline{w}^{prepay}(x) &= \max_{b',m'} \underline{w}(b',m',y,r^m,r) \quad \text{s.t. prepayment constraint (4) holds}\\
\underline{w}^{refi}(x) &= \max_{b',m'} \underline{w}(b',m',y,r+\omega^m,r) \quad \text{s.t. refinancing constraint (5) holds}
\end{aligned}
$$

Finally, we take the limit as $\Delta \to 0$. Using the property that $\underline{w}^*(x) = w^*(x)$ in the limit as $\Delta \to 0$ – which one can see by inspection since the left branch of the first line converges to the left branch of the second line – we recover equation (16).

C Proofs

C.1 Proof of Corollary 1

Recall that, with naiveté, the perceived continuation-value function of a $\beta < 1$ household equals the value function of an exponential $\beta = 1$ household and solves (8″). Assume that the household does not refinance at time t so that the perceived continuation-value function $v(x_t)$ is characterized by a standard HJB equation. This HJB equation is given by the left branch of (8″), which we write here as

$$\rho v(x) = \max_c \; u(c) + \frac{\partial v(x)}{\partial b}(y + rb + \omega^{cc} b^- - (r^m + \xi)m - c) + (\mathcal{B}v)(x) \tag{17}$$

where the operator $(\mathcal{B}v)(x)$ is short-hand notation for lines two to seven of (8″). Recall that we use hat-notation to denote the policy functions that naive households perceive for future selves, and denote by $\widehat{c}(x)$ and $\widehat{s}(x) = (y + rb + \omega^{cc} b^- - (r^m + \xi)m - \widehat{c}(x))$ the corresponding perceived consumption and liquid saving policy functions. In contrast, denote by $c(x)$ (from Proposition 1) and $s(x) = (y + rb + \omega^{cc} b^- - (r^m + \xi)m - c(x))$ the *actual* policy functions.

The following observation is important in the proof below: the HJB equation for the perceived continuation-value function (17) features the *perceived* policy functions $\widehat{c}(x), \widehat{s}(x)$, rather than the *actual* policy functions. But what determines the evolution of liquid wealth b are the *actual* policy functions.

Differentiate (17) with respect to b and use the envelope theorem:

$$(\rho - r(b))\frac{\partial v(x)}{\partial b} = \frac{\partial^2 v(x)}{\partial b^2}\widehat{s}(x) + \frac{\partial}{\partial b}(\mathcal{B}v)(x). \tag{18}$$

Define the marginal continuation-value of wealth $\eta(x) \equiv \frac{\partial v(x)}{\partial b}$. From (18) it satisfies

$$(\rho - r(b))\,\eta(x) = \frac{\partial \eta(x)}{\partial b}\widehat{s}(x) + (\mathcal{B}\eta)(x). \tag{19}$$

If $\beta = 1$, from Itô's formula, the right-hand side of (19) also governs the expected change in the marginal value of wealth: $\mathbb{E}_t[d\eta(x_t)] = \left[\frac{\partial \eta(x_t)}{\partial b}\widehat{s}(x_t) + (\mathcal{B}\eta)(x_t)\right] dt$. But with $\beta < 1$ this is no longer true: the evolution of b is governed by the *actual* drift $s(x)$ rather than the *perceived* drift $\widehat{s}(x)$ and so

$$\mathbb{E}_t[d\eta(x_t)] = \left[\frac{\partial \eta(x_t)}{\partial b}s(x_t) + (\mathcal{B}\eta)(x_t)\right] dt. \tag{20}$$

Therefore, evaluating (19) along a particular trajectory x_t, we have

$$(\rho - r(b_t))\,\eta(x_t) = \frac{1}{dt}\mathbb{E}_t[d\eta(x_t)] - \frac{\partial \eta(x_t)}{\partial b}(s(x_t) - \widehat{s}(x_t)).$$

Rearranging

$$\begin{aligned}
\frac{1}{dt}\mathbb{E}_t[d\eta(x_t)] &= (\rho - r(b_t))\,\eta(x_t) + \frac{\partial \eta(x_t)}{\partial b}(s(x_t) - \widehat{s}(x_t)) \\
&= (\rho - r(b_t))\,\eta(x_t) + \frac{\partial \eta(x_t)}{\partial b}(\widehat{c}(x_t) - c(x_t)) \\
&= (\rho - r(b_t))\,\eta(x_t) + \frac{\partial \eta(x_t)}{\partial b}\left(\beta^{\frac{1}{\gamma}} - 1\right) c(x_t)
\end{aligned}$$

Finally, recalling that $\eta(x) \equiv \frac{\partial v(x)}{\partial b}$, the first-order condition is $u'(c(x)) = \beta\eta(x)$ and therefore

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

where going from the second line to the third line uses that, with CRRA utility, the coefficient of relative risk aversion is $\gamma = \frac{-u''(c(x_t))c(x_t)}{u'(c(x_t))}$. Dividing by $u'(c(x_t))$, we have (12). ■

C.2 Proof of Proposition 2

When proving Proposition 2, we refer to Appendix B.3 which spells out the full set of equations for the model with a stochastic effort cost satisfying Assumption 1.

We also note that clauses 1, 2a, and 2b do not rely on the instantaneous low-effort-cost period used in Assumption 1. The purpose of Assumption 1 is to create the sorts of deadlines that incentivize present-biased agents to complete effortful tasks (clause 2c).

C.2.1 Proof of Proposition 2, Clause 1

The proof of clause 1 follows from equation (16). Equation (16) shows that we can rewrite w^{prepay} and w^{refi} as:

$$\begin{aligned}
w^{prepay}(x) &= \max_{b',m'} \ \beta v(b', m', y, r^m, r) \quad \text{s.t. prepayment constraint (4) holds} \\
w^{refi}(x) &= \max_{b',m'} \ \beta v(b', m', y, r + \omega^m, r) \quad \text{s.t. refinancing constraint (5) holds}
\end{aligned}$$

These are exactly the same formulas as for v^{prepay} and v^{refi} in (7), except that there is an additional β discount factor. Since the additional β discount factor has no effect on the optimal choice of (b', m'), we recover clause 1 of Proposition 2 — the choice of (b', m') is independent of β. ■

Since $v^*(x) = \max\left\{v^{prepay}(x), v^{refi}(x)\right\}$ and $w^*(x) = \max\left\{w^{prepay}(x), w^{refi}(x)\right\}$, the above proof also implies that:

$$w^*(x) = \beta v^*(x). \tag{21}$$

This property will be used in the proof of clause 2 of Proposition 2.

C.2.2 Proof of Proposition 2, Clause 2

For clause 2a, when $\beta = 1$ the assumption that $\bar{\varepsilon}$ and $\underline{\varepsilon}$ are vanishingly small (Assumption 2) implies that $v(x)$ is arbitrarily close to $\underline{v}(x)$. Accordingly, policy function $\mathfrak{R}(x)$ converges pointwise to $\underline{\mathfrak{R}}(x)$ as the effort cost vanishes.

To prove clause 2b (procrastination when $\beta < 1$ and $\varepsilon = \bar{\varepsilon}$), consider the self in control at point x in the state space. Recall from (16) that the current-value function is given by $w(x) = \max\{\beta v(x), w^*(x) - \bar{\varepsilon}\}$. Therefore the current self will not adjust their mortgage when the value from not adjusting, $\beta v(x)$, is larger than the value from adjusting, $w^*(x) - \bar{\varepsilon}$.

The value of not adjusting is given by

$$\beta v(x) \geq \beta(v^*(x) - \bar{\varepsilon}), \tag{22}$$

where the inequality $v(x) \geq v^*(x) - \bar{\varepsilon}$ follows directly from equation (8′′).

Alternatively, adjusting requires the household to incur the effort cost $\bar{\varepsilon}$ in the current period and the value of adjusting is given by

$$w^*(x) - \bar{\varepsilon} = \beta v^*(x) - \bar{\varepsilon}, \tag{23}$$

where the equality follows from equation (21).

Comparing the two alternatives (22) and (23) shows that the $\beta < 1$ household will always prefer to procrastinate whenever $\varepsilon_t = \bar{\varepsilon}$, since

$$\beta(v^*(x) - \bar{\varepsilon}) > \beta v^*(x) - \bar{\varepsilon}.$$

Procrastination enables the effort cost $\bar{\varepsilon}$ to be discounted by β, while there is at most an infinitesimal cost to delaying refinancing for an instant.

To prove clause 2c (no procrastination when $\beta < 1$ and $\varepsilon = \underline{\varepsilon}$), consider the self in control at point x in the state space. Following the second line of equation (16), it will be (weakly) optimal for the current self to adjust their mortgage if and only if:

$$w^*(x) - \underline{\varepsilon} \geq \beta v(x).$$

Above, the left-hand side is the current-value from refinancing at effort cost $\underline{\varepsilon}$, and the right-hand side is the current-value from not refinancing and having the effort cost reset immediately to $\bar{\varepsilon}$. Since $w^*(x) = \beta v^*(x)$ (see equation (21)), this can be rewritten as

$$\beta v^*(x) - \underline{\varepsilon} \geq \beta v(x). \tag{24}$$

First, consider the case in which the next self is expected to adjust the mortgage if the current self procrastinates.[56] Since the next self is expected to have $\beta = 1$, this means $\widehat{\mathfrak{R}}(x) > 0$. In this case, equation ($8''$) implies that $v(x) = v^*(x) - \bar{\varepsilon}$. Plugging this into (24) shows that the current self will adjust their mortgage whenever $\beta v^*(x) - \underline{\varepsilon} \geq \beta v^*(x) - \beta\bar{\varepsilon}$ or

$$\underline{\varepsilon} \leq \beta\bar{\varepsilon},$$

which is satisfied because Assumption 1 imposes that $\underline{\varepsilon} < \beta\bar{\varepsilon}$. Intuitively, this says that the current self will adjust their mortgage now if the cost of doing so, $\underline{\varepsilon}$, is less than the discounted cost of adjusting next period, $\beta\bar{\varepsilon}$. Thus, if $\widehat{\mathfrak{R}}(x) > 0$ then $\underline{\mathfrak{R}}(x) = \widehat{\mathfrak{R}}(x)$, meaning that the household does not procrastinate.

Next, consider the case in which a $\beta = 1$ household would not refinance at point x, even in the low-effort-cost state $\varepsilon_t = \underline{\varepsilon}$, i.e., $\widehat{\underline{\mathfrak{R}}}(x) = 0$. In that case, equation (15) implies that $v(x) \geq v^*(x) - \underline{\varepsilon}$. Multiplying by β, this also implies that $\beta v(x) \geq \beta v^*(x) - \beta\underline{\varepsilon}$, and therefore

$$\beta v(x) > \beta v^*(x) - \underline{\varepsilon}.$$

Comparing this to equation (24) shows that it will not be optimal for the naive present-biased household to refinance. This is intuitive — if it is not optimal for a $\beta = 1$ household to refinance, there is no reason for it to be optimal for a naive $\beta < 1$ household to refinance. Thus, if $\widehat{\underline{\mathfrak{R}}}(x) = 0$ then $\underline{\mathfrak{R}}(x) = \widehat{\underline{\mathfrak{R}}}(x)$.

Tying these two cases together, we have shown that:

1. If $\widehat{\mathfrak{R}}(x) > 0$ then $\underline{\mathfrak{R}}(x) = \widehat{\mathfrak{R}}(x)$
2. If $\widehat{\underline{\mathfrak{R}}}(x) = 0$ then $\underline{\mathfrak{R}}(x) = \widehat{\underline{\mathfrak{R}}}(x)$

Since clause 2a of Proposition 2 implies that $\widehat{\underline{\mathfrak{R}}}(x)$ converges pointwise to $\widehat{\mathfrak{R}}(x)$ as the effort

[56]Note that the next self will face the high-effort-cost $\bar{\varepsilon}$ if the current self procrastinates when $\varepsilon_t = \underline{\varepsilon}$.

cost vanishes, the first bullet above can be rewritten as: if $\widehat{\underline{\mathfrak{R}}}(x) > 0$ then $\underline{\mathfrak{R}}(x)$ converges pointwise to $\widehat{\underline{\mathfrak{R}}}(x)$. This completes the proof of clause 2c of Proposition 2. ■

D Model Extensions and Robustness

D.1 The Distributional Effects of Policy

The top row of Figure 6 breaks down the consumption response to fiscal policy for our two benchmark cases. Each panel plots the consumption response to fiscal stimulus on impact as a function of pre-shock consumption. In the Exponential Benchmark (left) the consumption response is relatively evenly spread across the consumption distribution. In the Present-Bias Benchmark (right) this is not the case — the lowest consumption households experience a drastic consumption boom from fiscal policy. These households are borrowing-constrained, and sharply increase consumption following the liquidity shock.

The bottom row of Figure 6 breaks down the consumption response to monetary policy on impact. In the Exponential Benchmark, the largest consumption response comes from low-consumption households. These households are near $\underline{b}$, and implement a cash-out refinance following the rate cut. Thus, in the Exponential Benchmark the refinancing channel of monetary policy endogenously targets itself to constrained households.

Alternatively, in the Present-Bias Benchmark the low-consumption households respond very little to monetary policy on impact. The largest response now comes from households with intermediate levels of pre-shock consumption. Low-consumption households are constrained on impact, and because they procrastinate on refinancing they cannot immediately adjust consumption. Households with intermediate consumption are not liquidity constrained on impact. These households will typically be in either a refinancing region following the rate cut (in which case they expect to refinance in the next instant), or near a refinancing region (in which case they expect to refinance soon). In both cases, consumption smoothing implies that these households will increase consumption today in expectation of the cash-out refinance that they plan to conduct in the near future. This ability to smooth consumption relies on pre-existing liquidity at the time of the monetary policy shock, which households at $\underline{b}$ do not have.

As discussed in the main text, the key takeaway from Figure 6 is that present bias reverses the distributional consequences of fiscal versus monetary policy. In the Exponential Benchmark, monetary policy is an effective way to stimulate the consumption of low-consumption households. In the Present-Bias Benchmark, procrastination hampers the ability for monetary policy to stimulate the short-run consumption of constrained households. However, fiscal policy instead becomes highly effective at increasing these households' consumption.

