

Windfall Income Shocks with Finite Planning Horizons

Michael Boutros*

October 2022

Abstract

How do households respond to unanticipated income shocks? I build and estimate a quantitative model of bounded rationality in which reoptimization is costly. Households respond to windfall income shocks by choosing a finite planning horizon over which to reoptimize. The optimal horizon is increasing in income, wealth, and the magnitude of the income shock. In the estimated model, the distribution of consumption responses is consistent with two motivating facts: highly liquid households have large consumption responses out of income shocks that cannot be driven by borrowing constraints, and larger income shocks induce smaller consumption responses.

*Bank of Canada. Email: michael.boutros@bankofcanada.ca. I am grateful to Jason Allen, David Berger, Francesco Bianchi, Laura Castillo-Martinez, Mory Elsaify, Daniel Garrett, Michael Gelman, Cosmin Ilut, Sasha Indarte, Oleksiy Kryvtsov, Andrea Lanteri, John Leahy, Peter Maxted, Alistair Macaulay, David Min, Ben Moll, James Pinnington, Christian Wolf, and participants at numerous institutions for their helpful comments. Version 2022.10.1.

How do households adjust consumption in response to changes in income? Since most shocks directly or indirectly affect household income, this question lies at the center of economics. In macroeconomics, the statistic that often matters is the average propensity to consume:¹

$$\iint PC_i(\Delta) \cdot \partial F_i(i) \cdot \partial F_\Delta(\Delta).$$

This statistic integrates the propensity to consume function, $PC_i(\Delta)$, which measures the consumption response of household i to income shock Δ , over the distributions of households and income shocks, capturing the importance of heterogeneity for the amplification and propagation of macroeconomic shocks (Krueger, Mitman and Perri, 2016).


Standard models built on the Permanent Income Hypothesis generate a consumption response function that depends almost exclusively on each household's wealth. Wealthy households perfectly smooth income shocks and have a near-zero consumption response, while households with low wealth have a propensity to consume near one. Seminal work in Kaplan, Violante and Weidner (2014) and Kaplan and Violante (2014) focuses on measuring the distribution of wealth using high quality micro-data and building models that endogenously match the distribution of low-wealth households with large propensities to consume. But, features of the consumption response function itself are at odds with recent empirical findings that show a significant consumption response for even wealthy households and an important role for the size of the income shock.

In this paper, I develop a new theory of the consumption response function that incorporates bounded rationality in the form of finite horizon planning in response to income shocks. Under full rationality, households adjust lifetime consumption plans in response to every income shock. This level of sophistication does not reflect the constraints that households may face on their ability to reason and make financial plans deep into the future. With bounded rationality, the response to an income shock includes not only new plans, but an intentional decision regarding the extent of reoptimization.

To model this decision, I propose a new constrained-optimal mechanism, bounded intertemporal rationality (BIR), as the household's mechanism to respond to income shocks within an otherwise standard consumption-saving model. The household forms lifecycle plans according to the Permanent Income Hypothesis, but in response to an income shock, reoptimizes and forms new consumption and savings plans over an endogenously selected planning horizon, returning afterwards to its existing plans. This two-layer approach is modeled after the two layers of decision making studied by psychologists:

¹The marginal propensity to consume is the limiting case of the consumption response function as income shocks approach zero. Generally, the propensity to consume may vary with the size of the income shock.

Figure 1: Marginal Propensity to Consume by Liquidity and Shock Size



Notes: (a) Marginal propensity to consume is survey response to hypothetical question. Cash-on-hand is the sum of income and liquid wealth. (b) Estimated marginal propensity to consume out of dividend payments from the Alaska Permanent Fund with 95% confidence intervals. Relative dividend size is payment divided by income. *Source:* (a) 2010 Italian Survey of Household Income and Wealth. Replicates Figure 2 of Jappelli and Pistaferri (2014). (b) Columns (5) and (6) of Table 4 in Kueng (2018).

global and local processing. Local processing is less cognitively demanding than global processing and leads to behavior that oftentimes approximates global processing (Navon, 1977; Forster and Dannenberg, 2010).

I connect the scope of processing to the households optimal planning horizon: a longer horizon corresponds to more global processing. This choice is motivated by empirical evidence in Gabaix, Laibson, Moloche and Weinberg (2006) and Spears (2012) that agents choose finite forecasting horizons to balance the costs and benefits of reducing uncertainty. In the present context, the household selects the optimal planning horizon to trade off the benefits of intertemporal consumption smoothing against planning costs increasing in the horizon, and the optimal horizon will depend on both the shock and the characteristics of the household. Small income shocks endogenously induce local thinking and less intertemporal smoothing, while large income shocks induce global thinking and more consumption smoothing. In the limiting case, a sufficiently large shock induces lifetime reoptimization, as it would in the standard model.

The main contribution of the estimated model is a consumption response function that reflects the observed behavior of households along the entire wealth distribution, especially the majority of households that are liquid and unlikely to be financially con-

strained.² A growing literature finds a significant consumption response for both illiquid and liquid households. For example, panel (a) of Figure 1 plots empirical evidence on self-reported MPCs from Jappelli and Pistaferri (2014). The MPC is decreasing in cash-on-hand, starting at 0.75 for the least liquid households and decreasing to 0.30 for the most liquid households.³ Popular models in this literature cannot generate large consumption responses for liquid households. In both one- and two-asset models built on the Permanent Income Hypothesis, only borrowing constrained households have non-zero consumption responses. Models with hyperbolic discounting generate more constrained households due to higher relative impatience, but the consumption response for unconstrained households is similar to that in the standard exponential model because the “effective discount factor” is approximately equal to exponential discounting for high-wealth households (Harris and Laibson, 2001). In these types of models, the aggregate consumption response misses the contribution of unconstrained households, who constitute the majority of the population.

A second contribution of the model is a consumption response function with a negative intensive-margin effect. Income shocks vary across households in absolute and relative amounts, and there is a large empirical literature documenting the negative relationship between the size of an income shock and the size of the consumption response. For example, panel (b) of Figure 1 plots the estimated consumption response in Kueng (2018), which studies fixed payments from the Alaskan Permanent Fund across households sorted by income.⁴ I estimate the consumption response to Economic Stimulus Payments in 2008 sorted by relative shock size and find a similar negative relationship: the MPC decreases from 0.35 for the smallest relative shocks to 0.12 for the largest. The model also generates a positive extensive-margin effect, consistent with, for example, survey evidence in Fuster et al. (Forthcoming) that households who report ignoring small shocks also report a positive consumption response for larger shocks.⁵ In Section 3, I show that

²Kaplan et al. (2014) estimate that the fraction of households who are not hand-to-mouth, as measured by their balance sheets, is 70% or more in eight advanced economies, including the USA, United Kingdom, and Canada.

³Similar examples of elevated MPCs for highly liquid households can be found in Olafsson and Pagel (2018), Lewis, Melcangi and Pilossoph (2021), Gelman (Forthcoming), Fagereng, Holm and Natvik (2021), Baugh, Ben-David, Park and Parker (2021), and McDowall (2020).

⁴Other examples include Fuster, Kaplan and Zafar (Forthcoming), who use survey questions to identify within-household intensive-margin effects, Fagereng et al. (2021), who use Norwegian lottery winnings and find an intensive-margin effect even when explicitly controlling for liquidity, and Hsieh (2003), who finds that households who have large consumption responses out of small income tax refunds have much smaller consumption responses out of larger dividend payments from the Alaskan Permanent Fund.

⁵Misra and Surico (2014) and Lewis et al. (2021) estimate the distribution of consumption responses out of stimulus payments in 2008 and similarly find a mass of households with no consumption response and a mass of households with a significant response.

standard models generate consumption response functions that are largely inelastic in the size of the shock.⁶ With bounded intertemporal rationality, the consumption response inherits the size elasticity of the optimal planning horizon. Capturing these margins of the consumption response function is important because the aggregate marginal propensity to consume also integrates over the (non-degenerate) distribution of shocks.

The model consists of two layers. In the outer layer, the household forms consumption and savings plans while taking into account typical fluctuations in income, such as seasonal variation in hours worked and temporary spells of unemployment. The outer layer of the model is the standard consumption-savings model and is calibrated using standard values from the literature so that the model's stationary distribution of wealth matches the distribution of liquid wealth in the 2007 Survey of Consumer Finances. In the inner layer, which is the main contribution of the model, the household is subject to an exogenous windfall income shock and reoptimizes over the short-run. The choice of windfall income shock and focus on finite planning horizons are motivated by a literature documenting that households distinguish between typical and windfall income shocks and opt to expend windfall income shocks over relatively short horizons.⁷

The key mechanism in the model is costly reoptimization due to planning costs. When planning costs are zero, the household will always opt to smooth any income shock over the remainder of its lifetime. This choice is trivial because the marginal benefit of smoothing consumption over an additional period is strictly positive under standard preferences. Introducing planning costs generates a meaningful tradeoff between smoothing and planning, which induces shorter planning horizons. Planning costs represent the cognitive costs to form new plans and any cost to adjust away from existing plans. Relative to no planning costs and full rationality in the standard model, the household in my model exhibits bounded intertemporal rationality.

Taking the model to the data, I estimate the planning costs, which are the key driver of the model's dynamics, using the Generalized Method of Moments and Economic Stimulus Payments (ESPs) in 2008. Consistent with the model, households receiving smaller relatively sized payments had the largest consumption responses. Households in the first tercile received an ESP equal to approximately 11% of monthly income and spent all of it

⁶In the seminal two-asset model with fixed transaction costs of Kaplan and Violante (2014), the consumption function is kinked and households exhibit large positive extensive-margin effects around the kink. However, conditional on adjusting, the consumption response function is again inelastic with respect to the size of the shock. In Section 3, I show that the consumption response function in two-asset models with smooth transaction costs, such as Kaplan, Moll and Violante (2018) or Auclert, Bardóczy, Rognlie and Straub (2020), is also size inelastic.

⁷See, for example, Fogel (1999), Auclert, Rognlie and Straub (2018), Gelman (Forthcoming), Fagereng et al. (2021), and McDowall (2020).

within three months of receipt. Households in this group were, on average, both the highest earners and most wealthy. On the other hand, households in the third tercile received an ESP equal to roughly half of monthly income and spent only half of it within three months of receipt. These households had the lowest incomes and least liquid wealth, ruling out the standard borrowing constraint explanation. Using these estimates as targets, I estimate the planning costs and verify their validity using external data from Gelman (2021). The planning costs are non-parametrically estimated but are very closely approximated by an increasing logarithmic function.

In the last part of the paper, I compare the distribution of consumption responses in the estimated model to a household with full rationality in similarly calibrated one- and two-asset models, highlighting the model's main contributions: a consumption response function that captures the behavior of liquid households and has realistic extensive- and intensive-margins with respect to shock size. Applying the model to stimulus transfer programs, policymakers can leverage bounded rationality to maximize aggregate spending. The ideal program targets households based on income and structures payments to induce large consumption responses. In practice, this means designing payments as a percentage of income below the estimated thresholds that induce long-term consumption smoothing. Since these thresholds are higher for lower-income households, the program can be progressive, yet still make smaller transfers to high-income households that are typically excluded from such programs.⁸

Literature This paper adds to a large literature that studies household finance and departures from perfect rationality. Campbell (2006) surveys the literature on household finance and discusses settings in which households make financial decisions that depart from full rationality but can be explained by frictions otherwise ignored in standard finance theory. Focusing on income shocks, Fuchs-Schuendeln and Hassan (2016) survey more than two dozen papers studying the consumption response of income shocks and conclude that “households tend to behave consistently with the Permanent Income Hypothesis when the stakes are high, that is, when dealing with large or repeated changes in their income,” while “for households that are not constrained, near-rationality is a likely candidate to explain their excess sensitivity to small anticipated income changes.” I model finite planning horizons induced by reoptimization costs as the friction that generates near-rational behavior consistent with the empirical evidence.

Cochrane (1989) shows that the welfare penalty of deviating from fully rational con-

⁸This type of transfer program is feasible and resembles existing programs. For example, in 2008 and 2021, US transfer programs were designed to phase out by 5% of adjusted gross income above a certain threshold. The program suggested by the model includes several levels of phase-outs at much lower thresholds, but requires no additional information on recipients.

sumption behavior is typically small, motivating bounded rationality as a means of explaining households who set consumption equal to income (i.e., hand-to-mouth households). Building on this, I focus specifically on bounded rationality with respect to making plans into the future, while Ilut and Valchev (2019) build a model in which the household is boundedly rational with respect to the endogenous state variables. Laibson (1997) and Gabaix (2019) show that larger discounting of the future can generate larger propensities to consume. Relative to these papers, my contribution is to model the endogenous decision of precisely how much to discount the future, and how this decision can vary across households and depending on the shock.

This paper also adds to a large literature that studies household heterogeneity and the propensity to consume out of income shocks. My contribution is a mechanism that focuses on limits to financial planning instead of limits to borrowing. My mechanism generates large propensities to consume for households along the entire distribution of wealth. This is crucial for generating an aggregate marginal propensity to consume in line with the data, with the ultimate goal of studying the macroeconomic impact of household-level heterogeneity (Krueger et al., 2016). In models centered around borrowing constraints, the aggregate marginal propensity to consume depends crucially on the mass of constrained households. In one-asset models, the fraction of households that may be plausibly categorized as constrained using total wealth is too small Jappelli (1990). More recent models have introduced elements to directly or indirectly increase the fraction of constrained households. For example, in the term saving model of Campbell and Hercowitz (2019), households with high wealth are effectively constrained because they have earmarked their savings for a large future expenditure. In the seminal two-asset model of Kaplan and Violante (2014), households may be wealthy in the illiquid asset but constrained as measured by holding of the liquid asset.⁹

At the intersection of these two literatures is a small subset of papers that generate plausible consumption response functions using structural behavioral models (DellaVigna, 2018). Laibson, Maxted and Moll (2021) study monetary and fiscal policy in a two-asset model with present bias, while Lian (2021) develops a more general framework that can accommodate a number of behavioral frictions to generate large MPCs. Most closely related to this paper is McDowall (2020), who builds a model of mental accounts that nests the standard one-asset model with near-zero MPCs at one extreme and a hand-to-mouth consumption model at the other. My use of mental accounts is only to separate windfall

⁹Households with low liquidity may be unable to borrow, or may simply have preferences that generate low levels of liquidity and high propensities to consume (Aguiar, Bils and Boar, 2020; Andreolli and Surico, 2021).

income shocks that require reoptimization from typical income shocks that do not. In his model, preferences over an aversion to saving drives behavior, with higher aversion to saving generating larger MPCs out of any income shock. In my model, the household has standard preferences and chooses the planning horizon taking as given planning costs, and this endogenous decision will cause the MPC to vary depending on the shock.

Outline In Section 1, I detail the baseline model of bounded intertemporal rationality, and in Section 2, I estimate the planning costs that drive bounded intertemporal rationality using data from the Economic Stimulus Act of 2008. Using the estimated model, I discuss the distribution of MPCs in Section 3, and some implications for the design of stimulative fiscal policy in Section 3.3. Section 4 concludes and discusses avenues for future research.

1 Model

In this section, I build a partial equilibrium model of household consumption and savings. The model consists of two layers. The outer layer is a baseline model of consumption and savings in which the household abides by the Permanent Income Hypothesis (PIH). The household forms state-contingent plans over a stochastic stream of labor income. The household anticipates the potential of unemployment and forms precautionary savings, and it may also have plans for vacations, child-bearing, or other large expenses. The inner layer of the model activates when the household faces a completely unanticipated windfall income shock. In response, the household reoptimizes over a finite planning horizon, the length of which is subject to planning costs.

1.1 Baseline Model of Consumption and Saving

I model a household's working life over $T \leq \infty$ periods. The household enjoys consumption, c , and leisure, $\ell = 1 - h$, according to utility function $u(c, \ell)$. The household inelastically supplies a fraction of time, $h \in (0, 1)$, which combines with stochastic productivity, z , to form income, $y = zh$. Consumption and savings plans, c and s , respectively, are jointly formed to maximize lifetime utility. The household's time preference, λ , determines the discount factor, $\beta = (1 + \lambda)^{-1}$. Financial markets consist of a single bond, s_t , that costs R^{-1} per unit and pays one unit in the next period. In this environment, the household chooses state-contingent streams of consumption and saving to maximize

expected lifetime utility,

$$E_0 \sum_{t=0}^T \beta^t u(c_t, \ell),$$

subject to the budget constraint that consumption and saving sum to income and wealth,

$$c_t + R^{-1}s_t = y_t + w_t,$$

and the exogenous borrowing constraint on wealth, $s_t \geq -\underline{b}$.

These plans are fully intertemporally rational because the household uses financial planning to intertemporally smooth lifetime consumption and obtain maximal lifetime utility. If the household were required to pay a planning cost to make long-term plans, then the long-term plans would also be subject to bounded intertemporal rationality. Such an environment could be seen as a microfoundation for the spender-saver model in Mankiw (2000). Households that pay the cost of forming state-contingent plans are the savers that are PIH households, while households with no intertemporal plans live hand-to-mouth in every single period.

The household's problem can be written recursively as

$$V_t(w) = \max_{c_t(x), s_t(x)} u(c_t(x), \ell) + \beta E_t V_{t+1}(w'),$$

subject to the above borrowing and budget constraints. As usual, the budget constraint can be rearranged to reduce the state-space from both income and savings to wealth, w , which evolves according to exogenous income and endogenous savings:

$$w' = y' + s_t(x) = y' + R(x - c_t(x)).$$

In the last period of life, T , the household chooses zero savings, $s_T(w) = 0$, and consumes all wealth, $c_T(w) = w$.

1.2 Windfall Income Shocks

In this section, I model the household's response to a windfall income shock. While abiding by its lifecycle plans, which may account for some stochasticity of income, the household may be subject to a windfall income shock. In response, the household reoptimizes and forms short-term plans to accommodate for the income shock. I model the household's joint choice of the planning horizon, consumption plan, and savings plan, subject to bounded intertemporal rationality on the household in the form of planning costs that

are increasing in the planning horizon. As such, the household still aims to smooth the windfall income shock over many periods, but only up to a limit determined by the trade-off between the benefits of intertemporal smoothing and the planning costs.

1.2.1 Setup and Implementation

I now setup and solve the household's reoptimization problem. Formally, suppose that in some period t , the household learns of an income shock path,

$$\Delta = \{\Delta_t, \Delta_{t+1}, \dots, \Delta_{t+N_\Delta-1}\}.$$

The income shock lasts for N_Δ periods (including period t) and is perfectly anticipated once the household initially learns of it. In response, the household chooses both the length of the reoptimization horizon, k , and new state-contingent consumption and savings plans in each of those periods to maximize its expected lifetime utility:

$$\max_{\{c_\tau, s_\tau\}_{\tau=t}^{t+k-1}} E_t \left\{ u(c_t, \ell - \Phi(k)) + \sum_{\tau=t+1}^{t+k-1} \beta^{\tau-t} u(c_\tau, \ell) + \sum_{s=t+k}^T \beta^{s-t} u(c_s, \ell) \right\},$$

subject to periodic budget constraints and the total borrowing constraint.¹⁰ The household chooses new plans over the k periods of the reoptimization horizon, represented by the first two terms of this expression. The difference between the first and second terms are the planning costs, $\Phi(k)$, that are fully paid in the first period of the reoptimization. These planning costs depend exclusively on the length of the endogenous planning horizon, k .