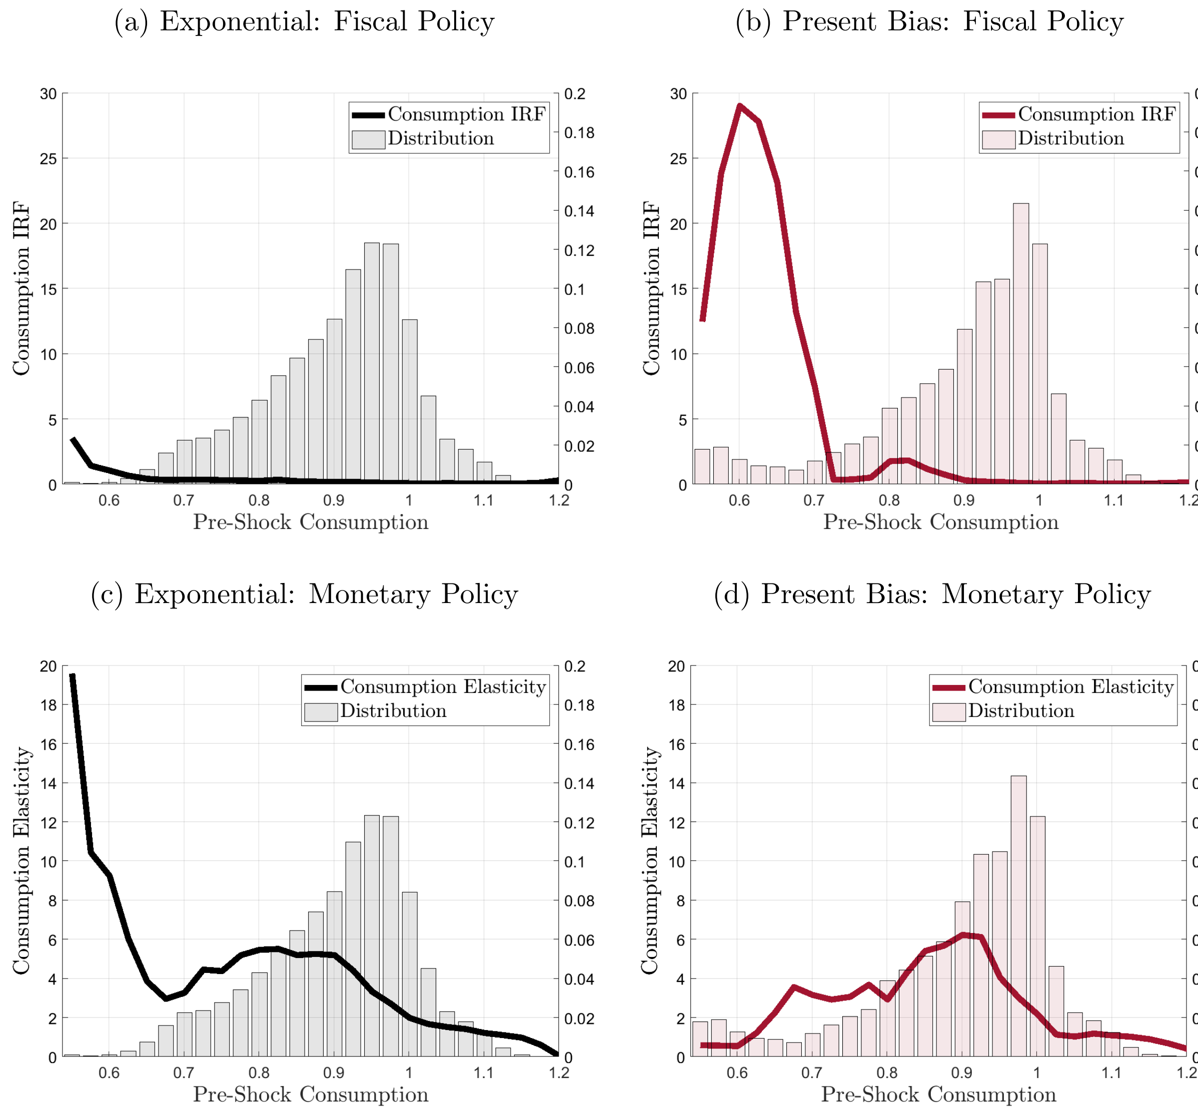

Figure 6: Heterogeneity Analysis.

Notes: This figure plots the on-impact consumption response to fiscal (top row) and monetary (bottom row) policy as a function of households' pre-shock consumption. The solid line plots the consumption response, and the bars show the distribution of households over pre-shock consumption.

D.2 Adding Aggregate House Price and Income Shocks

Macroeconomic stabilization policy typically responds to shocks hitting the economy. In particular, expansionary monetary and fiscal policy are often used in recessions. Recessionary shocks, by definition, correspond to a temporary decline in aggregate income. Recessions can also coincide with declining house prices. This section examines the ways in which shocks to house prices and aggregate income affect our results.

D.2.1 House Price Shocks

Present bias amplifies monetary policy by producing a consumption boom driven by home-equity extractions. However, this cash-out channel of monetary policy is limited to households with enough home equity to actually conduct a cash-out refinance. This makes monetary policy sensitive to house price shocks, which can quickly create or destroy home equity.

Understanding the effect of house price shocks on macroeconomic policy is particularly important when considering the three most recent recessions: the COVID-19 Recession, the Great Recession, and the Early 2000s Recession. Home prices collapsed during the Great Recession, but boomed throughout the Early 2000s Recession and the COVID-19 Recession.

Our baseline analysis in Section 5 corresponds to the case where home prices are stable before the cut to interest rates. To examine the effect of house price shocks we exogenously shock the home value h by $\pm$ 25%.[57] The negative 25% shock corresponds to the Great Recession. The positive 25% shock corresponds to the early 2000s, where house prices boomed while the Federal Reserve adopted a multi-year path of low interest rates. Second, policymakers immediately respond to this house price shock with either monetary or fiscal policy. As in Section 5, the monetary policy experiment is a rate cut from 1% to 0%, and the fiscal policy experiment is a $1,000 liquid transfer.

Figure 7 plots the consumption response to monetary policy after a negative (left panel) or positive (right panel) 25% shock to house prices. The solid curves plot the consumption response to monetary policy in the shocked economy. For reference, the transparent lines mark the baseline case in Figure 5. Though the magnitude of the consumption response is sensitive to house price shocks, our main result that present bias amplifies the consumption response to monetary policy holds in both cases.

The left panel of Figure 7 shows that monetary policy is significantly weakened by the collapse in house prices. The negative shock wipes out home equity and prevents many homeowners from refinancing. This result is consistent with recent research documenting that negative house price shocks undermined monetary policy following the Great Recession (e.g., Beraja et al., 2019).

The right panel of Figure 7 plots the positive 25% shock case. Now, the consumption boom generated by the rate cut is even larger than in the baseline case. The positive shock generates additional home equity, strengthening the cash-out channel of monetary policy. This is consistent with the boom in home-equity extractions that was observed in the mid-2000s (Khandani et al., 2013; Bhutta and Keys, 2016).

It is also important to explore whether house price shocks affect fiscal policy. Figure 8 plots the consumption response to fiscal stimulus in the negative (left) and positive (right)

[57] The economy starts in the "steady state" before the shock to h. For simplicity we assume that the shock is permanent. However, we only study the short-run consumption response to monetary and fiscal policy.

(a) -25% House Price Shock (b) +25% House Price Shock

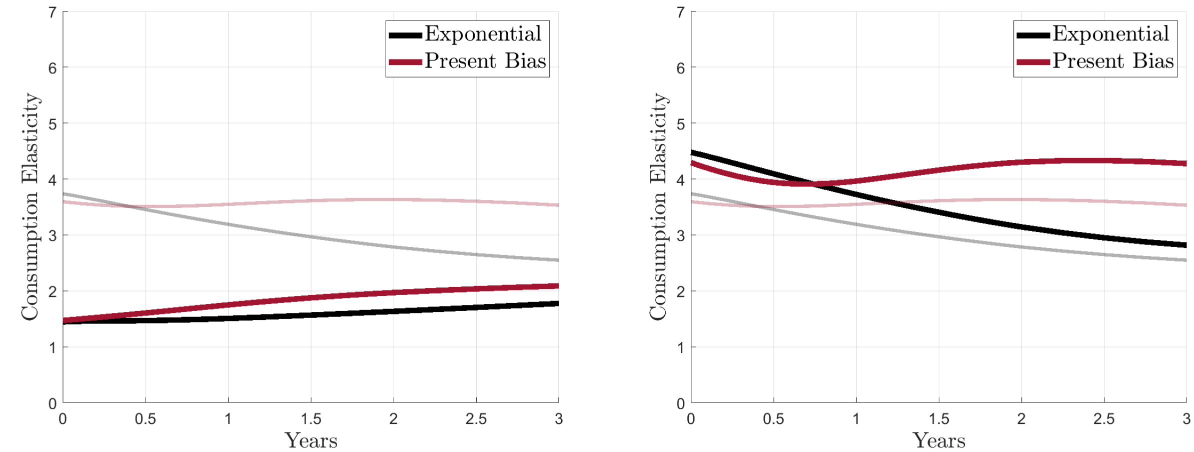

Figure 7: Monetary Policy and House Price Shocks.
Notes: This figure plots the consumption response to an interest rate cut that immediately follows a house price shock of -25% (left) or +25% (right). The transparent lines plot the baseline case in Figure 5, and are included for reference.

shock case. For both positive and negative house price shocks, we find that present bias continues to strongly amplify the consumption response to fiscal policy.

(a) -25% House Price Shock (b) +25% House Price Shock

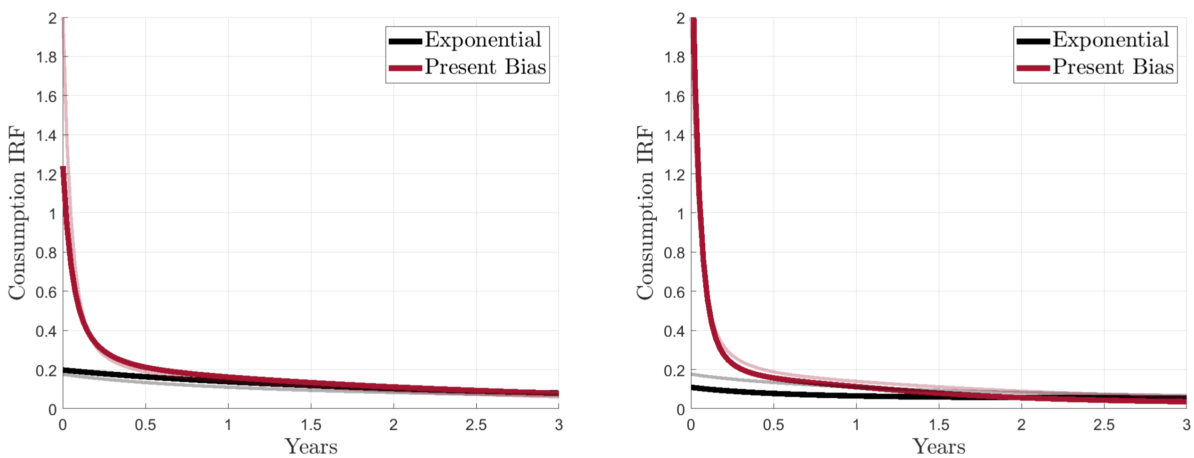

Figure 8: Fiscal Policy and House Price Shocks.
Notes: This figure plots the IRF of aggregate consumption to a $1,000 fiscal transfer that immediately follows a house price shock of -25% (left) or +25% (right). The transparent lines plot to the baseline case in Figure 4, and are included for reference.

D.2.2 Income Shocks

We also evaluate the effect of recessionary income shocks on monetary and fiscal policy. We generate a temporary 5% fall in aggregate income by shifting a share of high-income households to the middle-income state, and a share of middle-income households to the low-income state.[58] Policymakers immediately respond to this recessionary income shock with either monetary or fiscal policy.

The left panel of Figure 9 plots the consumption response to monetary policy and the right panel plots the consumption response to fiscal policy. Though this recessionary income shock leads to an immediate decline in aggregate consumption, Figure 9 shows that the subsequent consumption response to monetary and fiscal policy is almost identical to the baseline results in Section 5. This is because liquidity, not income, is the key driver of the consumption response to these policies.

(a) Monetary Policy (b) Fiscal Policy

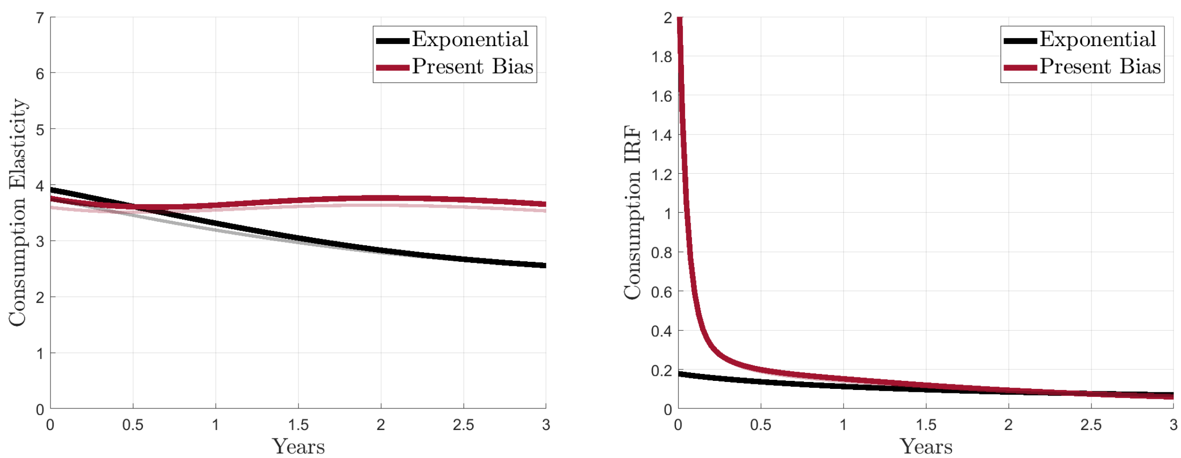

Figure 9: Fiscal and Monetary Policy Following a Negative Income Shock.
Notes: This figure plots the consumption response to monetary (left) and fiscal (right) policy that is implemented immediately following a transitory 5% decline in aggregate income.

D.3 A Call to ARMs?

In order to reflect the typical features of the U.S. mortgage market our paper studies macroeconomic stabilization policy under the assumption that households have fixed-rate mortgages (FRMs). Since the 2007-08 Financial Crisis, many economists have argued that downwardly

[58] We shift 9.5 percentage points of high-income households to middle income, and 9.5 percentage points of middle-income households to low income. The share of households across low, middle, and high income goes from 31%, 39%, 31%, respectively, to 40%, 39%, 21%.

flexible mortgages, such as adjustable-rate mortgages (ARMs), improve macroeconomic stability (e.g., Eberly and Krishnamurthy, 2014; Andersen et al., 2020; Campbell et al., 2021; Guren et al., 2021). If the monetary authority cuts interest rates in a recession, ARM payments automatically adjust downward, thereby increasing households' disposable income. This creates a fast and direct transmission of monetary policy to household balance sheets.

Section 5.2 shows that refinancing procrastination slows down the transmission of monetary policy. A natural policy question is whether such procrastination implies that downwardly flexible mortgages would improve the potency of monetary policy in our model with present bias.

To study this question, we re-solve our Present-Bias Benchmark under the assumption that all households have ARMs instead of FRMs. The model from Section 2 remains the same, except that mortgage rate r_t^m automatically adjusts with interest rate r_t. We also recalibrate the mortgage wedge from 1.7% to 0.9%.[59] This corresponds to the average difference between a 5/1 hybrid ARM and the 10-year treasury yield from 2015 – 2017.

In our monetary policy experiment with ARMs we cut r_t from 1% down to -1%. This doubles the magnitude of the rate cut from our earlier FRM analysis, where interest rates were reduced from 1% to 0%. This change ensures comparability across the two experiments, since ARMs are more sensitive to monetary policy than long-duration FRMs.[60]

We find that the consumption response to monetary policy is almost identical with ARMs versus FRMs (see Appendix Figure 17 for details). This result highlights that there is a tradeoff between ARMs and FRMs that arises when households are present biased. On the one hand, ARMs produce a fast pass-through of monetary policy that applies to all mortgage holders. This is particularly important for constrained households who procrastinate on refinancing. On the other hand, ARMs reduce the liquidity injection features of monetary policy because ARMs imply that households no longer need to refinance when the interest rate is cut. Present bias generates a powerful cash-out channel of monetary policy, but this channel is stifled by ARMs. Overall, the stimulative effect of ARMs accrues quickly and to all households, but is small. The stimulative effect of FRMs accrues slowly and only to households who plan to refinance, but is large. These two effects are of similar magnitude in our model.

An important factor explaining these offsetting effects is the large size difference between ARM payment adjustments versus home-equity extractions. Recall that the home value is calibrated to 3.29 times permanent income, and the average LTV ratio is 0.51. With ARMs, a 2% reduction in mortgage rates is equivalent to a $2\% \times 0.51 \times 3.29 = 3\%$ increase in income

[59] FRMs are typically more expensive than ARMs because FRMs lead to lower payments if interest rates rise, and come with the option to refinance if rates fall. The borrower has to pay ex-ante for this insurance.

[60] Empirically there is roughly a 50% pass-through from monetary policy to the 30-year mortgage interest rate. See Gertler and Karadi (2015), Gilchrist et al. (2015), and Eichenbaum et al. (2022) for details.

for the average household (roughly $3,000 per year). Alternatively, with FRMs the typical cash-out is about 30% of permanent income ($30,000) – an order of magnitude larger than the typical ARM payment reduction – and present-biased households have large MPCs out of these liquidity injections. Such large cash-outs are consistent with the data. The average cash-out amount from 1999–2010 was $40,000 (Bhutta and Keys, 2016), and there is little evidence that the majority of these home-equity extractions are kept as savings (Greenspan and Kennedy, 2008; Bhutta and Keys, 2016). However, ARMs prevent this large stock of dry powder from ever being ignited.

While we do not find that ARMs increase the power of monetary policy, we note that our model is too stylized to make rigorous quantitative claims. Our analysis also assumes that house prices are fixed. Negative house price shocks, such as those observed following the financial crisis, can significantly reduce the cash-out channel of FRMs (see Section D.2.1). Our results nevertheless highlight a new tradeoff between FRMs and ARMs that policymakers should be aware of when considering different mortgage contract designs. Our results also suggest that monetary policy is most powerful if mortgage contracts feature a fast pass-through (like ARMs) while simultaneously allowing for cash-outs (like FRMs). Appendix Figure 18 shows just how powerful monetary policy can be in a FRM environment if policymakers are able to reduce procrastination concurrently with a monetary expansion.

In addition to being stylized this section ignores important welfare considerations. For example, FRMs produce a consumption boom by encouraging overconsumption out of home equity. Monetary policy also appears to be more equitable under ARMs than FRMs. In Section D.1 we showed that low-consumption households procrastinate on refinancing a FRM, whereas ARMs provide immediate payment relief to low-consumption households.[61]

D.4 Alternate Calibrations: Intermediate Cases

As mentioned in Section 4.2, we also examine various Intermediate Cases that differentially allow for present bias and/or refinancing inertia. We introduce four additional cases here which, when combined with the two Benchmark calibrations in the main text, mean that we study six calibration cases overall. These six calibration cases are summarized in Table 5 below, and we describe them more fully in the next paragraph. Overall, we consider three types of refinancing inertia (no inertia, rational inertia, and procrastination), and two types of time preferences (exponential preferences and present-biased preferences). In Table 5, the cases marked Benchmark are the cases that we already studied in the main text, and the cases marked IC are additional intermediate cases that we present below.

[61] As discussed in Harris and Laibson (2013) and Maxted (2023), the IG model is well-suited for studying welfare because it features a single welfare criterion despite preferences being dynamically inconsistent. As mentioned in the conclusion, we view welfare analyses as an interesting pathway for future research.

	Exponential	Present Bias
No Refinancing Inertia	**Benchmark**	IC: $\varepsilon_t \equiv 0$
Rational Inertia	IC: $\bar{\varepsilon} \to \infty$	IC: $\bar{\varepsilon} \to \infty$
Procrastination	IC: $\beta \uparrow 1$	**Benchmark**

Table 5: Summary of Calibration Cases.

Notes: The case with exponential time preferences and no refinancing inertia is our Exponential Benchmark. The case with present-biased time preferences and procrastination is our Present-Bias Benchmark. The other four cases marked IC are intermediate cases that we examine here.

Starting with the first row, the Exponential Benchmark studied in the main text has exponential time preferences and no refinancing inertia. Moving to the second column of that row, our first intermediate case is a model with present bias ($\beta < 1$) but *without* refinancing inertia, which we achieve by setting all effort costs equal to zero ($\varepsilon_t \equiv 0$). Note that present bias without refinancing inertia can also be viewed as the case of present bias under full sophistication (further details in Appendix D.5).

In the second row, we study an alternate setup that generates what we refer to as rational inertia. Specifically, we break Assumption 2 that effort costs are vanishingly small, and instead take the typical effort cost $\bar{\varepsilon} \to \infty$ so that households *optimally* do not refinance in the high-cost state (since the effort cost of doing so is exceedingly onerous). Below, we study this alternate setup of rational inertia both for exponential and present-biased households.

In the third row, we return to our original assumption that effort costs are vanishingly small (Assumption 2). While we cannot generate slow refinancing in this case when $\beta = 1$, any amount of naive present bias is sufficient to generate procrastination.[62] Hence, the case in column one of the third row is not a true "exponential" case. Instead, we numerically set $\beta = 0.999$: by setting $\beta < 1$ we introduce procrastination, but this tiny amount of present bias also has essentially no effect on households' consumption choices. Finally, our Present-Bias Benchmark studied in the main text features both present bias and procrastination.

In all four of the intermediate cases presented below, we recalibrate the discount function to fit the same steady state targets as the Benchmark calibrations. Specifically, for all cases we recalibrate ρ to fit the LTV moment, and for the present-bias cases we also recalibrate β to fit the credit card borrowing moment. The purpose of these intermediate cases is not necessarily to be realistic, but rather to provide various stepping stones that help the reader traverse between the Exponential Benchmark and the Present-Bias Benchmark. Additionally, the comparison between the Rational Inertia cases and the Procrastination cases highlights the extent to which naive present bias can generate *unexpected* refinancing inertia, which differentiates present-bias-driven procrastination from the sorts of "rational"

[62] I.e., there is a discontinuity in the limit as $\beta \to 1$; with any amount of naive present bias, the effort costs will always be small enough under Assumption 2 that households choose to procrastinate.

refinancing inertia that could arise under exponential time preferences.

D.4.1 Intermediate Cases: Exponential Discounting

We start by presenting the intermediate cases with exponential discounting. The left panel of Figure 10 plots the consumption response to fiscal policy, and the right panel plots the consumption response to monetary policy. Both panels plot the two Benchmark calibrations (opaque lines), and the two new intermediate cases (transparent lines) in order to provide stepping stones between the two benchmarks. Solid lines correspond to No Refinancing Inertia, dashed lines to Rational Inertia, and dotted lines to Procrastination.

(a) Fiscal Policy: Exponential Cases (b) Monetary Policy: Exponential Cases

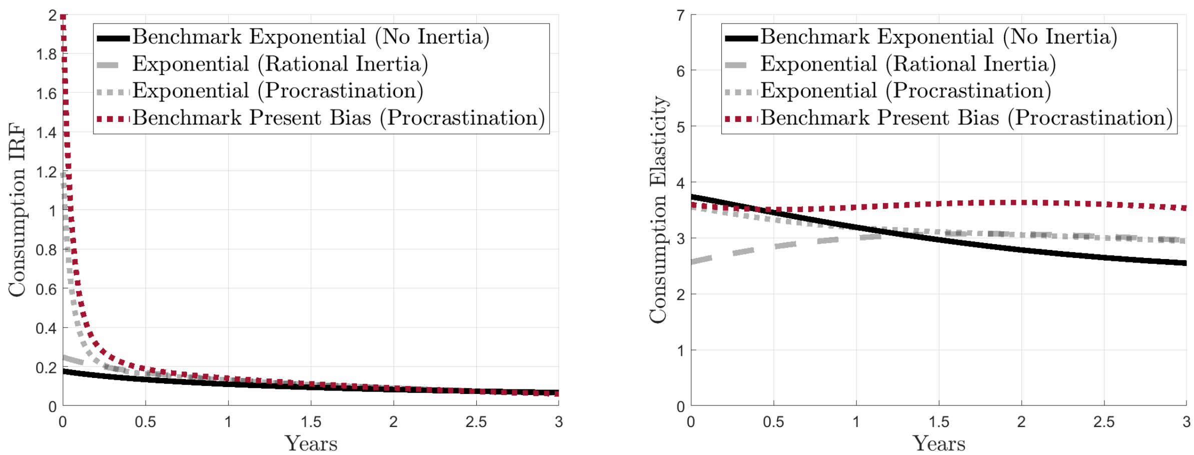

Figure 10: Fiscal and Monetary Policy with Exponential Intermediate Cases.
Notes: This figure plots the consumption response to fiscal (left) and monetary (right) policy for the two Benchmark calibrations and the two Intermediate Cases with exponential preferences (see Table 5).

Starting with fiscal policy, the consumption response is comparable in the Exponential Benchmark and the exponential case with rational inertia, but is much larger in the cases with procrastination. This relates to the share of borrowing-constrained households. In the Exponential Benchmark, there are very few households at $\underline{b}$ because they refinance as soon as they hit the constraint. In the case with rational inertia, households know that they will be slow to refinance, so they endogenously refinance before hitting $\underline{b}$ if they are in a stochastic low-cost-state. Only in the case where households *unexpectedly* procrastinate on refinancing do we get a buildup of households at $\underline{b}$ that generates a sizable short-run consumption response.

Turning to monetary policy, the notable feature is that we see less of an on-impact consumption response when households are rationally inertial. In this case, households are fully aware that they will be slow to refinance following the rate cut, which makes them more cautious about increasing consumption before that refinance is actually enacted.

D.4.2 Intermediate Cases: Present-Biased Discounting

Next, Figure 11 presents the intermediate cases with present-biased discounting.

(a) Fiscal Policy: Present-Bias Cases

(b) Monetary Policy: Present-Bias Cases

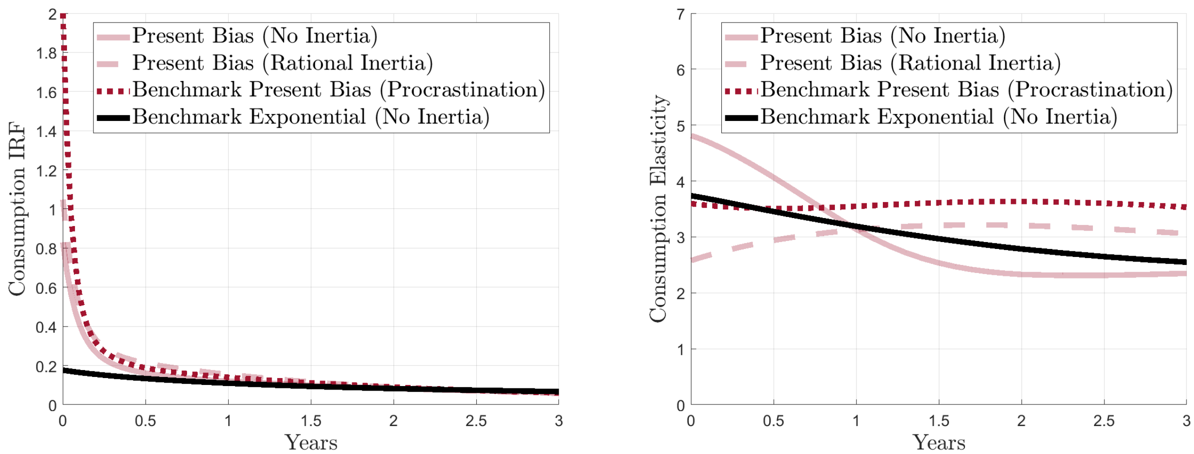

Figure 11: Fiscal and Monetary Policy with Present-Bias Intermediate Cases.
Notes: This figure plots the consumption response to fiscal (left) and monetary (right) policy for the two Benchmark calibrations and the two Intermediate Cases with present-biased preferences (see Table 5).

Starting with fiscal policy, we first see that present bias generates a larger consumption response than the Exponential Benchmark in all cases. Second, we again see that the consumption response is broadly comparable for the two intermediate cases of no inertia and rational inertia, but jumps up in the Present-Bias Benchmark with procrastination. The intuition here is broadly similar to our discussion of the exponential cases, where the key to getting a buildup of households at $\underline{b}$ is for them to unexpectedly procrastinate on refinancing.