I assume that the entire income shock is expended over the planning horizon and that the household returns to its original plans after completion of the planning horizon. This is in line with the motivating evidence regarding the short window over which income shocks are spent (Gelman, Forthcoming; Fagereng et al., 2021; Gelman, 2021). As such, in the third term, which includes the periods after the planning horizon, $s \geq t+k$, the household uses the consumption and savings plans it had previously formed. In reality, income shocks are likely not expended exactly over a finite planning horizon, but this assumption makes the model technically tractable without limiting the qualitative mechanism.

The household's problem consists of jointly choosing a discrete planning horizon, k , and new consumption and savings plans over the planning horizon. In the next two sections, I separate the household's reoptimization into two subproblems and describe

¹⁰For ease of exposition, I omit notation indicating that these plans are state-contingent.

each in more detail: first, for a given horizon, the choice of consumption and savings plans, and, second, the choice of the optimal planning horizon.

1.2.2 Windfall Consumption and Savings Plans

In response to a windfall income shock, the household forms short-term consumption and savings plans. In this section, I take as given the reoptimization horizon, k , and study how the household makes new plans for the given horizon. I assume that the household is subject to a one-period positive income shock, $\Delta > 0$. The exposition can easily be extended to multi-period positive or negative shocks.

I denote the household's windfall policy functions by $c_\tau(w, \Delta)$ and $s_\tau(w, \Delta)$ for all periods τ in the reoptimization horizon. This notation emphasizes that in contrast to the initial policy functions, which are defined only over wealth, w , the windfall plans are formed over both wealth and the income shock. However, similarly to the original plans, the new plans are state contingent over the lifecycle income process, y . The household's problem can be further divided into three parts: the first period of the planning horizon, the intermediate periods, and the final period.

First Period of the Planning Horizon At time t , the household faces income shock Δ and also takes as given its wealth prior to the shock, w . Also taking as given the value function from the next period in the reoptimization horizon, $V_{t+1}(w', \Delta')$, the household solves

$$V_t(w, \Delta) = \max_{c_t(\cdot), s_t(\cdot)} u(c_t(w, \Delta), \ell - \Phi(k)) + \beta E_t V_{t+1}(w', \Delta'),$$

subject to the same exogenous borrowing constraint, $s_t(w, \Delta) \geq -\underline{b}$, and the budget constraint

$$c_t(w, \Delta) + R^{-1}s_t(w, \Delta) = w + \Delta.$$

The household bears the leisure costs of reoptimization contemporaneously with the reoptimization itself: in period t , the household incurs the planning cost $\Phi(k)$ as a loss of leisure, $\ell - \Phi(k)$. The planning cost depends only on the planning horizon, k , and I assume that $\Phi(k+1) > \Phi(k)$ for all $k \geq 0$ (which will be verified in the estimation).

To facilitate comparing the original and reoptimized plan, I define excess consumption and excess saving, respectively, as follows:

$$\begin{aligned} c_t^\Delta(w, \Delta) &\equiv c_t(w, \Delta) - c_t(w), \\ s_t^\Delta(w, \Delta) &\equiv s_t(w, \Delta) - s_t(w). \end{aligned}$$

Subtracting the initial budget constraint from the reoptimization budget constraint yields:

$$c_t^\Delta(w, \Delta) + R^{-1}s_t^\Delta(w, \Delta) = \Delta.$$

The intuition for this expression is intuitive: the shock, Δ , is partitioned into excess consumption and excess saving. The allocation between consumption and saving will depend on the household's preferences and, importantly, the number of periods in the reoptimization. Although this notation makes clearer the mental accounting between typical and windfall income, it is still the case that the windfall plans take into account the stochastic lifecycle income process. The windfall income shock smooths consumption over the k periods, but as with wealth, is also used to smooth consumption over different realizations of the stochastic income process in each period.

The value function for the next period, $V_{t+1}(w', \Delta')$, incorporates the household's saving decision. Letting $\Delta' = s_t^\Delta(w, \Delta)$, the household's mental account for the shock evolves according to excess saving:

$$\Delta' = R(\Delta - c_t^\Delta(w, \Delta)).$$

As before, the household's wealth excluding the income shock evolves according to the policy functions that were derived prior to the income shock: $w' = y' + R(w - c(w))$. The key assumption is that the household has different mental accounts for the two forms of wealth even though they are transactionally equivalent.

Intermediate Periods of the Planning Horizon In the intermediate periods $\tau \in \{t + 1, t + 2, \dots, t + k - 1\}$ until the final period of the planning horizon, the household's problem is almost identical to that in the first period. The only difference is that the mental account for the shock is the residual saving after consuming out of the shock in the previous period, and not the exogenous shock as in the first period. In each of these intermediate periods, because of the way the evolution of wealth was defined using the pre-shock policy functions, wealth is the same as it would have been had there been no income shock. Taking as given wealth, w , the mental account for the shock, Δ , and the value function for the next period, the household solves

$$V_\tau(w, \Delta) = \max_{c_\tau(\cdot), s_\tau(\cdot)} u(c_\tau(w), \ell) + \beta E_\tau V_{\tau+1}(w', \Delta'),$$

subject to the total borrowing constraint, $s_\tau(w, \Delta) \geq -\underline{b}$, and the budget constraint

$$c_\tau(w, \Delta) + R^{-1}s_\tau(w, \Delta) = w + \Delta.$$

Again subtracting away the initial budget constraint and using excess consumption and saving, this constraint can be rewritten as

$$c_t^\Delta(w, \Delta) + R^{-1}s_t^\Delta(w, \Delta) = \Delta.$$

As before, the household's wealth excluding the income shock evolves according to the policy function that was derived prior to the income shock, while the household's mental account for the shock evolves according to excess saving:

$$\Delta' = R(\Delta - c_t^\Delta(w, \Delta)).$$

Last Period of the Planning Horizon The final period of the planning horizon is in $t + k - 1$. The household's problem in this period is different due to the assumption that the household expends the entire income shock over the finite planning horizon. The implication of this assumption is that excess saving into the mental account for the shock in the next period, $s_{t+k-1}^\Delta(w, \Delta) = \Delta'$, is zero. Starting in the next period, the household returns to using the value and policy functions that were derived prior to the income shock.

Taking as given wealth, w , the mental account for the shock, Δ , and the value function for the next period, the household solves

$$V_{t+k-1}(w, \Delta) = \max_{c_{t+k-1}(\cdot), s_{t+k-1}(\cdot)} u(c_{t+k-1}(w, \Delta), \ell) + \beta E_{t+k-1} V_{t+k}(w'),$$

subject to the total borrowing constraint and the budget constraint:

$$c_{t+k-1}(w, \Delta) + R^{-1}s_{t+k-1}(w, \Delta) = w + \Delta.$$

Again subtracting the initial constraint and incorporating the assumption that the excess saving in this period amounts to zero, excess consumption is equal to the mental account for the income shock:

$$c_{t+k-1}^\Delta(w, \Delta) = \Delta.$$

That is, in the final period of the planning horizon, the difference between the reoptimized consumption plan and the original consumption plan is the entire balance of the mental account. In the next period, the household no longer mentally accounts for the income shock, and consumption and saving return to their pre-shock levels. Beginning from period $t + k$, the household continues as-if the windfall income shock had never occurred.

1.2.3 Choice of Planning Horizon

In addition to new consumption and savings plans, the household must also choose the length of its planning horizon. The household has access to costly technology that allows it to reoptimize in response to income shocks and to make new plans for a specified number of periods. The optimal planning horizon trades off the benefits of smoothing the shock over an additional period against the planning costs. A longer planning horizon divides the windfall income shock over relatively more periods and induces a smaller consumption response in each period, though the total consumption response is the same for any planning horizon.

The costly reoptimization technology represents the cognitive ability to make new plans and induces bounded intertemporal rationality. At the extreme case of zero planning costs, the household is fully intertemporally rational. In this case, the optimal planning horizon is always the remainder of the household's lifetime because the benefit to consumption smoothing over the remainder of the household's lifetime is always strictly positive. As planning costs become positive and increase in the horizon, the net benefit of smoothing far into the future decreases and the household optimally chooses shorter planning horizons.

Household Optimization Framework In a standard model without reoptimization costs, the optimal choice is a problem solved by the modeler, but in this model, the household's optimization framework is itself an aspect of the model. As such, I must specify how the household chooses the optimal planning horizon.

The optimal horizon trades off benefits of consumption smoothing against planning costs that vary with the horizon. The total planning cost, $\Phi(k)$, is the sum of planning costs for each period,

$$\Phi(k) = \sum_{s=1}^k \phi(k),$$

where $\phi(k)$ is the marginal cost of making plans for the k^{th} period. To find the optimal horizon, the household considers each discrete choice, beginning with the one-period planning horizon in which the household consumes the entire shock contemporaneously, and compares the marginal benefits of consumption smoothing to the associated marginal planning cost.

The household knows the entire schedule of marginal costs, while the benefit for each planning horizon is unknown until the marginal cost is paid and plans for that horizon are formed. If the marginal benefit is greater than the cost, the household continues to the next horizon and pays the next marginal cost. The search terminates when the marginal

cost is equal to or greater than the marginal benefit for the horizon under consideration.¹¹ At this point, the household has found the optimal horizon as well as the associated consumption and savings plans described above.

Optimal Planning Horizon and Shock Size Let $k_t^*(\Delta; w)$ denote the optimal planning horizon for a household at time t with wealth w facing income shock Δ . Proposition 1 states that the optimal planning horizon is increasing in the size of the shock.

Proposition 1. *Consider a household at time t with a given level of wealth, w . If $\Delta' > \Delta$, then $k_t^*(\Delta'; w) \geq k_t^*(\Delta; w)$.*

Proofs of both propositions are in Appendix A. The benefit of consumption smoothing increases with the size of the income shock. If the optimal planning horizon for a small shock is k , then for a larger shock, it will always be at least as beneficial to smooth for that many periods and pay the same planning cost. If the larger shock is sufficiently large, it might even be worth extending the planning horizon and paying a further planning cost. The proof to this proposition depends on the assumption that planning costs depend only on the length of the planning horizon, and I discuss the implications of this simplifying assumption in Section 1.4.2.