Turning to monetary policy, we again see a smaller on-impact consumption response in the case of rational inertia, similar to the exponential case with rational inertia described above. For the present-bias case with no inertia, we now see a larger short-run consumption response, but also one that decays more quickly. The intuition for this no-inertia case is largely similar to the "Present Bias 1Q No Proc." case in Figure 5, except that the initial response here is somewhat more mild because there are fewer constrained households ex-ante (since households never procrastinate on refinancing).

D.5 Generalization to (Partial and Full) Sophistication

To this point we have assumed that households are fully naive about their present bias. One benefit of our naiveté assumption is that it is theoretically and computationally easy to handle — starting from a model without present bias, it is relatively simple to back out the behavior of households with naive present bias (see Propositions 1 and 2). As such, naive present bias provides a bridge for the broader macroeconomics literature.

The case with (partial or full) sophistication is a much more complicated theoretical object. Let β^E denote the short-run discount factor that the current self expects all future selves to have. Full naiveté means that $\beta^E \equiv 1$, whereas (partial) sophistication sets $\beta^E \in [\beta, 1)$ so that the current self is at least partially aware that future selves will also face a self-control problem. Sophistication therefore implies that the current self is aware that preferences are dynamically inconsistent, thus making behavior the equilibrium outcome of a dynamic game. Despite being more complicated theoretically, the technical advances of Harris and Laibson (2013) and Maxted (2023) imply that solutions to models with sophistication are still available. We now utilize these advances to extend our analysis to the case of sophistication.

In this appendix we ask how our results vary with sophistication, and we provide two perspectives. On the one hand, we show that a model with partial sophistication can be recalibrated with a different β so that it produces household-level behavior that is analogous to that of full naiveté. That is, while we assumed full naiveté in the main text because it is theoretically and computationally easy to handle, our analysis is robust to all but the limiting case of complete sophistication ($\beta^E = \beta$).[63]

On the other hand, one could also ask about comparative statics with respect to partial sophistication (while not recalibrating any other model parameters). In this case, we show that households' propensity to procrastinate is (weakly) decreasing in their sophistication.

D.5.1 The Household Balance Sheet: A Slight Modification

Formalizing this analysis is, admittedly, complex, and we provide only a heuristic analysis here. The material presented below likely requires familiarity with Harris and Laibson (2013) and Maxted (2023).

We begin by slightly modifying the model of the household balance sheet in Section 2. Note that this modification is not necessary. However, we make it so that we can use the theoretical results in Maxted (2023) to study the effects of sophistication *in closed form*, which we believe provides clearer economic insights.[64]

To utilize the results in Maxted (2023), we need to respecify the model so that households always remain in the interior of the liquid wealth space rather than occasionally facing a

[63] As we detail further below, the case of full sophistication does not generate procrastination under Assumption 2 that effort costs are vanishingly small. Intuitively, without some scope for incorrect expectations, we cannot use vanishingly small effort costs to generate non-vanishing spells of procrastination.

[64] Without this modification one could still numerically solve the model under sophistication by using the $\hat{u}$ technology developed in Harris and Laibson (2013), though numerical implementation could be difficult.

binding hard borrowing constraint at $\underline{b}$ (see the "Key Assumption" of Maxted (2023)).[65] To do so we make two (realistic) changes, as presented below.

First, we now allow households to borrow beyond the ad-hoc limit of $\underline{b}$ to a lower limit of $\underline{\underline{b}} < \underline{b}$. For all debt beyond $\underline{b}$, we assume that the household must incur a very large borrowing wedge of $\omega^{\uparrow} \gg \omega^{cc}$. For example, $\omega^{\uparrow} \approx 400\%$ is a typical payday loan interest rate (Lee and Maxted, 2023), but one could imagine $\omega^{\uparrow}$ to be taken even larger. Following the "Key Assumption" of Maxted (2023), we impose in Assumption 3 below that $\omega^{\uparrow}$ is made large enough that households always keep their liquidity strictly above $\underline{\underline{b}}$ in equilibrium.[66]

Second, we introduce a small delay to refinancing. In the model in the main text we assumed for simplicity that there are no delays to refinancing; refinancing occurs instantly after a household completes its application. Without refinancing delays, households that expect to refinance in the next instant will not care about borrowing at interest rate $\omega^{\uparrow}$. So long as households only expect their borrowing to persist for one instant, the rate at which they borrow will not affect their overall wealth. Thus, we additionally assume that, after a household fills out its refinancing paperwork, the new mortgage only closes at a Poisson rate denoted λ^C.[67] Such delays are realistic, as mortgage underwriting often takes well over a month.

With these updates, we now have a model that can be calibrated so that households will always remain in the interior of their liquid wealth space. That is, we assume:

Assumption 3 *The model is calibrated such that borrowing limit $\underline{\underline{b}}$ does not bind in equilibrium. Formally, if $b_0 > \underline{\underline{b}}$ then $b_t > \underline{\underline{b}}$ for all $t \geq 0$.*

We maintain Assumption 3 throughout the remainder of Appendix D.5, which will allow us to utilize the results of Maxted (2023) in this modified model.

We also briefly mention one other complexity, which is that a full-sophisticate's refinancing policy function takes the form of a mixed-strategy equilibrium where the household refinances probabilistically (e.g., O'Donoghue and Rabin, 2001). However, this added complexity effectively drops out under Assumption 2 that $\bar{\varepsilon}$ is vanishingly small, since in this case

[65] The basic issue which these modifications aim to address is that sophisticated present bias interacts with binding hard borrowing constraints like $\underline{b}$, because present-biased agents value such constraints as a commitment device of sorts that limits the overconsumption of future selves (see equation (3)). This new interaction – which does not arise under naiveté – limits our ability to directly map between sophistication and naiveté (see Maxted (2023) for a fuller discussion). Note too that had we instead kept this interaction effect, it would have led to *larger* consumption discontinuities at $\underline{b}$ and hence *larger* MPCs at $\underline{b}$ under sophistication. In other words, this interaction effect would have led to *larger* differences relative to exponential discounting.

[66] Note that this will always be possible once we introduce refinancing delays (next paragraph). The basic argument here is that as $\omega^{\uparrow}$ increases, then households must keep their debt closer and closer to $\underline{b}$ in order to avoid the possibility of a zero-consumption state (and hence $-\infty$ utility).

[67] Specifically, the household pays the effort cost up front, but does not finalize its new mortgage terms nor receive the cash out (or pay the cash in) until the mortgage closes (details in Appendix D.5.3).

the probability of refinancing over any discrete interval of time converges to one. Further details are left to Appendix D.5.4.

D.5.2 The Effect of Sophistication on Policy Functions and Results

We can now generalize Propositions 1 and 2 to allow for sophistication. To avoid complicating the exposition with additional technicalities that arise when households are sophisticated, we state the basic results here and then add details later in Appendices D.5.3 and D.5.4.

We continue to use hat-notation to denote the policy functions of the "comparable household" that has $\beta = 1$. However, in this more general setup that allows for sophistication, that comparable household is an otherwise-identical household that has $\beta = 1$ *and* faces a typical refinancing effort cost of $H^E \times \bar{\varepsilon}$ (rather than just $\bar{\varepsilon}$), where $H^E = \left(\frac{\gamma-(1-\beta^E)}{\gamma}\right)^{-\gamma}$. The intuition behind this rescaling is that, once agents are at least partially sophisticated ($\beta^E < 1$), they are aware that their self-control problems will cause them to act as if they face a higher hurdle rate to refinance.[68] Further details are provided in Appendix D.5.3 below. Note that this additional step of rescaling $\bar{\varepsilon}$ was not needed in the main text since we assumed that $\beta^E \equiv 1$ (full naiveté) so that $H^E = 1$. Nonetheless, this rescaling effect remains trivial here, due to Assumption 2 that $\bar{\varepsilon}$ is vanishingly small.

Using this agent with $\beta = 1$ and a typical refinancing effort cost of $H^E \times \bar{\varepsilon}$ as a point of comparison, we can start by characterizing the consumption decisions of present-biased agents using a result from Maxted (2023):

Proposition 3 (Continuous Control) *For all $b > \underline{b}$, the household sets*

$$c(x) = \left(\frac{\beta^E}{\beta}\right)^{\frac{1}{\gamma}} \frac{\gamma}{\gamma - (1-\beta^E)} \widehat{c}(x),$$

where $\widehat{c}(x)$ is the consumption policy function of an exponential $\beta = 1$ household but with effort cost $H^E \times \bar{\varepsilon}$ in place of $\bar{\varepsilon}$ where $H^E = \left(\frac{\gamma-(1-\beta^E)}{\gamma}\right)^{-\gamma}$.

Proof. See Maxted (2023), with additional details in Appendices D.5.3 and D.5.4. ∎

Proposition 3 is just like Proposition 1 in the main text, except that the consumption scaling factor is generalized to allow for sophistication. For our purposes, the key implication of Proposition 3 is the following corollary, which implies that there is a limiting observational equivalence between the consumption decisions of sophisticates and naifs:

[68] Accordingly, this perceived hurdle rate increases as the agent gets more sophisticated, i.e., as β^E decreases.

Corollary 2 *The consumption function of a naif with short-run discount factor β converges pointwise (as the effort cost $\bar{\varepsilon}$ vanishes) to the consumption function of a (partial) sophisticate with perceived $\beta^E \in [\beta, 1)$ and a true short-run discount factor of $\beta' = \beta\beta^E \left(\frac{\gamma}{\gamma-(1-\beta^E)}\right)^{\gamma}$. This means that there is an observational equivalence between the consumption policy of naifs and (partial) sophisticates in the limit as $\bar{\varepsilon} \to 0$.*

Proof. This corollary would follow directly from Proposition 3, except that the consumption function $\widehat{c}(x)$ changes with β^E (since we use hat-notation to denote an otherwise-identical household that has $\beta = 1$ *and* faces a typical refinancing effort cost of $H^E \times \bar{\varepsilon}$, where $H^E = \left(\frac{\gamma-(1-\beta^E)}{\gamma}\right)^{-\gamma}$ depends on β^E). However, Assumption 2 that $\bar{\varepsilon} \to 0$ means that the effect of β^E on $\widehat{c}(x)$ becomes arbitrarily small. ■

Next, we show that our results on refinancing procrastination in Proposition 2 also continue to hold for all but the limit case of full sophistication.

Proposition 4 (Optimal Stopping)

1. *Adjustment targets m' and b' are independent of β. Thus, $m'(x) = \widehat{m}'(x)$, $b'(x) = \widehat{b}'(x)$, $\underline{m}'(x) = \widehat{\underline{m}}'(x)$, and $\underline{b}'(x) = \widehat{\underline{b}}'(x)$ for all x. These adjustment-target functions may still vary with β^E, though this effect vanishes under Assumption 2. This effectively means that neither β nor β^E affects the adjustment targets.*

2. (a-1) *For $\beta = 1$, the refinancing policy function $\mathfrak{R}(x)$ converges pointwise to $\underline{\mathfrak{R}}(x)$ as the effort cost vanishes. This effectively means that the $\beta = 1$ household's mortgage adjustment behavior does not depend on the state of the effort cost.*

   (a-2) *For $\beta < 1$ and $\beta^E = \beta$ (full sophistication), it is effectively the case that the fully sophisticated household's mortgage adjustment behavior does not depend on the state of the effort cost (details are provided in Appendix D.5.4).*[69]

   (b) *For $\beta < 1$, $\beta^E \in (\beta, 1]$, and $\varepsilon = \bar{\varepsilon}$, $\mathfrak{R}(x) = 0$ for all x. This means that for all but the limit of full sophistication, the present-biased household procrastinates and will not adjust its mortgage when $\varepsilon = \bar{\varepsilon}$.*

   (c) *For $\beta < 1$, $\beta^E \in (\beta, 1]$, and $\varepsilon = \underline{\varepsilon}$, $\underline{\mathfrak{R}}(x)$ converges pointwise to $\widehat{\underline{\mathfrak{R}}}(x)$ as the effort cost vanishes. This effectively means that the present-biased household does not procrastinate when $\varepsilon = \underline{\varepsilon}$.*

Proof. The first step of this proof is to show that full sophisticates do not procrastinate (clause 2a-2). This is an important first step, since partial sophisticates believe themselves

[69] For now, a formal statement of this result is complicated by the fact that a sophisticate's refinancing policy function takes the form of a mixed-strategy equilibrium (details in Appendix D.5.4).

to be fully sophisticated (with a perceived short-run discount factor of β^E) starting next instant. Though we leave details to Appendix D.5.4, the intuition is fairly straightforward. Because a full sophisticate is aware of future selves' behavior, the current self would rather refinance now at cost $\bar{\varepsilon}$ than let future selves procrastinate for long enough that the cost of such procrastination amounts to more than $\bar{\varepsilon}$ of effort. Under Assumption 2 that $\bar{\varepsilon}$ is vanishingly small, this effectively implies that a fully sophisticated household will never procrastinate (conversely, this argument also implies that if $\bar{\varepsilon}$ was not vanishingly small then there would still be scope for a fully sophisticated household to procrastinate).

Next, we turn to clause 1 that the adjustment targets are independent of β and (effectively) β^E. The proof here is in some sense much more difficult than under full naiveté in Appendix C.2, because it is no longer the case that the continuation-value function of a (partially) sophisticated household is identical to that of a $\beta = 1$ household. But, in Appendix D.5.3 we show that the continuation-value function of a (partially) sophisticated household is still an affine transformation of the continuation-value function of a $\beta = 1$ household, building on arguments from Harris and Laibson (2013) and Maxted (2023). Thus, it is still the case that the m' and b' that maximize the value function of a present-biased household will maximize that of the corresponding $\beta = 1$ household, and vice versa. The one slight complication to this proof is that the corresponding $\beta = 1$ household denoted by the hat-notation faces a refinancing effort cost of $H^E \times \bar{\varepsilon}$ instead of just $\bar{\varepsilon}$. However, this effect is trivial under Assumption 2 that $\bar{\varepsilon}$ is vanishingly small.

Third, we discuss clause 2b regarding the procrastination decision of a less-than-fully-sophisticated household when $\varepsilon = \bar{\varepsilon}$. The basic intuition follows from the fact that so long as a household is at least partially naive, then they expect to be less present biased next instant than they are now (i.e., $\beta^E > \beta$). Accordingly, any time that a household with a short-run discount factor of β would want to refinance, a less-present-biased household with a short-run discount factor of $\beta^E > \beta$ would *also* want to refinance. But then, the current self will always choose to procrastinate (again for one instant in expectation) in order to push the effort cost off into the future. In short, we do not need full naiveté to generate procrastination. Rather, all that is needed to generate procrastination is for the current self to think that the next self is less present biased than they are, which is the case for all but full sophistication.

More formally, a partially naive household will consider refinancing only if $\beta v(x) \leq \beta v^*(x) - \bar{\varepsilon}$, or equivalently $v(x) \leq v^*(x) - \frac{1}{\beta}\bar{\varepsilon}$. However, recall that the household also expects that next instant it will have a short-run discount factor of $\beta^E > \beta$. And, a sophisticated household with a short-run discount factor of β^E will refinance whenever $v(x) < v^*(x) - \frac{1}{\beta^E}\bar{\varepsilon}$. Thus, the household perceives that $v(x) \geq v^*(x) - \frac{1}{\beta^E}\bar{\varepsilon}$, which also implies that the current self will never consider refinancing now (since $v^*(x) - \frac{1}{\beta^E}\bar{\varepsilon} > v^*(x) - \frac{1}{\beta}\bar{\varepsilon}$).

Finally, the proofs of clauses 2a-1 and 2c are similar to those in Appendix C.2. ■

Robustness of Results to Partial and Full Sophistication. In the modified model described above in Section D.5.1, we are now prepared to present our main results on the effects of sophistication. Namely, as the effort cost vanishes: *(i) there is a limiting observational equivalence between the behavior of partial sophisticates and naifs, and (ii) in the case of full sophistication ($\beta^E = \beta$), there is a limiting observational equivalence between the behavior of full sophisticates and naifs without any refinancing inertia.* These two statements follow from Corollary 2 and Proposition 4.

Explaining clause 1 in more detail, Corollary 2 implies that the consumption policy function generated by discount-function parameters β and ρ under full naiveté is analogous to that produced by discount-function parameters β' and ρ under partial sophistication. Similarly, Proposition 4 implies that the refinancing policy functions generated by discount-function parameters β and ρ under full naiveté are analogous to those produced by discount-function parameters β' and ρ under partial sophistication. What this tells us is that, if we were to assume that households were partially sophisticated instead of fully naive, then by calibrating the discount function with parameters β' and ρ we would still hit the same LTV moment, the same credit card borrowing moment, and generate the same responses to fiscal and monetary policy as in the full-naiveté case with discount-function parameters β and ρ.

The argument for clause 2 is similar, except that full sophisticates do not procrastinate. Thus, the full-sophistication case can be calibrated to produce behavior that is analogous to the case of naive present bias with no refinancing inertia.

As a final step, we note that the two clauses just discussed only apply in the modified model of Section D.5.1, not the main-text model of Section 2. However, we conjecture that there exist calibrations of the modified model studied here in which the naif's equilibrium behavior in the modified model is comparable to their equilibrium behavior in our main-text model. In particular, by taking $\omega^{\uparrow} \to \infty$ we can approximate the main-text model with a hard constraint at $\underline{b}$. Then, by taking $\lambda^C \to \infty$ we minimize refinancing delays, again as in the main-text model. Under such calibrations, our analysis then implies that a model with partial sophistication will provide comparable predictions about LTVs, credit card borrowing, and consumption responses to fiscal and monetary policy as those in the full-naiveté case. Accordingly, *our analysis in the main text is robust to all but full sophistication.* Similarly, *the full-sophistication case can be viewed as an alternate motivation for the intermediate case of naive present bias without refinancing inertia that was already presented in Appendix D.4*.

Generalization: Comparative Statics for β^E. Proposition 4 implies that procrastination is unaffected by β^E except for the limit case of full sophistication. We emphasize, however, that this is partially due to our simple two-state effort cost in Assumption 1, which can be generalized for additional richness. We summarize one possible extension in the following remark:

Remark 2 *In Proposition 4 above, we get the stark result that households procrastinate homogeneously for all but the limit case of full sophistication. This result can be modified for additional richness by generalizing the structure of effort costs in Assumption 1.*

For example, consider the case where the low-cost state $\underline{\varepsilon}$ is stochastic, such that conditional on drawing an instantaneous low-cost state, the effort cost in that state is given by $\tilde{\underline{\varepsilon}} \sim Uniform[0, \bar{\varepsilon}]$. In this case, households will procrastinate until they draw a low-cost-state in which $\tilde{\underline{\varepsilon}} < \frac{\beta}{\beta^E}\bar{\varepsilon}$.[70] Since $\frac{\beta}{\beta^E}\bar{\varepsilon}$ is increasing as households become more sophisticated, this implies that procrastination decreases as agents become more sophisticated.

D.5.3 Additional Details: Value Functions with Sophistication

Turning to adding further technical details to the arguments above, we now spell out the full system of value functions for fully and partially sophisticated households. Similar to the main text, we will use $v(x)$ to denote a household's (perceived) continuation-value function, and $w(x)$ to denote its (perceived) current-value function.

Step 1: Defining Continuation-Value Function $v(x)$ for Full Sophisticates. We start by expressing the continuation-value function $v(x)$ for a fully sophisticated household. As will be discussed in Step 2, this also provides the (perceived) continuation-value function for a partially sophisticated household, since a partial sophisticate perceives themselves to be fully sophisticated in the next instant with a short-run discount factor of β^E.

Step 1a: Defining $v^{refi}(x)$ for Full Sophisticates. To begin, we need to pin down the continuation-value function from choosing to incur the refinancing effort cost in order to enter the "intermediate" state where the household has filled out its refinancing paperwork, but is waiting for the new mortgage to close. We denote the household's continuation-value

[70] Because a partially sophisticated agent expects their future selves to be fully sophisticated (with short-run discount factor β^E), the current self expects future selves to refinance such that $v(x) = v^*(x) - \frac{1}{\beta^E}\bar{\varepsilon}$. So, by not refinancing the current-value is $w(x) = \beta v^*(x) - \frac{\beta}{\beta^E}\bar{\varepsilon}$. Alternatively, by refinancing the current-value is $w(x) = \beta v^*(x) - \tilde{\underline{\varepsilon}}$. Thus, the current self will refinance whenever $\tilde{\underline{\varepsilon}} < \frac{\beta}{\beta^E}\bar{\varepsilon}$.

function in this intermediate state by $v^{refi}(x)$, and it is expressed as follows:

$$\begin{aligned}\rho v^{refi}(x) = u(c(x)) + &\frac{\partial v^{refi}(x)}{\partial b}(y + rb + \omega^{cc}\left(\max\{b^{-}, \underline{b}\}\right) + \omega^{\uparrow}(b - \underline{b})^{-} - (r^m + \xi)m - c(x)) \\ &- \frac{\partial v^{refi}(x)}{\partial m}(\xi m) \\ &+ \sum_{y' \neq y} \lambda^{y \to y'}\left[v^{refi}(b, m, y', r^m, r) - v^{refi}(x)\right] \\ &+ \sum_{r' \neq r} \lambda^{r \to r'}\left[v^{refi}(b, m, y, r^m, r') - v^{refi}(x)\right] \\ &+ \lambda^{R}\left[v^{R}(x) - v^{refi}(x)\right] \\ &+ \lambda^{C}\left[\max_{b', m'}\left\{v(b', m', y, r + \omega^m, r)\right\} - v^{refi}(x)\right], \text{ s.t. refi. constraint (5) holds,}\end{aligned}$$

subject to the optimality condition $u'(c(x)) = \beta\frac{\partial v^{refi}(x)}{\partial b}$. In the first line of this equation, we use $(b - \underline{b})^{-} = \min\{b - \underline{b}, 0\}$ to denote that the consumer pays the borrowing wedge of ω^{cc} up to $\underline{b}$, and then any debt beyond $\underline{b}$ incurs the higher borrowing wedge of $\omega^{\uparrow}$.

This equation for v^{refi} is similar to equation (8′), but with three main changes. First, in the equation above for v^{refi}, consumption is pinned down by the condition $u'(c(x)) = \beta\frac{\partial v^{refi}(x)}{\partial b}$.[71] This is in contrast to equation (8′), where the time-consistent household with $\beta = 1$ chooses consumption optimally to maximize v.

Second, in the equation for v^{refi} we remove the outer maximization that exists in equation (8′) over whether or not to adjust the mortgage, since the household in this intermediate state is *already* in the process of refinancing. Instead, the equation for v^{refi} adds a new line relative to (8′), which is line six: $\lambda^{C}\left[\max_{b', m'}\left\{v(b', m', y, r + \omega^m, r)\right\} - v^{refi}(x)\right]$. This line captures the household's continuation-value conditional on having its mortgage close, which occurs at Poisson rate λ^{C}.[72] Conditional on closing, the household chooses b' and m' to maximize current-value function $w(x)$. However, since $w(x) = \beta v(x)$, we work directly with $v(x)$ in line six for notational simplicity.

Third and relatedly, the equation for v^{refi} removes the possibility of forced refinancing. We make this simplification since the household is already in the process of adjusting its mortgage.

Step 1b: Continuation-Value Function for Full Sophisticates. Next, we characterize when a sophisticated agent will choose to adjust their mortgage in the typical high-effort-

[71] For further details on the Bellman equation for the continuation-value function of households with present bias, see Harris and Laibson (2013).

[72] As discussed above, by taking $\lambda^{C} \to \infty$ we effectively remove refinancing delays and get back to $v^{refi}(x) \to \max_{b', m'} v(b', m', y, r + \omega^m, r)$.

cost state of $\bar{\varepsilon}$. Let $v^*(x) = \frac{1}{\beta} w^*(x)$ denote the continuation-value from adjusting.

In the typical high-cost state, a sophisticated agent with present bias will adjust their mortgage whenever $v(x) < v^*(x) - \frac{1}{\beta}\bar{\varepsilon}$, and will not adjust whenever $v(x) > v^*(x) - \frac{1}{\beta}\bar{\varepsilon}$.[73] Using this property, we are now prepared to define the continuation-value function $v(x)$ for full sophisticates, following similar methods as in Harris and Laibson (2013) and Maxted (2023). Specifically, the sophisticate's continuation-value function v can be expressed as an HJBQVI, as follows:

$$\rho v(x) = \max \left\{ u(c(x)) + (\mathcal{A}v)(x) \ , \ \rho \left(v^*(x) - \frac{1}{\beta}\bar{\varepsilon} \right) \right\}, \tag{25}$$

subject to the optimality condition $u'(c(x)) = \beta \frac{\partial v(x)}{\partial b}$. For notational compactness, equation (25) uses the same infinitesimal generator notation as in equation (8).

Step 2: Continuation-Value Function for Partial Sophisticates. A partially sophisticated agent perceives themselves to be fully sophisticated in the next instant, with a short-run discount factor of β^E. Thus, a partial sophisticate's continuation-value function is identical to equation (25), except that β is replaced with β^E. That is, the partial sophisticate's continuation-value function v can be expressed as $\rho v(x) = \max \Big\{ u(c(x)) + (\mathcal{A}v)(x) \ , \ \rho \left(v^*(x) - \frac{1}{\beta^E}\bar{\varepsilon} \right) \Big\}$, subject to the optimality condition $u'(c(x)) = \beta^E \frac{\partial v(x)}{\partial b}$.

Step 3: Current-Value Function. We now express the current-value function $w(x)$. In this modified model, current-value function $w(x)$ remains similar to equation (16) in Appendix B.3 (which is itself just the generalized version of equation (9) in the main text), but with one key modification: we update the definition of $w^{refi}(x)$ (i.e., the current-value function from refinancing) because the modified model presented in Section D.5.1 imposes delays between the start of a refinance and its close. Fully, $w(x)$ is now given by:

$$\begin{aligned} w(x) &= \max \left\{ \beta v(x), \ w^*(x) - \bar{\varepsilon} \right\} \quad \text{and} \\ \underline{w}(x) &= \max \left\{ \beta v(x), \ w^*(x) - \underline{\varepsilon} \right\} \quad \text{with} \\ w^*(x) &= \max \left\{ w^{prepay}(x), w^{refi}(x) \right\} \\ w^{prepay}(x) &= \max_{b',m'} \ w(b', m', y, r^m, r) \quad \text{s.t. prepayment constraint (4) holds} \\ w^{refi}(x) &= \beta v^{refi}(x) \qquad \text{where } v^{refi}(x) \text{ is defined above} \end{aligned} \tag{26}$$

[73] Note the rescaling of the effort cost by $\frac{1}{\beta}$, which we explained intuitively in Appendix D.5.2 above and which follows from the first line of equation (26) in Step 3 below.

Step 4: Defining the Comparable $\beta = 1$ Household. We end this subsection by adding further details on mapping the (partially or fully sophisticated) present-biased household to a corresponding household with $\beta = 1$ *and* a refinancing effort cost of $H^E \times \bar{\varepsilon}$. Specifically, equation (25) above helps us to see that we can make use of the following result from Harris and Laibson (2013).

Proposition 5 (Value Function Equivalence) *The continuation-value function $v(x)$ of a fully sophisticated present-biased agent is an affine transformation of the value function of an otherwise-identical agent with $\beta = 1$ and a refinancing effort cost of $H \times \bar{\varepsilon}$, where $H = \left(\frac{\gamma-(1-\beta)}{\gamma}\right)^{-\gamma}$.*

Corollary 3 (Extension to Partial Sophistication) *The continuation-value function $v(x)$ of a partially sophisticated present-biased agent is an affine transformation of the value function of an otherwise-identical agent with $\beta = 1$ and a refinancing effort cost of $H^E \times \bar{\varepsilon}$, where $H^E = \left(\frac{\gamma-(1-\beta^E)}{\gamma}\right)^{-\gamma}$.*

Proof. See Harris and Laibson (2013) and Maxted (2023) for the methods necessary to prove Proposition 5. Corollary 3 then follows, since a partial sophisticate perceives themselves to be fully sophisticated in the next instant with a short-run discount factor of β^E.