Optimal Planning Horizon and Household Wealth Proposition 2 states that the optimal planning horizon is increasing in the household's wealth.

Proposition 2. *Consider a household at time t , facing income shock Δ . If $w' > w$, then $k_t^*(\Delta; w') \geq k_t^*(\Delta; w)$.*

This proposition builds on the fact that the benefits of consumption smoothing over additional periods are increasing in the household's wealth. Since wealthier households benefit more from additional consumption smoothing, they will optimally select a longer planning horizon than a poorer household will for a given income shock. This derives directly from the household's assumed prudence; that is, it derives from the convexity of the marginal utility function (Kimball, 1990). Intuitively, wealthier households have higher consumption and lower marginal utility. It benefits them more to increase consumption marginally over many future periods than to increase consumption by the same total amount but over fewer periods. As the household's wealth and consumption decrease, its marginal utility increases and its returns to smoothing consumption further into the

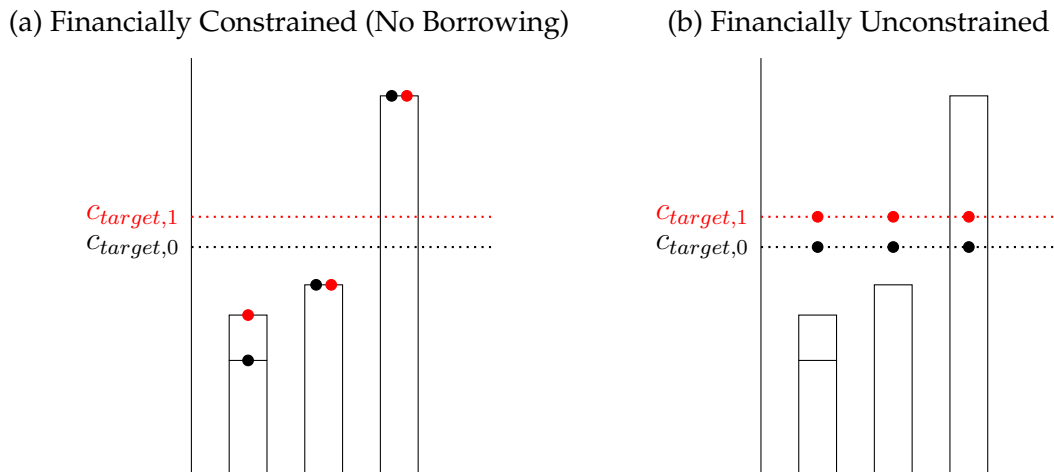
¹¹In a continuous-time model with a choice over continuous-time planning horizons, the household would conclude its search when the marginal costs and benefits are equal. In discrete time, the household cannot know *ex ante* whether the benefit of the next horizon will exceed the cost until it has paid the cost, at which point it becomes optimal to stop searching.

future decrease. Alternatively, one can frame the household with higher wealth as being relatively more patient and, therefore, deriving additional benefits from consumption in later periods relative to a poorer and less-patient household.

1.3 Stylized Examples of Finite Planning Horizons

I develop the intuition for the finite planning horizon model, and especially the second proposition relating the optimal horizon to wealth, with a stylized model of consumption smoothing. Consider first a simple three-period model illustrated in Figure 2. The vertical bars show the upward-sloping income profile of a household that lives for three periods. I assume that preferences and interest rates are such that the household's consumption target is represented by the horizontal line labeled $c_{target,0}$, which is a function of total lifetime income. Actual consumption is given by the solid markers. In the first two periods, the consumption target is greater than income and the household aims to smooth consumption in the current period by borrowing from the future.

Figure 2: Consumption Smoothing in Stylized Model



Notes: Illustrative three-period model of consumption smoothing. Vertical bars depict income, dashed lines represent consumption targets, and markers show actual consumption. Black pattern illustrates initial household behavior and red pattern illustrates new behavior after $\epsilon > 0$ income shock in first period.

In panel (a), I assume the household is financially constrained and cannot borrow. As such, in the first two periods, the household sets consumption equal to income, well below the consumption target. In the final period, the household also consumes all of its income. At the other extreme, in panel (b), I assume the household can borrow without limit. Consumption in each period is equal to the consumption target. In the first two periods, the household borrows from the future to increase consumption above income.

As a result, in the final period, consumption is below income. This is the household's optimal lifetime plan for consumption because of its desire to smooth consumption.


Suppose the household receives an unanticipated $\Delta > 0$ income shock. The household reoptimizes to accommodate for the additional income, increasing the consumption target commensurate to $c_{target,1}$. Consider first the constrained household that was not reaching its initial consumption target. This household opts to consume the entire income shock, bringing it closer to its consumption target in the first period. Because the household is so far from its consumption target in the first period, the benefit from smoothing consumption for one period is greater than the benefit from smoothing consumption for two (or three) periods. Through the lens of my model, the household optimally selects a one-period planning horizon, even before considering the planning costs.

Now consider instead the unconstrained household. Absent planning costs, the marginal propensity to consume is roughly $1/3$ because the household opts to smooth the income shock equally across every period of its life. In each period, the household again meets its (increased) consumption target. In the first period, because of the income shock, the household borrows less than it previously had. In the second period, the household borrows slightly more and, in the third period, the household consumes slightly more of its income. For the unconstrained household, the benefit from smoothing consumption for three periods is greater than the benefit from smoothing for two periods, which is in turn greater than the benefit from smoothing for only one period. When planning costs are introduced, the benefits from additional smoothing remain the same but the household selects a shorter horizon since the planning costs are increasing in the horizon.

The intuition for unconstrained households is similar when the household has a flat income profile. To illustrate this, consider another example, in Figure 3, of a household that lives for 10 periods and faces a constant stream of income. The household consumes its endowment in every period and net saving is zero. Suppose again that the household receives an unanticipated $\Delta > 0$ income shock. Absent planning costs, the household smooths the income shock across every period of its life and its marginal propensity to consume is roughly $1/10$ in each period. With planning costs, the household must choose the optimal planning horizon. Table 1 shows the MPCs over time for each choice of planning horizon in this stylized example. Suppose that the household's optimal choice of planning horizon is four periods, or $k^* = 4$. In the period of the shock and the next three periods, consumption increases by $1/4$ of the income shock. Beginning in the fifth period, consumption returns to its original level, as-if the shock had never occurred.

The mechanism in this simple model extends directly to the model with $T \leq \infty$ periods, stochastic income, and an occasionally binding borrowing constraint. Constrained

Figure 3: Consumption Smoothing in Stylized Model



Notes: Illustrative three-period model of consumption smoothing. Vertical bars depict income, dashed lines represent consumption targets, and markers show actual consumption. Black pattern illustrates initial household behavior and red pattern illustrates new behavior after $\epsilon > 0$ income shock in first period.

Table 1: MPCs with Finite Planning Horizons in Stylized Example

Horizon, k	Cost	MPC in Period				
		1	2	3	4	5
1	$\phi(1)$	1	0	0	0	0
2	$\phi(2)$	$1/2$	$1/2$	0	0	0
3	$\phi(3)$	$1/3$	$1/3$	$1/3$	0	0
4	$\phi(4)$	$1/4$	$1/4$	$1/4$	$1/4$	0

Notes: Stylized example of MPCs and corresponding planning costs for planning horizon $k \in \{0, 1, 2, 3, 4\}$.

households have a desire to smooth consumption by borrowing from the future but are unable to do so; when subject to a positive income shock, they spend a large fraction of it to immediately increase consumption, generating a large marginal propensity to consume. Unconstrained households have a desire to smooth consumption by saving for the future, but planning costs subtract from the benefits of smoothing deep into the future. The household finds it optimal to front-load their consumption of the shock and this generates a larger marginal propensity to consume.

1.4 Discussion of Key Mechanisms

In this section, I discuss the two key mechanisms that generate the two-layered model: mental accounts and planning costs. Together, these features generate a separation between local and global processing. Mental accounting forces the household to reoptimize in response to windfall income shocks, generating two layers of plans, while planning costs for reoptimization are the constraint that bound rationality and govern the degree of local vs. global processing.

1.4.1 Mental Accounting of Windfall Income Shocks

In response to a windfall income shock, the household reoptimizes and forms new consumption and savings plans. The household cannot use its existing plans over its typical income because income is non-fungible; the household distinguishes between typical and windfall income. More generally, fungibility is “the notion that money has no labels” and “in the context of the life-cycle theory, the fungibility assumption is what permits all the components of wealth to be collapsed into a single number” (Thaler, 1990, p. 194). Windfall income shocks are the focus of this paper, while, for example, Ganong, Jones, Noel, Farrell, Greig and Wheat (2020) study “typical income shocks” associated with fluctuations in household income due to firm-level labor demand shocks and Blundell, Pistaferri and Preston (2008) study variation in total income using longitudinal survey data. These typical income shocks are taken into account in the outer layer of the model when the household makes its lifecycle consumption and savings plans.

Windfall income shocks are the focus of a large and established literature in psychology. There are two main criteria for what constitutes a windfall income shock: anticipation and source. Arkes, Joyner, Pezzo, Nash, Siegel-Jacobs and Stone (1994) demonstrate that the unanticipated nature of windfalls is an important part of what separates them from anticipated or typical changes in income. In labeling anticipated shocks as typical or regular income shocks, it is important to note the use of “anticipated” in the economic, not statistical, sense. Statistically, an income shock is anticipated if the household assigns a nonzero probability to its realization. Economically, there are many events that are unanticipated despite having nonzero likelihoods of occurring. One classic example of a windfall income shock is the sudden death of a relative and the associated wealth inheritance. From a technical standpoint, the likelihood of a sudden death and early inheritance is strictly positive, but households neither fully internalize nor make plans for such events. In that sense, the income shock is unanticipated and is labeled as a windfall.