Briefly sketching out the proof of Proposition 5, the methods presented in Harris and Laibson (2013) and Maxted (2023) show that the continuation-value function $v(x)$ of the fully sophisticated present-biased agent in equation (25) is equivalent to the value function $\hat{v}(x)$ of a time-consistent (i.e., $\beta = 1$) "$\hat{u}$ agent" with a typical refinancing effort cost of $\frac{1}{\beta}\bar{\varepsilon}$. In particular, this time-consistent $\hat{u}$ agent has a modified utility function of $\hat{u}(\hat{c}) = \frac{\psi}{\beta}u\left(\frac{1}{\psi}\hat{c}\right) + \frac{\psi-1}{\beta}$, where $u(c)$ is the standard CRRA utility function and $\psi = \frac{\gamma-(1-\beta)}{\gamma}$. From inspection one can see that $\hat{u}(\hat{c})$ is an affine transformation of standard CRRA utility, such that $u(c) = \frac{\beta}{\psi^\gamma}\hat{u}(c) + (\text{other constants})$. Thus, by applying this affine transformation to both the $\hat{u}$ utility function and to the $\hat{u}$ agent's refinancing effort costs, we see that the behavior of this $\hat{u}$ agent will be equivalent to the behavior of a "standard exponential agent" with $\beta = 1$, standard CRRA utility $u(c)$, and a typical refinancing effort cost of $\frac{\beta}{\psi^\gamma} \times \frac{1}{\beta}\bar{\varepsilon} = H \times \bar{\varepsilon}$, exactly as stated in the Proposition. ■

D.5.4 Additional Details: Policy Functions with Sophistication

As a final step, we provide further details on the policy functions of households with partially or fully sophisticated present bias.

Consumption. Having argued in Proposition 5 that there is a value-function equivalence (up to an affine transformation) between a present-biased household and the corresponding $\beta = 1$ household, Proposition 3 follows similar arguments as in Maxted (2023).

Mortgage Adjustment (Full Sophistication). The mortgage-adjustment decision of partially sophisticated households was described above in Proposition 4. However, the mortgage-adjustment decision of fully sophisticated households features the added technical complexity that the refinancing policy function becomes a mixed-strategy equilibrium, which we discuss further here.

Specifically, so long as $\varepsilon_t = \bar{\varepsilon}$ then in equilibrium it cannot be the case that a sophisticate adjusts their mortgage with probability-one at any point x. As with naiveté, if the current self believes with certainty that the next self will adjust, then the current self will simply procrastinate. Thus, in equilibrium we instead assume that sophisticates adjust their mortgage at a Poisson rate. Let $\lambda^{adjust} : x \to [0, \infty)$ denote this rate.

The fact that sophisticates adjust their mortgage probabilistically also means that there is something "going on behind the scenes" in the right branch of equation (25). While it is the case that $v(x) = v^*(x) - \frac{1}{\beta}\bar{\varepsilon}$ when the household is in the adjustment region, it is not necessarily the case that this adjustment happens immediately. Rather, in an adjustment region the agent's continuation-value function can be expressed as:

$$\rho v(x) = u(c(x)) + (\mathcal{A}v)(x) + \lambda^{adjust}(x)\left(v^*(x) - \bar{\varepsilon} - v(x)\right), \tag{27}$$

where $\lambda^{adjust}(x)$ is determined in equilibrium precisely to ensure that $v(x) = v^*(x) - \frac{1}{\beta}\bar{\varepsilon}$.[74] Rearranging this equation using the property that $v(x) = v^*(x) - \frac{1}{\beta}\bar{\varepsilon}$ gives:

$$\lambda^{adjust}(x) = \frac{\rho v(x) - u(c(x)) - (\mathcal{A}v)(x)}{\bar{\varepsilon}\left(\frac{1}{\beta} - 1\right)}.$$

On the one hand, when $\bar{\varepsilon}$ is large then one can see how λ^{adjust} could be relatively low and hence that slow refinancing can arise under full sophistication. On the other hand, under Assumption 2 that $\bar{\varepsilon}$ is vanishingly small, λ^{adjust} gets arbitrarily large and hence refinancing occurs arbitrarily quickly.

[74] Alternatively, $\lambda^{adjust}(x) = 0$ whenever $v(x) > v^*(x) - \frac{1}{\beta}\bar{\varepsilon}$ because the household does not want to adjust.

E Supplements to Sections 4

E.1 MPCs and MPXs out of Discrete Wealth Shocks

In Section 4 the MPC and the MPX are defined over infinitesimal wealth shocks. Following Achdou et al. (2021), this section extends these definitions to discrete wealth shocks.

Let $C_\tau(x) = \mathbb{E}\left[\int_0^\tau c(x_t)dt \mid x_0 = x\right]$ denote total expected consumption from time 0 to time τ. Recall that $x = (b, m, y, r^m, r)$. Let $x + \chi$ be shorthand for the vector $(b + \chi, m, y, r^m, r)$, i.e. $x + \chi$ is point x plus a liquid wealth shock of size χ.

For a discrete liquidity shock of size χ the MPC is defined as:

$$MPC_\tau^\chi(x) = \frac{C_\tau(x+\chi) - C_\tau(x)}{\chi}.$$

The MPX is defined as (see Laibson et al. (2021) for details):

$$MPX_\tau^\chi(x) = \left(1 - s + \frac{\nu s}{r_0 + \nu}\right) MPC_\tau^\chi(x) + \frac{s}{\nu + r_0}\left(\frac{\mathbb{E}\left[c(x_\tau) \mid x_0 = x + \chi\right] - \mathbb{E}\left[c(x_\tau) \mid x_0 = x\right]}{\chi}\right).$$

Total consumption $C_\tau(x)$, which is used in the MPC calculation, can be calculated numerically using a Feynman-Kac formula (see Lemma 2 of Achdou et al. (2021) for details). To calculate the MPX we also need to solve for the expected consumption rate at time τ, $\mathbb{E}\left[c(x_\tau)|x_0 = x\right]$. Again, a Feynman-Kac formula can be used to solve for this directly.[75] Numerically, we solve the Feynman-Kac formula for the sample path $r_t = 1\%$ for all t (i.e., no aggregate interest rate shocks) since these calculations are conducted in the steady state.

[75]The Feynman-Kac formula for $C_\tau(x)$ is provided in Achdou et al. (2021). The Feynman-Kac formula for $\mathbb{E}\left[c(x_\tau)|x_0 = x\right]$ is specified slightly differently. Here, $\mathbb{E}\left[c(x_\tau)|x_0 = x\right]$ is given by $\Gamma(x, 0)$, where $\Gamma(x, 0)$ satisfies the PDE $0 = (\mathcal{A}\Gamma)(x, t)$ subject to the terminal condition $\Gamma(x, \tau) = c(x)$.

E.2 Model Solution Details: MPCs

(a) Quarterly MPCs Across Transfer Amounts

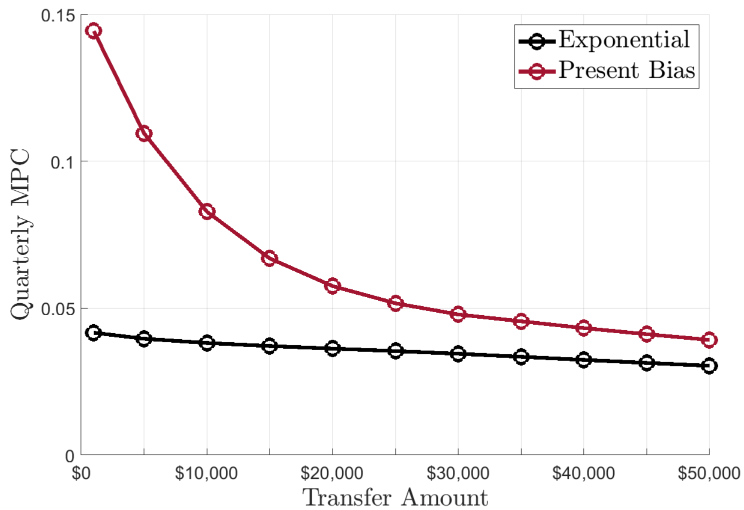

(b) Present-Bias Benchmark: MPCs over Liquid Wealth

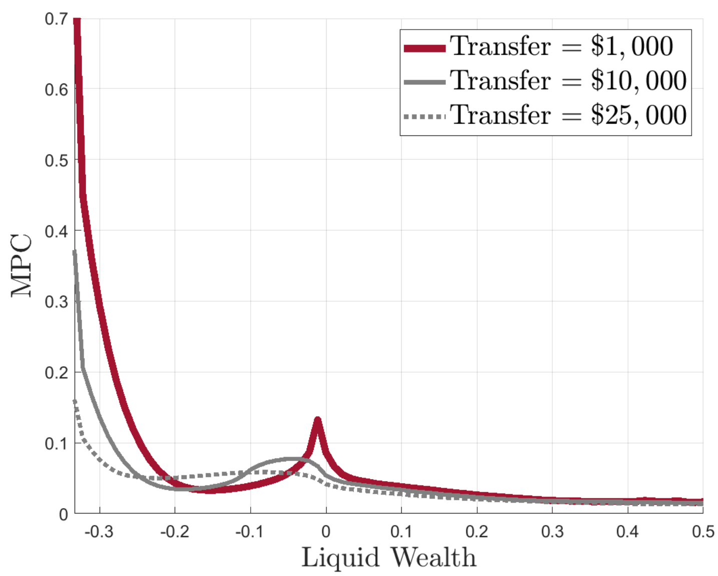

Figure 12: MPCs Across Transfer Amounts.

Notes: The top panel plots quarterly MPCs out of transfers ranging from $1,000 to $50,000 for the two calibration cases. The bottom panel replicates the MPC analysis in Figure 3 for the Present-Bias Benchmark calibration across transfer amounts of $1,000 (benchmark), $10,000, and $25,000.

E.3 Model Solution Details: Steady State Distributions

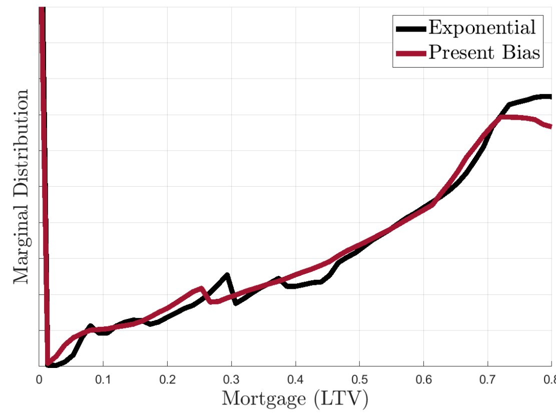

Figure 13: LTV Distribution.

Notes: This figure shows the steady state distribution of households over the LTV ratio.

(a) Exponential Benchmark

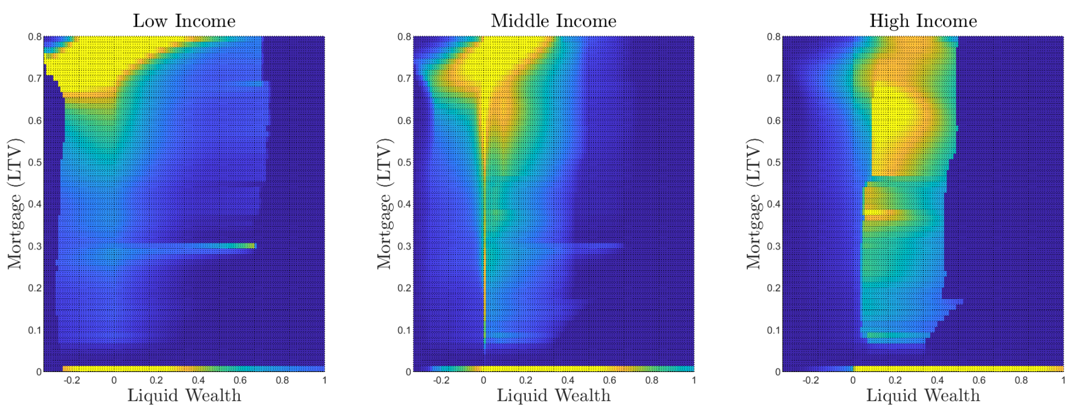

(b) Present-Bias Benchmark

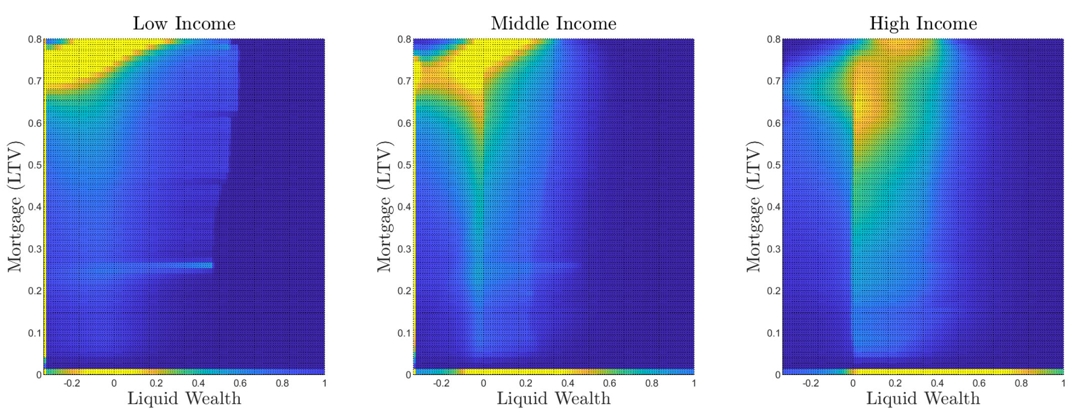

Figure 14: Steady State Distribution.

Notes: This figure presents the full steady state distribution over income, liquid wealth, and mortgage debt. Dark blue regions are rarely encountered, while light yellow regions feature large masses of households.

F Supplements to Section 5

F.1 Fiscal Policy: Financing of Fiscal Stimulus with Future Taxes

As stated in Section 5.1, the government finances the initial \$1,000 stimulus payment with a flow income tax on all households in perpetuity that is chosen so as to satisfy the government budget constraint. This appendix spells out the details of the fiscal rule the government uses to achieve this in our environment with a stochastic interest rate on government debt, building on work by Bohn (1998) and Blanchard (2023).

Before spelling out this fiscal rule, we spell out the government budget constraint. To this end, B_t denotes real government debt in per-capita terms which the government can issue at interest rate r_t, the same stochastic interest rate as that on households' liquid savings. In our continuous-time model, we allow for fiscal stimulus payments not just in the form of continuous flow payments but also as lumpy wealth transfers. To this end, denote by T_t the government's *cumulative* fiscal stimulus from time 0 to time t and denote by dT_t the fiscal stimulus at time t (so that $T_t = \int_0^t dT_s$). Finally, denote by τ_t the tax revenues from a flow income tax that the government uses to pay for such stimulus (more on this below).