Arkes et al. (1994) and Fogel (1999) present evidence that the source or effort in ac-

quiring additional income is another important determinant of windfall income shocks. They find that earned income is relegated to more utilitarian expenses, while unearned income is spent on more recreational expenses. For example, consider the case where a household earns an additional week of income due to a temporary and unanticipated increase in hours worked or wins a raffle equal to the same amount. In the latter case, the income is treated as a windfall since it is unearned. Relatedly, the labels used to describe an income shock play a role in how they are mentally accounted for. Epley, Mak and Idson (2006) analyze the framing of tax rebate payments and find that referring to them as “bonuses” increases the propensity to consume, which can be attributed to a change in the way respondents mentally account for the extra income. Beatty, Blow, Crossley and O’Dea (2014) study the UK Winter Fuel Payment, a cash transfer with the label “fuel payment” in its name, and find that almost half of the payment was spent on fuel despite the fact that there was no monitoring or enforcement. The authors suggest this is the behavioral effect of labeling and estimate that only 3% of the payment would have been spent on fuel had there been no labeling effect.

Finally, an anticipated but unearned income shock may also be considered a windfall income shock. Again, the classic example of a windfall income shock is the receipt of a wealth bequest after the expected passing of an elderly or ill relative. Despite the anticipated nature of this shock, households treat the income differently from typical income because of the unusual source. Similarly, payments from the Alaska Permanent Fund may be considered annual windfall income shocks, despite the fact that they are “large, regular, predetermined, and salient payments” (Kueng, 2018).

1.4.2 Finite Planning Horizons and Planning Costs

The household, faced with an unanticipated windfall income shock that it differentiates from a typical income shock, must reoptimize and form new consumption and saving plans. The household is boundedly rational. In information processing problems, this means the household faces a cost of processing signals about an unknown quantity. In the present context, bounded rationality imposes a cost on the household’s ability to make new plans. Specifically, I model this as the household choosing the number of periods over which to reoptimize, and I impose a cost based on this choice. Under full rationality in the standard model, there is no cost to making new plans, and the household reoptimizes over the remainder of its lifetime.

I focus on the time dimension of the choice because of the empirical evidence that the consumption response of households to income shocks decays to zero within a short time period. In the US, Parker, Souleles, Johnson and McClelland (2013) estimate that

the total consumption response out of stimulus checks in 2008 was 50-90% within three months of receipt, and Gelman (Forthcoming) estimates that income tax returns were spent in their entirety within six months of receipt. Fagereng et al. (2021) estimate that the consumption response out of Norwegian lottery winnings decays to zero after four years, and Auclert et al. (2018) find agreement for this estimate using Italian survey data. In standard PIH models, even for households with large initial consumption responses, the decay is gradual. In the limiting case of fully unconstrained households, the income shock is annuitized and consumption increases in every remaining period.

Planning Cost Function I model the planning costs as a draw on the household's limited time endowment. Each household is endowed with a unit of time that is initially (exogenously) divided between leisure, ℓ , and labor, h . Planning costs are represented as a function, $\phi(k)$, which depends on the length of the planning horizon, k . Households derive utility from leisure, and planning costs subtract from leisure:

$$\ell = 1 - h - \phi(k).$$

I model these costs as the foregone leisure required to make new consumption and savings plans and interpret them as exerting effort in two broad categories. First, the cognitive effort required in dealing with unexpected changes in income (Browning and Collado, 2001). Here, the household must exert some cognitive effort, which reduces utility, to process new information and make new choices (Reis, 2006; Ergin and Sarver, 2010), especially with respect to financial planning and budgeting (Ameriks, Caplin and Leahy, 2003). The second category is the effort exerted to adjust the household's consumption basket (Chetty and Szeidl, 2007) or savings behavior (Grossman and Laroque, 1990; Huang and Caliendo, 2011; Kaplan and Violante, 2014).

In my analysis, I assume that the leisure cost of forming plans depends only on the length of the planning horizon, k . This assumption is akin to focusing exclusively on the extensive margin of forming plans over a specified horizon. However, both the extensive and intensive margins of planning likely depend on the characteristics of the household, such as preferences or budgeting ability (Ameriks et al., 2003), and the characteristics of the shock, such as its size. I abstract from these factors because I will be unable to account for them in the estimation.

This simplifying assumption is relied upon in the proofs to Propositions 1 and 2, which, respectively, study the optimal horizon as the characteristics of the household (i.e., wealth) and the shock (i.e., size) vary. If planning costs varied with either one, then I would require additional assumptions or restrictions for these proofs. The weakest re-

striction I must make for the main mechanism to remain intact is that high-wealth unconstrained households face planning costs sufficiently high that their optimal planning horizons are shorter than those of households in the standard model. Given that planning costs are zero in the standard model, this requires assuming that high-wealth households face positive planning costs for all shocks. This is a reasonable assumption since although it may or may not be that high-wealth households have an inherent ability for financial planning, the opportunity cost of leisure is increasing in wealth and, thus, planning costs for even high-wealth households are likely net positive.

Finite Planning Horizons and Present Bias To isolate the impact of finite planning horizons on household behavior, I make as few departures as possible from the standard model. It is straightforward to incorporate bounded intertemporal rationality into a model with present bias since the two mechanisms are complimentary.¹² Present bias models have been used to generate larger aggregate consumption responses, but unlike bounded intertemporal rationality, these models will not generate large consumption responses for unconstrained households because the degree of the present bias endogenously and negatively covaries with wealth.

To demonstrate this, Harris and Laibson (2001) derive a generalized Euler equation under hyperbolic preferences and show that a household’s “effective discount factor” is a weighted average between the standard exponential discount factor and the present bias discount factor. The weight on the present bias discount factor is the expected marginal propensity to consume in the next period, which depends exclusively on expected wealth in the next period. In this class of models, wealth is highly persistent. Unconstrained households anticipate continuing to be unconstrained and their effective discount factor places almost all weight on the standard exponential factor. As a result, these households’ consumption responses are small and observationally equivalent to those in the standard model. Constrained households anticipate continuing to be constrained and have effective discount factors that are larger than in those in the standard model, generating even larger consumption responses out of income shocks for constrained households. Altogether, for a given fraction of constrained households, the aggregate consumption responses will be larger than in the standard model, and the aggregate response, again, is driven by constrained households.

Empirically, Gelman (2021) estimates a consumption response function that is in line with predictions from the generalized Euler equation. Gelman presents estimates of the consumption response to a positive income shock for households sorted by quintiles of

¹²Present bias is typically modeled using quasi-hyperbolic discount factors. See E for a detailed discussion.

liquidity. Unconstrained households in the upper quintiles consume evenly across many periods, which is consistent with both present bias and exponential discounting. Constrained households in the lower quintiles unevenly tilt consumption towards earlier periods, which is a telltale sign of present bias. However, the estimated level of the consumption responses for unconstrained households is too large to be explained by either type of discounting on its own, but these responses can be explained by finite planning horizons.

2 Quantitative Model, Calibration, and Estimation

In this section, I bring the model to the data using standard techniques for calibration and estimation. The key object in the estimation are the reoptimization costs that drive bounded intertemporal rationality. After introducing two important quantitative features, I calibrate standard parameters and estimate the planning costs using the Generalized Method of Moments and a natural experiment resembling a windfall income shock. For external validity, I show that behavior in the estimated model is comparable to empirical estimates from Gelman (Forthcoming), an unrelated study of a separate form of windfall income shocks.

2.1 Quantitative Model Extensions

I enrich the model with three additional features that are important quantitatively but do not qualitatively affect the main mechanisms driving bounded intertemporal rationality.

Epstein-Zin Preferences I use Epstein-Zin preferences to separate the roles of risk aversion and the elasticity of intertemporal substitution (EIS). Following Rudebusch and Swanson (2012), I use the following form of recursive preferences:

$$V_t(x) = \max_{c_t(\cdot), s_t(\cdot)} u(c_t(x), \ell) + \beta E_t(V_{t+1}(x')^{1-\alpha})^{\frac{1}{1-\alpha}}.$$

This formulation of recursive preferences is chosen since the kernel for utility includes both consumption and leisure. When $\alpha = 0$, risk aversion and the EIS are inversely related, while a choice of $\alpha > 0$ can yield any combination of risk aversion and EIS. Correctly calibrating the degree of risk aversion is important for generating realistic precautionary saving. As noted by Olafsson and Pagel (2018) and Gelman (2021), and discussed extensively by Aguiar et al. (2020), correctly calibrating the elasticity of intertemporal substitution is crucial for discussion of the marginal propensity to consume. Regardless

of financial constraints, a preference for less intertemporal substitution generates a high propensity to consume and less liquid wealth. If low liquid wealth is used as a proxy for financial constraints, then a researcher may attribute the high propensity to consume to financial constraints, when consumption decisions are based solely on preferences. By separating risk aversion from the intertemporal elasticity of substitution, the model can generate realistic precautionary savings and marginal propensities to consume.

I assume standard separable preferences between consumption and leisure:

$$u(c, \ell) = \frac{c^{1-\gamma}}{1-\gamma} + \frac{\ell^{1+\chi}}{1+\chi}.$$

In line with this literature and my focus on consumption-savings plans, I assume that when making long-term plans, the household inelastically supplies a fraction of its unitary time endowment to labor, h . Leisure is fixed to $\ell = 1 - h$ and the leisure component of utility is irrelevant for the maximization of long-term utility. As such, I do not need to calibrate either the Frisch elasticity of labor supply, χ^{-1} , or fraction of hours worked, h , to solve the long-term problem. However, in forming short-run plans, the household must allocate its time between leisure, labor, and forming plans, and the choice of leisure is endogenous. I will discuss in Section 2.4.2 my choice to estimate the leisure component of utility nonparametrically in order to avoid taking a stand on preferences over leisure.

Differential Saving and Borrowing Rates To further aid in generating a realistic distribution of liquid wealth, I assume that households save and borrow at different rates. Between both mental accounts, if the household is a net borrower, the interest rate is r_{borrow} , and if the household is a net saver, the interest rate is r_{save} .

Default Planning Horizon I modify the optimal horizon selection process to allow for a zero-period planning horizon that is the first choice considered by the household in its optimization framework. With a zero-period planning horizon, the household ignores the shock and freely disposes of it, yielding a marginal benefit of zero. The planning cost is set to zero, $\Phi(0) = 0$, yielding a zero net benefit when the shock is ignored. I label the zero-period horizon as the default behavior because this is the first horizon considered and, if chosen, the household's consumption and saving plans do not change.

This addition allows the estimation to match the empirical finding that households report an inactivity region for smaller positive income shocks, that is, the positive extensive margin effect (Hsieh, 2003; Kueng, 2018; Fagereng et al., 2021; Fuster et al., Forthcoming). Propositions 1 and 2 imply that this structure will yield an inactivity region below some size threshold that is increasing in the household's wealth. The size threshold will depend on the unrestricted estimate for the one-period planning cost. If the estimated one-period

planning cost is sufficiently small, then it is never optimal to dispose of the shock and there will be no inactivity region.

2.2 External Estimates and Calibrations

Income The model is estimated at the monthly frequency and the evolution of income is approximated using a discretized AR(1) process. Gelman (2021) uses monthly transaction-level data for a long panel of households to separate permanent and temporary fluctuations in income. I use his estimates of:

$$y_{it} = (1 - \rho)\mu_y + \rho y_{i,t-1} + \sigma_y \epsilon_{it},$$

in which $(\rho_y, \mu_y, \sigma_y) = (0.883, 0.096, \sqrt{0.039})$.¹³

Preferences Following Kaplan and Violante (2014), I set the annualized discount factor to 0.941, which is similar to the estimated annualized discount factor of 0.935 in Gelman (2021). I also set the coefficient of (constant) relative risk aversion to 4 and the elasticity of intertemporal substitution to $1/2$.

Aguiar et al. (2020) demonstrate that households with high marginal propensities to consume also have high average propensities to consume, implying that their behavior may be driven by preferences in addition to liquidity constraints, and the authors suggest a different calibration for the elasticity of intertemporal substitution. In my baseline specification of the model, I use the same calibration as Kaplan and Violante (2014) to facilitate a comparison. In an alternate specification using the calibration in Aguiar et al. (2020), all households indeed have larger marginal propensities to consume, but the planning mechanism in my model remains crucial for generating a realistic relationship between the propensity to consume and wealth.

Financial Markets Using Table H.15 from the Federal Reserve Board, I calculate that the annualized interest rate on a 3-month certificate of deposit in 2007 was 2.73% and use this as the interest rate for savings. According to the Survey of Consumer Finances, the median interest rate on credit cards was 9.10% and the median credit card borrowing limit was 1.51 times monthly income. I use these as values for the annualized interest rate on borrowing and the borrowing limit, respectively.

Table 2 summarizes the parameters governing the model's long-term plan layer. Figure 4 plots the stationary distribution of wealth in the long-term model compared to the

¹³For more details on the procedure he uses to reach these estimates, see Section 3.3.3 of Gelman (2021).

Table 2: Summary of Long-Term Model Parameters

Parameter	Description	Value	Source
Regular Income Process			
ρ_y	Persistence	0.096	Gelman (2021)
μ_y	Unconditional Mean	0.883	Gelman (2021)
σ_y^2	Variance	0.039	Gelman (2021)
Preferences			
β	Annualized Time Preference	0.941	Kaplan and Violante (2014)
γ	Risk Aversion	4	Kaplan and Violante (2014)
	Elasticity of Intertemporal Substitution	1/2	Kaplan and Violante (2014)
Financial Markets			
r_a	Annualized Saving Rate	2.73%	Federal Reserve Board
r_d	Annualized Borrowing Rate	9.10%	Survey of Consumer Finances (2007)
\underline{a}	Borrowing Limit (\times monthly income)	1.51	Survey of Consumer Finances (2007)


Notes: Summary of the calibrated parameters governing dynamics of outer long-term layer of model.

distribution of liquid wealth from the 2007 Survey of Consumer Finances. Given its parsimony, the model does a fairly good job of fitting the distribution. By construction, the minimal value of wealth in the model is $-1.51 \times$ monthly income, but approximately 10% of households in the Survey of Consumer Finance reported liquid wealth of less than this amount. Similarly, the model is unable to capture roughly the top 10% of the liquid wealth distribution. The model also does not attempt to replicate the mass of households that report holding zero wealth. Despite all of this, the average level of wealth in the model is approximately $0.59 \times$ monthly income, which is approximately the average level of wealth in the data when the lower and upper 10% of the distribution are excluded.

2.3 Model Solution Technique

The model is solved numerically using standard techniques in dynamic programming. The income process is discretized into five gridpoints using the Tauchen method. The baseline model of consumption and saving is solved over 500 gridpoints for wealth. For a given planning horizon, the household solves a finite horizon problem, taking as given the value and policy functions from the baseline model. This finite horizon problem is discretized over the same grid for wealth and an additional grid of 200 points for the windfall income shock. The household's continuation value in the final period of the reoptimization is given by the value function from the baseline model; this reflects that, after the finite planning horizon, the household continues to abide by its initial plans as if the income shock had never occurred. I use a dense grid over lower levels of wealth

Figure 4: Stationary Distribution of Wealth Using Long-Term Plans



Notes: In blue, histogram of liquid wealth to monthly income in the data, censored from below at -1.51 , the borrowing constraint in the model, and from above. Black line shows the stationary distribution of liquid wealth to monthly income in the model. Vertical lines depict the average level of liquid wealth to monthly income in the data and model.

where the value and policy functions have more curvature and I linearly interpolate the value or policy functions off-grid.

2.4 Estimation of Planning Costs

Planning costs are the key driver of the household's short-term response to windfall income shocks. I estimate the planning cost function, $\Phi(k)$, using the Generalized Method of Moments and the consumption response of households to Economic Stimulus Payments in 2008.

2.4.1 The Economic Stimulus Act of 2008

The Economic Stimulus Act of 2008 transferred almost \$100 billion directly into the pockets of households. Economic Stimulus Payments (ESPs) ranged from \$300 to \$600 per

adult, depending on income, and additional payments were made to households with dependents. Parker et al. (2013) use the 2008 wave of the Survey of Consumer Expenditure to estimate that households increased non-durable spending by between 12 and 30 percent of the ESP within three months of receipt. They find that low-income households spent the largest fraction of their ESPs but high-income households spent nearly as much. Reflecting holdings of wealth, they find some relationship between age and homeownership.¹⁴ Shapiro and Slemrod (2009) use an insert in the University of Michigan’s Survey of Consumers to ask households whether they used the majority of their ESPs to increase spending, increase savings, or repay debt. Approximately 20% of households responded that they used the majority of the rebate to increase spending. High-income households most frequently reported that they would spend the majority of their ESPs, but again, the differences between the income groups were small.

Overall, evidence from both revealed and reported preferences suggests violations of the standard PIH model. Borrowing constraints may be part of the explanation, but still cannot account for high propensities to consume of households with high income and/or liquid wealth that are traditionally believed to be financially unconstrained. In Appendix C, I show that the presence of hand-to-mouth households defined in Kaplan and Violante (2014) and Kaplan et al. (2014) increases the number of constrained households theoretically and empirically, but, again, cannot account for high propensities to consume for the remaining and presumably unconstrained households.

Economic Stimulus Payments as Windfall Income Shocks I use the consumption responses of households to ESPs to estimate the model of short-run plans formed over windfall income shocks. Following the discussion in Section 1.4.1, the ESPs meet the criteria to be considered windfall income shocks. These direct payments to households were unanticipated, unearned, and explicitly labelled as “stimulus” payments.

Economic Stimulus Payments moved from idea to implementation in roughly three months, leaving little time for households to anticipate and incorporate them into their lifecycle plans. As detailed in Boutros (2019), Economic Stimulus Payments were suggested by Federal Reserve Chairman Ben Bernanke in a January 17, 2008, speech before the U.S. House of Representatives. Less than one month later, the Economic Stimulus Act was signed into law, and the first payments were distributed in April 2008, less than two

¹⁴Lewis et al. (2021) propose a novel econometric method to study the relationship between household characteristics and the MPC. In their model, instead of *ex ante* grouping households by a given characteristic, they optimally weight households into various groups to maximize model fit. This allows the data to *ex post* reveal underlying patterns between household characteristics and the MPC. Their main findings are that households with high income and/or mortgages have larger MPCs and that households’ MPCs and average propensities to consume are related.

months later.

The IRS distributed payments to all households below certain income thresholds, requiring no opt-in or even knowledge of the program. In his speech, Bernanke suggested that the goals of fiscal policy should be “maximizing the amount of near-term stimulus” and “explicitly temporary . . . to avoid unwanted stimulus beyond the near-term horizon.” The resulting fiscal transfers were explicitly labelled Economic Stimulus Payments and were clearly structured as one-time payments.¹⁵

Sorting Households by Relative Payment Size Payments from the Economic Stimulus Act of 2008 were made to households with joint income of up to \$150,000, almost three times median annual income in the United States.¹⁶ In this section, I present motivating evidence consistent with my model’s prediction that smaller relative income shocks induce less intertemporal smoothing and therefore higher MPCs. I construct Relative ESP by dividing the ESP into either monthly income or cash-on-hand, defined as the sum of monthly income and liquid assets.

The distribution of relative ESPs is driven by variation in both income and ESPs, which may vary due to non-income factors such as household composition. This is an important feature of the data because in the model, the consumption response is driven by both the size of the shock and the household’s income and wealth. In Kueng (2018), which also studies the relation between the consumption response and the household’s characteristics, every household receives the same dividend payment.