With this notation in hand, the flow government budget constraint is:

$$dB_t = (r_t B_t - \tau_t)dt + dT_t. \tag{28}$$

This differential equation needs to hold for all t and for *any* realizations of interest rates, including the possible negative interest-rate realizations (recall from Section 4.1 that interest rates in our calibration follow a Markov process with values $r_t \in \{-1\%, 0\%, 1\%, 2\%\}$). The fact that interest rates can go negative introduces a difficulty with writing the government budget constraint in present value form or, equivalently, with writing the appropriate no-Ponzi condition. To resolve this difficulty we adopt the approach developed by Reis (2021) to analyze present value budget constraints when $r < g$ (here with $g = 0$). Because this difficulty is not central to our choice of fiscal rule satisfying the government budget constraint, we postpone its discussion until the end of this appendix section.

Starting from the steady state at time $t = 0$, the government unexpectedly pays each household a lumpy "helicopter drop" fiscal stimulus payment of \$1,000, i.e. $T_{0^+} = \$1,000$. Thereafter, there are no more fiscal stimulus payments, i.e. $dT_t = 0$ for all $t > 0$ and so also $T_t = \$1,000$ for all $t > 0$. We assume that, in the initial steady state, government debt is zero, $B_{0^-} = \bar{B}$. Given the time path for fiscal policy, government debt initially jumps up to $B_{0^+} = \bar{B} + T_{0^+} = \$1,000$.

Going forward, we then assume that, at each time $t > 0$, the government levies a stochastic tax of τ_t given by

$$\tau_t = \kappa \times r_t B_t, \quad \text{where } \kappa \text{ is a scalar satisfying} \quad \begin{cases} \kappa > 1 \text{ but } \kappa \downarrow 1 & \text{when} \quad r_t \geq 0 \\ \kappa < 1 \text{ but } \kappa \uparrow 1 & \text{when} \quad r_t < 0 \end{cases}. \tag{29}$$

This simple fiscal rule ensures that the government budget constraint is satisfied and that government debt eventually reverts to its initial steady state level. The intuition is as follows.

Suppose first that $\kappa = 1$ so that tax revenues are $\tau_t = r_t B_t$. Examining (28) and recalling that $dT_t = 0$ for all $t > 0$, we then have $dB_t = 0$ for all $t > 0$ and all interest-rate realizations so that the real value of the government's debt stays constant in perpetuity and hence $B_t = B_{0^+} = \$1,000$ all $t > 0$. If all interest rate realizations were strictly positive, then this would be enough to satisfy a standard no-Ponzi condition and hence to satisfy the present-value government budget constraint. However, as just discussed, they are not; in particular r_t can take the values -1% and 0%.

The alternative fiscal rule with $\kappa \approx 1$ as in (29) ensures that the present-value government budget constraint is satisfied even with the possibility of zero or negative interest rates. Intuitively, substituting (29) into (28) and recalling $dT_t = 0$ for all $t > 0$, we have

$$dB_t = (1 - \kappa) r_t B_t dt < 0 \quad \text{for all } t > 0 \text{ and all interest-rate realizations.}$$

That is, at each point in time, the government raises "just a little bit more" tax revenue than the interest payments resulting from the initial stimulus and thus repays a little bit of the initial debt at each point in time, i.e. $dB_t < 0$ for all $t > 0$ and all interest-rate realizations (when $r_t < 0$ we have $\tau_t = \kappa r_t B_t < 0$, i.e. the government makes a transfer to households but one that is a little bit less than the revenues from the negative interest rate on its debt). Therefore $B_T \to 0$ as $T \to \infty$, i.e. government debt eventually reverts to its initial steady state level $\bar{B} = 0$ regardless of the time path of interest-rate realizations.

The fiscal rule (29) is similar to that proposed by Blanchard (2023), who suggests "making the primary balance a function of debt service [...] with one-to-one pass-through" and points out that the rule is a natural extension of the "Bohn rule" (Bohn, 1998) in which the primary balance is an increasing function of the level of debt (rather than debt service).

Given that our model features heterogeneous households, there are some degrees of freedom in specifying how exactly to raise the tax revenues τ_t satisfying (29). In practice, we do this by levying a proportional tax $\bar{\tau}_t \times y_{it}$ on households' inelastically supplied labor income y_{it} so that higher-income households pay higher taxes in dollar terms. Tax revenues are then given by

$$\tau_t = \bar{\tau}_t \times \bar{y},$$

where $\bar{y} = \int y_{it} di$ is average household income (which is constant because we assume a stationary income process).

Finally, it is worth emphasizing just how small the per-period taxes can be while still satisfying the government budget constraint. Since households' average income of $95,718 has been normalized to 1, $\bar{\tau}_t$ is given by:

$$\bar{\tau}_t = \kappa \times r_t \times \frac{B_t}{\bar{y}} \approx r_t \times \frac{\$1,000}{\$95,718},$$

where the approximation uses that $\kappa \approx 1$ and $B_t \approx \$1,000$ because the government pays its debt down only slowly. That is, even in time periods with the highest interest-rate realization $r_t = 2\%$, the tax rate equals only about $\bar{\tau}_t \approx 2\% \times \frac{\$1,000}{\$95,718} \approx 0.02\%$ of labor income. This explains why the initial short-run consumption response with our fiscal rule (29) is very similar in size to the consumption response to the same fiscal stimulus but without imposing a government budget constraint at all.

Present-Value Budget Constraint with Negative Interest Rates (Reis, 2021). As already noted, a difficulty is that our model allows for the possibility of negative interest rates on government debt. We here show how to write an appropriate present-value government budget constraint corresponding to (28), and then show that the simple tax rule (29) satisfies this present-value constraint.

To write the present-value budget constraint, we follow Reis (2021) and use a strictly positive discount rate $\delta > 0$ in place of interest-rate realizations $\{r_t\}_{t\geq 0}$ to compute present values. Imposing the no-Ponzi condition

$$\lim_{T\to\infty} e^{-\delta T} B_T = 0, \tag{30}$$

the present-value budget constraint becomes[76]

$$\int_0^\infty e^{-\delta t} dT_t + B_0 = \int_0^\infty e^{-\delta t}\tau_t dt + \int_0^\infty e^{-\delta t}(\delta - r_t)B_t dt. \tag{31}$$

As discussed by Reis (2021), in principle, *any* discount rate $\delta > 0$ can be used, but a sensible choice is the private return from investing in productive capital (as opposed to government bonds) which, in the data, has historically exceeded the economy's growth rate. In this case (31) has the interpretation of the present-value government budget constraint but discounted at the return of private investors δ. The equation then states that the government's present value of spending dT_t (discounted at δ) plus its initial debt B_0 must not exceed the present value of future taxes τ_t (discounted at δ) plus a non-standard term that Reis names the "bubble premium revenue term" which is the present-value of the implicit government revenues that arise from paying an interest rate r_t on its debt that is below the private return δ (the convenience yield of government debt).

The only thing that remains is to show that the fiscal rule (29) satisfies the no-Ponzi condition (30) and the present-value budget constraint (31). This follows immediately from the fact that this rule implies that government debt converges to zero in the long run, $B_T \to 0$ as $T \to \infty$. In particular, the no-Ponzi condition is immediately satisfied. Similarly, in (31), the bubble-premium revenue term is bounded for any sequence of interest-rate realizations $\{r_t\}_{t\geq 0}$ because $(\delta - r_t)B_t \to 0$ as $t \to \infty$ for any $\{r_t\}_{t\geq 0}$.

Full General Equilibrium Analysis. Finally, we again note that our model is not a general equilibrium model and therefore leaves out a number of considerations that may be important in practice, in particular "Keynesian" multiplier effects of fiscal stimulus that work by stimulating aggregate demand. Future research should explore the effects of fiscal stimulus in full-blown general equilibrium models with present-biased households.

[76] The advantage of using $\delta > 0$ rather than $\{r_t\}_{t\geq 0}$ when writing the no-Ponzi condition and government budget constraint is that it sidesteps the issue that these conditions are ill-defined when r_t can go negative. For example, the standard approach of integrating the flow budget constraint while discounting at $\{r_t\}_{t\geq 0}$ involves the term $\lim_{T\to\infty} e^{-\int_0^t r_s ds} B_T$, which may converge to infinity for negative interest-rate realizations (in particular, one possible history is $r_t = -1\%$ for all t).

To derive (31), write the flow budget constraint (28) as

$$dB_t - \delta B_t dt = ((r_t - \delta)B_t - \tau_t)dt + dT_t.$$

Multiplying by $e^{-\delta t}$ and integrating between 0 and T we have

$$B_T e^{-\delta T} - B_0 = \int_0^T e^{-\delta t} dT_t - \int_0^T e^{-\delta t}\tau_t dt + \int_0^T e^{-\delta t}(r_t - \delta)B_t dt.$$

Taking $T \to \infty$ and using the no-Ponzi condition (30) yields the present-value budget constraint (31).

F.2 Fiscal Policy: Implementation Details

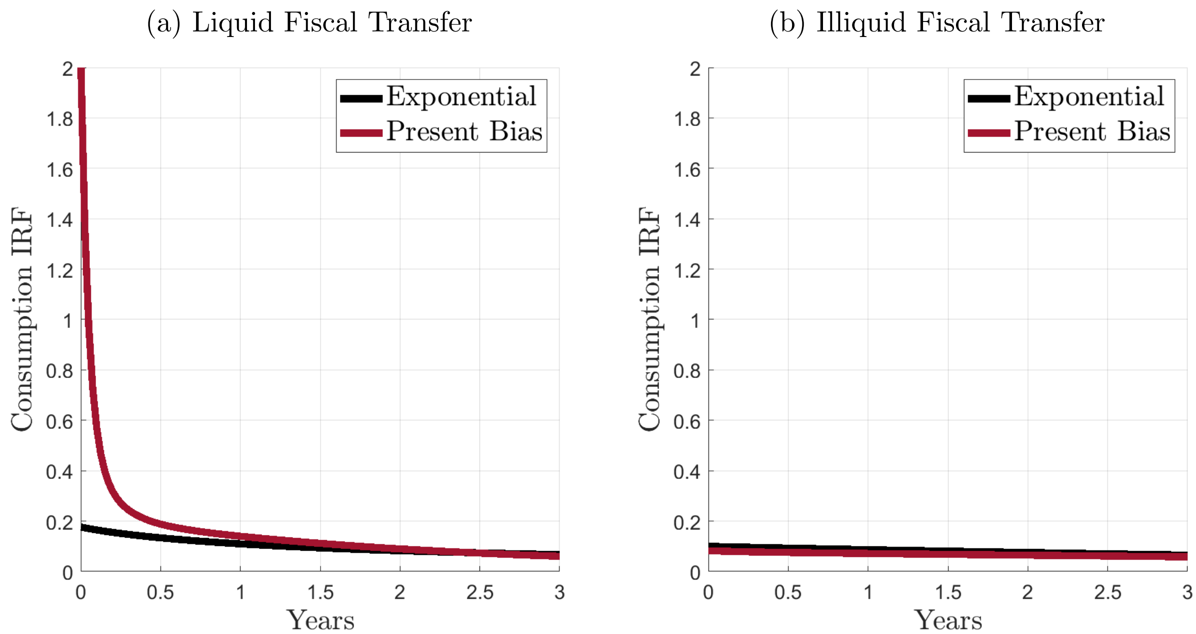

Figure 15: Liquidity of Fiscal Policy.

Notes: The left panel reproduces the benchmark fiscal policy analysis of Figure 4. The right panel plots the IRF of aggregate consumption to a $1,000 mortgage principal reduction.

F.3 Monetary Policy: Refinancing Dynamics

Figure 16 plots the adjustment regions following an interest rate cut from 1% to 0%. This figure replicates the phase diagrams in Figure 1, but now for the case of $r_t = 0\%$ and $r_t^m = 1\% + \omega^m$. Thus, Figure 16 plots the adjustment regions for households with a mortgage rate that is above the rate they can refinance into.

As in the main text, the red regions mark where households take a cash-out refinance and the blue regions mark where households prepay their mortgage. The gray regions indicate where households conduct a rate refinance, defined as the household increasing its mortgage balance by less than 5% during the refinance. Relative to the steady state adjustment regions, the interest rate cut causes the red/gray refinancing regions to expand drastically. In particular, households with larger LTVs are more likely to refinance, since households with larger mortgages have more to gain by reducing their mortgage interest payments.

Table 6 presents details of the refinancing decision. The first row lists the share of households who find themselves in a refinancing region at the time of the interest rate cut. Conditional on refinancing, the second row lists the share of households who extract equity when refinancing. The next four rows list the share of households who have actually refinanced within 1 quarter, 1 year, 2 years, and 3 years following the interest rate cut. While refinancing is instant in the Exponential Benchmark, procrastination means that refinancing occurs slowly in the Present-Bias Benchmark.

	Exponential	Present Bias
Share Refi Region (On Impact)	70.6%	72.3%
(Share Cash Out)	81.2%	77.7%
$\frac{1}{4}$ Year Realized Refi	72.5%	13.2%
1 Year Realized Refi	76.9%	40.9%
2 Year Realized Refi	81.3%	61.2%
3 Year Realized Refi	84.5%	72.0%
Average Refi Amount	0.31	0.28

Table 6: Refinancing Details.

Notes: This table summarizes details of household refinancing following an interest rate cut from 1% to 0%.

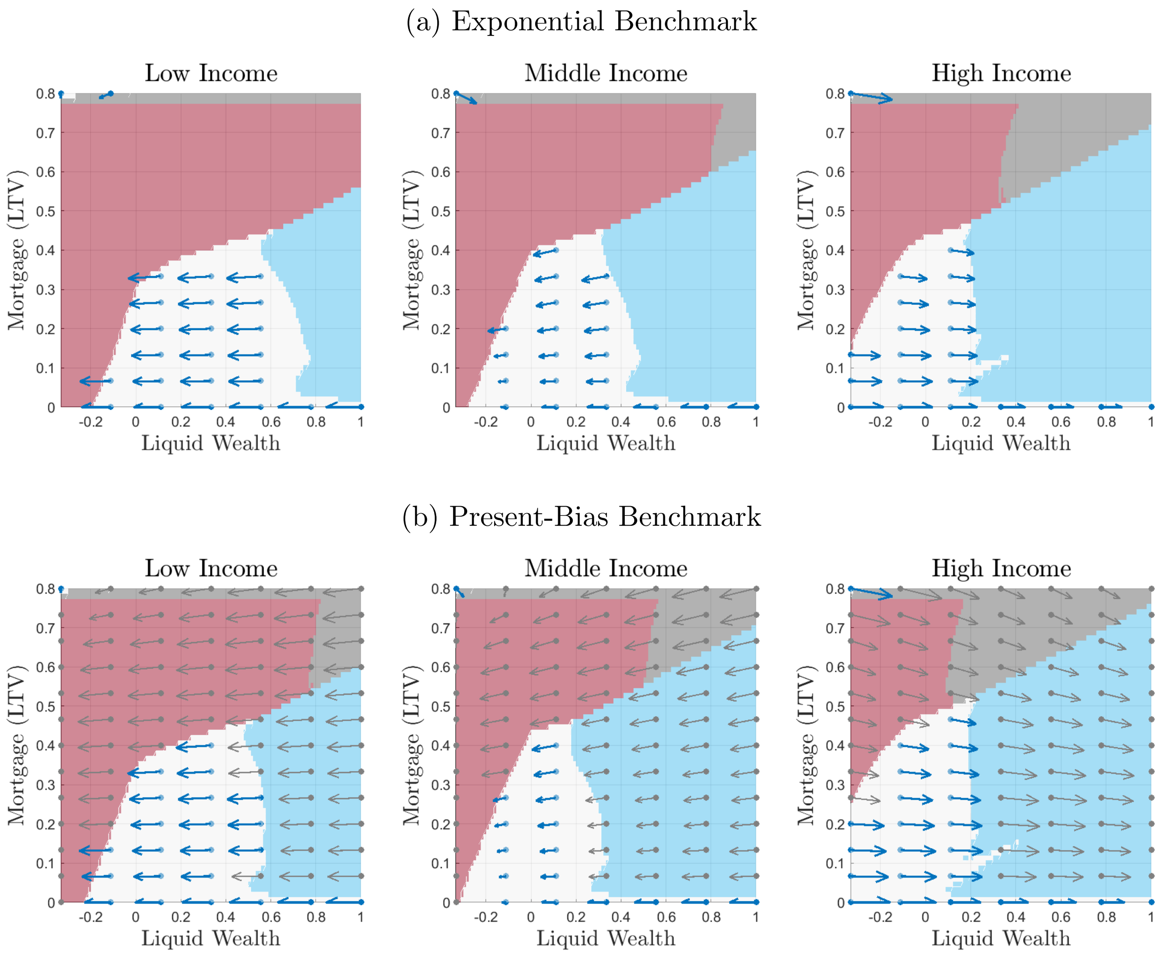

Figure 16: Rate-Cut Phase Diagrams.

Notes: This figure presents the phase diagrams for households who can refinance into a lower mortgage rate following an interest rate cut from 1% to 0% (see Figure 1 for phase diagram details).

F.4 Monetary Policy and Refinancing Procrastination

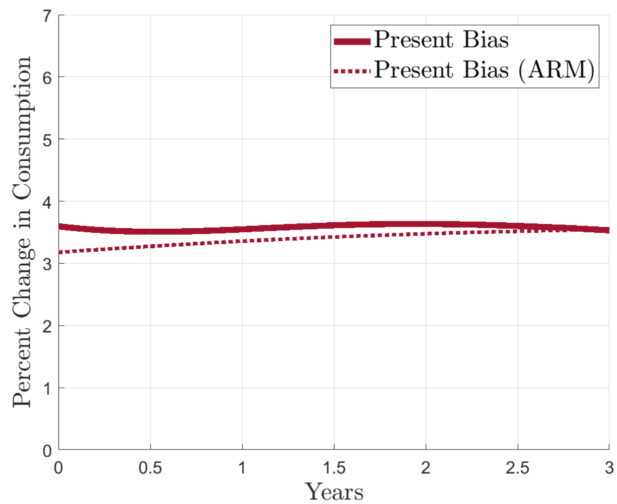

Figure 17: Monetary Policy Under FRMs Versus ARMs.

Notes: For the Present-Bias Benchmark calibration, this figure compares the consumption response to monetary policy under FRMs (solid line) versus ARMs (dotted line). The interest rate is cut by 2% in the ARM experiment, compared to 1% in the FRM experiment, since monetary policy produces larger movements in ARM rates than long-duration FRM rates.

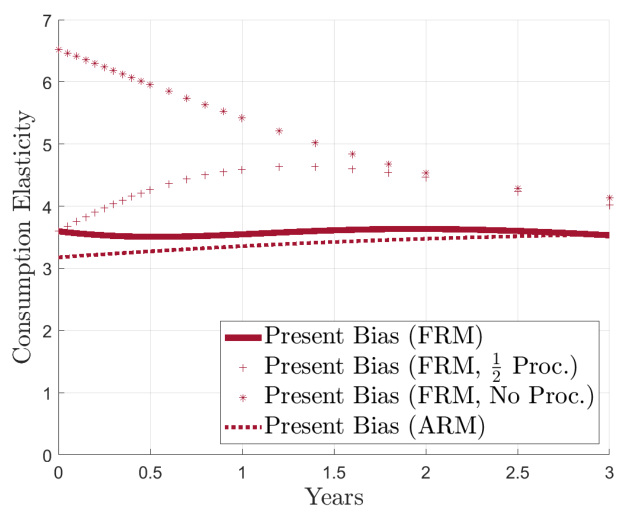

Figure 18: Monetary Policy with Procrastination Reduction.

Notes: This figure presents the consumption response to monetary policy in the Present-Bias Benchmark across varying levels of refinancing procrastination. The +'s assume that policymakers are able to halve the expected duration of procrastination at the time of the rate cut. The ∗'s make refinancing immediate at the time of the rate cut. The baseline consumption response under FRMs (solid red line) and ARMs (dotted red line) are presented for comparison.

G Discussion: Implications of Present Bias for Macroeconomic Policy in Fully-Fledged Macro Models

As we have already discussed, our model is set in partial equilibrium because abstracting from general equilibrium considerations allows for a richer, and more straightforward, investigation of the household problem. We also omit a number of model elements that could be important in principle, such as modeling residential investment. Finally, our calibration focuses on a specific subset of the population, namely homeowners. This raises the question of how present bias would affect the transmission of monetary and fiscal policy in a full general equilibrium analysis that relaxes these assumptions and models the entire population including non-homeowners.

G.1 Omissions from the Analysis and Restriction to Homeowners

We start by discussing certain omissions from the analysis other than general equilibrium effects. In many cases, while filling these gaps would clearly affect our monetary and fiscal policy results, it is less clear as to how doing so would affect our main results about present bias amplifying the impact of these policies (i.e., the comparison of the Exponential and Present-Bias Benchmarks). In others, the omission may materially affect our main results.

Residential Investment. As discussed in Section 2.1, we assume that each household is endowed with a home of fixed value h. That is, the housing size is completely fixed and cannot be adjusted (in contrast, housing *equity* can be adjusted via mortgage balances).

The omission of residential investment means that our model provides an incomplete picture of the effects of monetary policy, particularly as it relates to the spending response to cash-outs. While households often report using extracted home equity for residential investment (e.g., Greenspan and Kennedy, 2008), this channel is broadly missing from our model with fixed housing h. This channel could interact with present bias, thus affecting our main results.[77]

Yet another possibility that we assume away with our fixed-h assumption is that present-biased agents would buy different-sized houses in the first place. On the one hand, present-biased households may struggle to accumulate the liquid wealth required to make a down

[77] That said, the reduced-form MPX tool from Laibson et al. (2021) that we used in Section 4 does capture household spending on home improvements (which differs from residential investment as we explain momentarily). Specifically, we calibrate our mapping from MPCs to MPXs such that the MPX includes consumer spending on "furnishings and durable household equipment" like furniture, household appliances, and gardening equipment. In other words, while our MPX measure excludes residential investment, it does include other types of home improvements (like a new washing machine) that are included as Personal Consumption Expenditures by the BEA.

payment and hence will hold less housing. On the other hand, (partial) sophistication of the type modeled in Appendix D.5 in combination with binding financial constraints may lead present-biased households to buy an illiquid asset like housing as a commitment device (Laibson, 1997; Maxted, 2023), and may hence result in them buying larger houses. If present-biased and exponential households make different housing choices, they may of course also respond differently to monetary policy.

Supply Side of the Credit Card Market. While we calibrate the credit card wedge ω^{cc} to 10.3% in order to match the data on the commercial bank interest rate charged on credit cards, we do not actually model the supply side of the credit card market. An important question for future research is why such high interest rates arise in equilibrium. Default risk is certainly part of the story, though Dempsey and Ionescu (2023) suggest that interest rate spreads far exceed the risk of default. While default *levels* alone may not explain credit card interest rates, it also seems likely that issuers' profits covary positively with the business cycle, since loan loss provisions will generally peak during downturns. This suggests that credit cards are a "high-beta" product for credit card issuers, which would provide a risk-premium explanation for why credit card debt commands an elevated interest rate.

In light of this discussion, another simplification of our model is that it abstracts from credit card default. It is likely that a model with present-biased households will feature a different level of defaults, and perhaps also a different covariance of defaults with the business cycle, than a model with exponential households. Given that this will affect the equilibrium interest rates and borrowing limits that households face, it is again likely that modeling such considerations explicitly will affect our main results.

Supply Side of the Mortgage Market. We also abstract from the supply side of the mortgage market, though many of the questions above still apply. Indeed, we have already shown that present bias affects households' mortgage choices, so it is likely that present bias also influences both the product menus and equilibrium mortgage rates that households face.

Restriction to Homeowners. The subpopulation of homeowners that we calibrated our model to – and hence the individuals that populate our model – differs from the full U.S. population in a number of ways. Perhaps most importantly, our model population overrepresents debtors. More precisely, while our model does feature some households with substantial liquid savings (e.g., $100,000), the number of such households is relatively small (see Appendix Figure 14 and recall that we normalized average income of $95,718 to 1). Instead, most households in our model have substantial mortgage debt, credit card debt, or both.

As a result of these modeling and calibration choices, our results paint only a partial picture of the transmission of monetary and fiscal policy, and of how present bias affects this transmission. That is, we leave out a number of offsetting or amplifying effects of such policy changes. For example, given that our model is mostly populated by borrowers, our discussion of monetary policy in Section 5.2 shows primarily the effects on this subgroup of the population. This likely matters because the consumption of borrowers may respond more to such interest rate cuts than that of lenders due to standard income effects (similarly, borrowers may benefit more in welfare terms). A more fully-fledged macro analysis would also model the other side of the mortgage and credit card markets, in particular the households who lend as well as the financial institutions that facilitate this lending. While fully modeling the entire population would therefore clearly affect the economy's overall consumption response to monetary policy, it is again less clear as to how doing so would affect our main results about the impact of present bias on policy transmission.

G.2 General Equilibrium Effects

We next turn to the question of how present bias would affect the transmission of monetary and fiscal policy in a full general equilibrium analysis. Here we briefly discuss this question through the lens of the literature on Heterogeneous Agent New Keynesian (HANK) models. That is, we ask what the effect of present bias would be on the consumption response to monetary and fiscal policy in a general equilibrium version of the model with nominal rigidities (i.e., as before we focus on the comparison between the Exponential and Present-Bias Benchmarks, but now take into account general equilibrium considerations).[78]

In HANK models, macroeconomic stabilization policy can trigger a number of different indirect general equilibrium effects, particularly effects working through household labor income, asset prices, and returns (see e.g. Werning, 2015; Kaplan et al., 2018; Auclert, 2019; Alves et al., 2020; Slacalek et al., 2020). The size of these indirect effects depends on the size of these variables' movements as well as households' responsiveness to these changes, e.g. MPCs (and MPXs) out of labor income and asset price changes. In heterogeneous-agent models with idiosyncratic income risk and borrowing constraints of the type analyzed here these indirect effects can be important because such models often generate sizable MPCs (e.g., Kaplan et al., 2018).

As we have shown above, present bias increases both households' average MPC and the direct consumption effect of an interest rate cut. The likely implications for the transmission

[78]One could also imagine studying the impact of present bias on the consumption response to macroeconomic policy in models without nominal rigidities, and this may overturn our result that present bias amplifies this response. For example, in a model with a classical dichotomy, changes in nominal interest rates would have no effect on real consumer spending regardless of whether the economy features present bias. Similarly, one may be able to construct a general equilibrium version of our model in which an extreme form of Ricardian equivalence holds so that fiscal stimulus has no effect on consumer spending, again regardless of present bias. We view such exercises as less interesting and instead discuss environments in which policy affects consumer spending also in the absence of present bias.

of monetary and fiscal policy in a full general equilibrium analysis are as follows.

Fiscal Policy. We conjecture that, also in a general equilibrium HANK version of our model, present bias would continue to amplify households' spending response to fiscal policy. This follows from a simple "Keynesian cross" logic which takes as its starting point that the most potent general equilibrium effect triggered by fiscal policy is likely the one working through households' labor incomes: a fiscal transfer increases aggregate consumption demand (the impulse or direct effect); in equilibrium, firms hire more which increases households' labor incomes and leads to additional spending (the multiplier or indirect effect).[79] The key ingredient determining the size of both this impulse and multiplier are households' MPCs, which increase with present bias. Present bias therefore likely amplifies not only the direct effects but also the indirect general equilibrium effects of fiscal policy.

Monetary Policy. We conjecture that the situation is similar for monetary policy, namely that present bias would increase not only direct but also indirect effects and therefore the overall consumption response. Just like fiscal policy, monetary policy triggers indirect effects working through labor income and present bias would amplify these via higher MPCs. Monetary policy can also trigger indirect effects working through asset prices and returns (Gornemann et al., 2016; Kaplan et al., 2018; Alves et al., 2020; Slacalek et al., 2020).[80] However, since present bias does not significantly affect MPCs out of liquid wealth for high-liquidity households (see Figure 3), nor MPCs out of illiquid wealth (see Figure 15), it is natural to conjecture that present bias does little to impact the indirect effects working through asset prices and returns. Taken together, this discussion suggests that, in a HANK-version of our model, present bias would continue to amplify the effects of monetary policy once general equilibrium effects are taken into account.

Fully evaluating the impact of present bias on the economy's response to monetary and fiscal policy in a general equilibrium model is an important task for future work.

[79] For modern macro versions of this mechanism, see for example Auclert et al. (2018) and Wolf (2019).

[80] Since we study the effect of monetary policy on consumption at relatively high frequencies, we are mostly interested in asset price changes at those same frequencies. In this regard, empirical evidence usually points to interest rate cuts as increasing stock prices (e.g., Bernanke and Kuttner, 2005; Gürkaynak et al., 2005). There is also evidence that loose monetary policy increases house prices (e.g., Jordà et al., 2015), but this mechanism seems to operate at a lower frequency than we study.