Evidence from the Survey of Consumer Expenditures. The 2008 wave of the Survey of Consumer Expenditures asked respondents about the ESP. I use the publicly available Parker et al. (2013) dataset which aggregates responses to the household level. The dataset includes all households in the Survey of Consumer Expenditures that received exactly one ESP. The authors note the data reliability issues with respect to both income and, especially, liquid assets, which roughly half of households in the sample do not report. For more details on how the data is constructed, see Appendix C of Parker et al. (2013).

I divide households into terciles by relative ESP and present summary statistics in Table 3 for the relative ESP, the ESP amount, monthly income, and liquid assets. By construction, the median relative ESP is increasing by tercile, from 12% of monthly income for the first tercile to 41% of monthly income for the third tercile. Households in the first

¹⁵This fiscal program was designed as a stimulus program in the traditional sense: direct payments intended to prop up the economy and avoid a recession. In contrast, for example, Economic Impact Payments distributed in 2020 as part of the CARES Act were distributed after the pandemic-induced lockdown had begun. These, and a second round of transfers in early 2021, were more akin to insurance payments than stimulus.

¹⁶Median (nominal) income was \$52,397 in the 2007 Survey of Consumer Finances.

tercile have the smallest ESPs and most income and liquid assets, followed by households in the second tercile, then households in the third tercile. Using income and liquidity as standard proxies for borrowing constraints, households in the first tercile are those least likely to be financially constrained.

Table 3: CEX Terciles of ESP to Monthly Inc.

		Mean	Std. Dev.	25 th Perc.	Median	75 th Perc.
T1	Rel. ESP	0.109	0.036	0.084	0.115	0.138
	ESP Amount	803	413	600	600	1,200
	Monthly Inc.	7,862	3,903	4,681	7,507	10,310
	Liquid Assets	14,127	21,409	1,600	5,788	17,000
T2	Rel. ESP	0.218	0.034	0.189	0.212	0.245
	ESP Amount	1,023	495	600	1,200	1,200
	Monthly Inc.	4,778	2,360	2,833	4,583	6,393
	Liquid Assets	11,750	23,393	500	2,706	10,000
T3	Rel. ESP	1.187	10.302	0.334	0.405	0.560
	ESP Amount	1,048	575	600	1,030	1,200
	Monthly Inc.	2,398	1,536	1,250	2,000	3,388
	Liquid Assets	5,652	15,169	5	900	4,200

Notes: Summary statistics for households receiving exactly one ESP and reporting annual income, which is divided by 12 to yield monthly income. See Appendix C of Parker et al. (2013) for more details on how the sample was constructed.

Specifically, households in the first tercile received an average ESP of \$803, this group's average monthly income was \$7,862, and both average values were close to their medians. The average level of liquid assets for these households was \$14,127, but the distribution was highly skewed, and the median level of liquid assets was \$5,788. Relative to the first tercile, households in the second tercile had, on average, larger ESPs of \$1,023, less monthly income of \$4,778, and less liquid assets of \$11,750. Households in the third tercile of the relative ESP had the largest ESP payments of \$1,048, the smallest monthly incomes of \$2,398, and the smallest level of liquid assets of \$5,652. Again, in both the second and third terciles the median level of liquid assets was much less than the average.

To estimate the propensity to consume out of the ESP, Parker et al. (2013) regress changes in consumption on the amount of the ESP:

$$\Delta c_{it} = \alpha + \beta \cdot ESP_{it} + \delta \cdot z_{it} + \gamma_t + u_{it},$$

where Δc_{it} is the measured change in consumption for household i between t and $t - 1$,

ESP_{it} is the Economic Stimulus Payment at t for household i , z_t contains changes in family demographics, and γ_t is a monthly fixed effect. The coefficient of interest is β , which measures the propensity to spend out of the ESP in the same month of receipt.

The effect of the stimulus is identified by exploiting the randomized timing of ESP receipts among the non-random sample of households selected to receive these payments. Specifically, households received ESPs (either by check or direct deposit) based on the last two digits of their Social Security Numbers. To identify the causal impact of the ESPs on consumption, I compare consumption at t of households that received their ESPs at t against the consumption of households at t that received their ESPs at $t' \neq t$.

To measure the differential effect across relative ESP terciles, I interact the ESP amount with the relative ESP tercile:

$$\Delta c_{it} = \alpha + \beta_1 \cdot ESP_{it} + \sum_{j=2}^3 \beta_j \cdot ESP_{it} \times 1\{\text{Tercile } j\}_{it} + \delta \cdot z_{it} + \gamma_t + u_{it}$$

I instrument for the ESP amount (and interactions) using an indicator for households that received a payment and estimate the regression equation using 2SLS. Standard errors are clustered by household. The estimated coefficients are presented in Table 4.

Table 4: Spending Response of Consumption to Economic Stimulus Payments

	(a) Non-Durables		(b) Durables		(c) Total	
	Estimate	Implied MPC	Estimate	Implied MPC	Estimate	Implied MPC
ESP (Base: Tercile 1)		0.347** (0.168)		0.715 (0.537)		1.062* (0.576)
ESP \times Tercile 2	-0.137 (0.138)	0.210* (0.120)	-0.081 (0.423)	0.634 (0.392)	-0.217 (0.456)	0.845** (0.424)
ESP \times Tercile 3	-0.232* (0.136)	0.115 (0.109)	-0.260 (0.424)	0.455 (0.347)	-0.492 (0.454)	0.569 (0.368)
Observations		8,592		8,592		8,592
R^2		0.018		0.005		0.007

Notes: Standard errors in parentheses. *, **, *** denote significance at the 0.10, 0.05, and 0.01 levels under the assumption of a single test. Estimated using two stages least squares and instrumenting for ESP amount with an indicator for ESP receipt. See Parker et al. (2013) for more details.

Consistent with the model, the estimated marginal propensity to consume is decreasing in the relative ESP tercile for all measures of consumption. For reference, pooling all terciles together and estimating the baseline regression for nondurable consumption from Parker et al. (2013), the estimated MPC is 0.308. Sorted by relative ESP constructed using monthly income, the implied MPC for the first tercile is 0.347. The implied MPC for

the second tercile is 0.210, which is statistically significant at less than the 5% level, but not statistically different from the implied MPC for the first tercile. The implied MPC for the largest tercile is 0.115, which is not statistically different from zero, but is statistically different from the estimated MPC for the third tercile. A similar pattern emerges for both durable and total consumption, although the estimates are less precise.

Robustness These results suggest that relative ESP size, which takes into account the characteristics of the shock relative to those of the household, is an important determinant of the spending response. From Table 3, the standard deviations of both ESP amount and monthly income are large, and both drive variation in the relative ESP. To ensure this is the case, I estimate the model with households sorted into terciles by income and by $1/\text{Income}$, which is equivalent to assuming that the ESP is constant across households. The results are reported in Appendix B. In both cases, the patterns estimated above for relative ESP disappear. The estimates are largest for the low- and high-income groups, which is consistent with a similar estimation by income in Parker et al. (2013).

The 2008 Consumer Expenditure Survey has limited data on liquid wealth due to high nonresponse rates. In unreported results, I construct relative ESP using “cash-on-hand” defined as the sum of monthly income and liquid assets, and the patterns are largely the same as in Table : average ESP is increasing in terciles, average income is decreasing, and liquid assets are decreasing. The primary difference is that the average relative ESP in each group is much smaller than when the relative ESP is defined using only income in the denominator. In Parker et al. (2013), the estimated consumption responses sorted by liquid wealth are imprecise, and this remains the case when households are sorted by ESP relative to liquid wealth. Using more high-quality data, however, Fagereng et al. (2021) are able to precisely estimate consumption responses for a double-sort by liquidity and shock size. Consistent with the model, they find that conditional on shock size, the consumption response is decreasing in total liquid wealth, and conditional on total liquid wealth, the consumption response is decreasing in shock size.

2.4.2 Estimation Using the Generalized Method of Moments

Using the Generalized Method of Moments, I target the estimated propensities to consume in the regressions above. Since the model is monthly and the CEX estimates of consumption are over three-month periods, I target the cumulative MPC over three months in my model. In total, there are six targets for the MPCs, corresponding to a linear interpolation between the three estimates above. The median and maximum relative ESPs in the first tercile are 11% and 16% of monthly income, respectively, and both are targeted to

yield a cumulative MPC of 0.347. The median and maximum relative ESPs in the second tercile are 21% and 28% of monthly income, respectively, and both are targeted to yield a cumulative MPC of 0.210. The median and 75th percentile relative ESPs in the third tercile are 40% and 54% of monthly income, respectively, and are targeted to yield a cumulative MPC of 0.115. These targets are summarized in Panel A of Table 5.

The household's liquid wealth level in the model is an important determinant of its MPC and therefore is extremely relevant for the estimation procedure. Unfortunately, the 2008 wave of the CEX surveyed households on their liquid assets but did not ask about their liquid debt (i.e., unsecured credit card debt). Instead, I use data on liquid wealth from the Survey of Consumer Finances, merged to the CEX using monthly income profiles. See Appendix F for more details.

Table 5: Summary of Parameter Values (External Estimates and Calibrations)

(a) GMM Targets

#	Description		Target	Model
1	3M MPC for $\Delta = 0.11y$	(50 th Percentile of T1)	0.347	0.349
2	3M MPC for $\Delta = 0.16y$	(100 th Percentile of T1)	0.279	0.279
3	3M MPC for $\Delta = 0.21y$	(50 th Percentile of T2)	0.210	0.207
4	3M MPC for $\Delta = 0.28y$	(100 th Percentile of T2)	0.163	0.165
5	3M MPC for $\Delta = 0.40y$	(50 th Percentile of T3)	0.115	0.116
6	3M MPC for $\Delta = 0.56y$	(75 th Percentile of T3)	0.057	0.078

Notes: Targets for estimation using the Generalized Methods of Moments. Distribution of shock sizes, Δ , and three-month marginal propensities to consume (3M MPCs) are estimated from Economic Stimulus Payments in 2008 (see Section 2.4.1).

(b) External Validation

	$\Delta = 0.33y$		$\Delta = 0.45y$		$\Delta = 0.58y$	
	Data	Model	Data	Model	Data	Model
t = 1	0.083	0.077	0.059	0.045	0.038	0.038
t = 2	0.144	0.153	0.110	0.090	0.075	0.074
t = 3	0.173	0.228	0.138	0.133	0.096	0.110

Notes: Out-of-sample test for external validity of the estimated planning costs. Data columns contain estimates from Gelman (2021) of the one-, two-, and three-month cumulative marginal propensity to consume out of positive income shocks equal to 33%, 45%, and 58% of monthly income, respectively. Model columns contain marginal propensities to consume out of estimated model.

Implementing the Generalized Method of Moments For each target $n \in \{1, 2, \dots, N_{GMM}\}$, I find the planning horizon in the model, k_n^* , that yields the closest cumulative MPC. For

each target n , let $V_n(k, \Phi(k))$ denote the value from choosing horizon k and paying planning cost $\Phi(k)$:

$$V_n(k, \Phi(k)) \equiv \max_{\{c_\tau, s_\tau\}_{\tau=t}^{t+k-1}} E_t \left\{ u(c_t, \ell - \Phi(k)) + \sum_{\tau=t+1}^{t+k-1} \beta^{\tau-t} u(c_\tau, \ell) + \beta^k V_{t+k} \right\}.$$

The utility function is separable between consumption and leisure. In the estimation, I replace the term containing leisure with a scalar, $1 - \theta(k)$:

$$u(c, \ell - \Phi(k)) = \frac{c^{1-\gamma}}{1-\gamma} + \frac{(1-h-\Phi(k))^{1+\chi}}{1+\chi} = \frac{c^{1-\gamma}}{1-\gamma} + (1-\theta(k)),$$

where I use leisure and hours worked, h , and the planning cost must sum to the unit time endowment. I make this change for two reasons. First, this allows me to estimate the planning cost without calibrating the Frisch elasticity of labor supply, χ^{-1} , or hours worked, h . After $1 - \theta(k)$ is estimated, it is straightforward to calculate $\Phi(k)$ for a given calibration of the Frisch elasticity and hours worked. Second, from a technical perspective, the estimation is less computationally intensive when I introduce planning costs in this linear fashion instead of the curvature associated with standard utility over leisure.

To align the model with the targets, I impose a set of conditions such that the value from choosing k_n^* , inclusive of planning costs, is greater than the value from choosing any other $k \neq k_n^*$. That is, for $k \in \{1, 2, \dots, \bar{k}\} \setminus \{k_n^*\}$, the estimation searches for $\Phi_{k_n^*}$ and Φ_k such that

$$V_n(k_n^*, \Phi_{k_n^*}) - V_n(k, \Phi_k) > 0,$$

I implement these inequality constraints as equality constraints using the method described in Moon and Schorfheide (2009). Defining $V_n(k^*, k)$ as the difference in value between the targeted planning horizon, k_n^* , and some other planning horizon, k , this condition can be rewritten as

$$V_n(k_n^*, \Phi_{k_n^*}) - V_n(k, \Phi_k) = \varphi_{n,k},$$

with the parameter restriction $\varphi_{n,k} > 0$ representing the inequality constraint and entering the minimization problem directly. Letting \bar{k} denote the longest horizon considered in the estimation, each target generates one inequality constraint for each horizon other than k_n^* , for a total of $\bar{k} - 1$ inequality constraints per target.


Stacking each of the above restrictions in a vector, the minimization problem can be

written as

$$\min_{\Theta} \frac{1}{2} (V(k^*) - \varphi)' W (V(k^*) - \varphi),$$

where Θ contains the \bar{k} elements of the planning costs and $N_{GMM} \times (\bar{k} - 1)$ inequality constraint parameters are subject to the constraint that each element of φ is strictly positive. Although these conditions alone generally yield a strictly increasing cost function, I find large computational benefits by imposing that the planning cost is strictly increasing, that is, $\Phi_k > \Phi_{k-1}$.

Figure 5: Estimated Planning Costs as Fraction of Monthly Income



Notes: Dashed line plots estimated planning costs and solid line plots quadratic approximation. Shaded area represents confidence interval of estimated planning costs using 95% confidence interval of regressions in Table 4 as estimation targets.

Estimated Planning Costs The estimated planning costs are successful in bringing the model in line with the targets, as listed in Panel A of Table 5. Instead of taking a stance on the Frisch elasticity of labor, I plot the monthly income equivalent of the estimated planning cost in Figure 5. In response to an unanticipated income shock, the household pays just over 4% of monthly income to select even a one-month planning horizon and spends the entire shock. Planning costs increase slowly to around 4.5% of monthly income for one year and 5% of monthly income for three years. In 2008, median household income in the American Community Survey was \$52,029 (in 2008 dollars). Using this as the base-

line for annual income, reoptimizing for one month costs the household the equivalent of \$175, while increasing the planning horizon to one year increases planning costs to \$200.

The consumption-equivalent planning costs are in line with comparable studies in the literature. From a theoretical perspective, Cochrane (1989) demonstrates that only small planning costs are required to push households from perfectly rational to “near rational” behavior. Although the mechanisms behind bounded intertemporal rationality and two-asset models are not directly comparable, both introduce costs that induce less intertemporal smoothing than in the standard model. Reassuringly, the estimated planning costs for finite planning horizons are of the same order as transaction costs in two-asset models. Kaplan and Violante (2014), citing papers that estimate transaction costs on housing and other durable goods, use \$1,000 as the baseline transaction cost for the household to adjust its illiquid assets in response to a shock, which corresponds to approximately 2.1% of average consumption per adjustment. In Kaplan et al. (2018), the steady state transaction costs in the New Keynesian two-asset economy are equal to less than 4%.

In Panel B of Table 5, I perform an out-of-sample test by comparing the estimated model against external data kindly provided by Gelman (Forthcoming). He measures the consumption response of households to their annual tax refunds at a monthly frequency, and I compare the empirical marginal propensity to consume over each of the first three months to the model counterpart. The model is able to match the external targets fairly well, lending external validity to the estimates using the Economic Stimulus Act of 2008.

3 The Distribution of MPCs

With the estimated planning costs in hand, I construct each household’s optimal planning horizon as a function of income, wealth, and the sequence of shocks. I can then calculate the household’s consumption response and marginal propensity to consume. In the model, calculating the average MPC out of Economic Stimulus Payments in 2008 yields an estimate very similar to the baseline regression results discussed in Section 2.4.1. This follows directly from calibrating the long-run model to match the stationary distribution of wealth and estimating the short-run planning cost function using the 2008 Economic Stimulus Payments. Instead, I analyze the model’s distribution of MPCs across wealth and shock size.

3.1 MPC and Liquid Wealth

A large empirical literature documents the tight negative relationship between the MPC and liquid wealth. This relationship is present in standard models: low-wealth households are constrained and have near-one MPCs, while high-wealth households are unconstrained and have near-zero MPCs. However, a growing literature documents both this negative relationship and a significantly non-zero consumption response for unconstrained households (Jappelli and Pistaferri, 2014; Olafsson and Pagel, 2018; Lewis et al., 2021; Gelman, Forthcoming; Baugh et al., 2021; Fagereng et al., 2021; McDowall, 2020). Standard models have limited success in jointly matching these facts jointly, and Fagereng et al. (2021) conclude that “the high average MPC level [they] estimate cannot be explained by liquidity constraints alone.” The model of bounded intertemporal rationality (BIR) presented in this paper can match both facts.


Figure 6 plots the three-month cumulative MPC as a function of liquid wealth for an income shock equal to 28% of monthly income, the average relative size of an Economic Stimulus Payment in 2008. The black line in the figure is from the estimated BIR model developed in this paper. The red line represents the MPC calculated from a one-asset model that is calibrated the exactly same as the long-term model in Section 2.2. The two blue lines are the MPC calculated from the two-asset model developed in Auclert et al. (2020). In the two-asset model, the household can freely invest in a liquid asset or pay a transaction fee each time it adjusts its illiquid asset. The dashed blue line is the MPC for a household with low illiquid wealth and the dotted blue line represents a household with high illiquid wealth.¹⁷

In all four cases, the MPCs for households with low liquid wealth is high. As discussed in Section 1.2.3, households near their borrowing constraint have an unmet desire to smooth consumption by borrowing from the future. Faced with a positive income shock, they opt to increase consumption in the current period, generating large MPCs. This is true for all four models. As liquid wealth increases, the MPC decreases, but much more quickly in the one- and two-asset models. In these models, unconstrained households smooth the positive income shock over their entire lifetimes, consistent with the Permanent Income Hypothesis. The household saves most of the income shock in order to fund its increased consumption in every future period, generating a small marginal propensity to consume out of the shock.

In the BIR model, the MPC decreases more slowly because wealthier households opt

¹⁷Recall that the innovation in the two-asset model is that households with low liquid wealth behave similarly regardless of their illiquid wealth. For this reason, the two lines from the two-asset model are similar.

Figure 6: Marginal Propensity to Consume and Liquid Wealth




Notes: Marginal propensity to consume out of an income shock equal to 28% of monthly income, the average relative size of an Economic Stimulus Payment in 2008.

to smooth the income shock over relatively fewer periods. This is due to the combination of diminishing returns to consumption smoothing and increasing costs in the planning horizon. Again consistent with the Permanent Income Hypothesis, the wealthy household wishes to smooth the income shock over future periods, but doing so now incurs the planning costs, and this tradeoff induces shorter planning horizons. As a result, the BIR model generates MPCs for unconstrained households that are still smaller than for constrained households, but much more in line with the empirics.

3.2 MPC and Shock Size

In contrast to the relationship between the MPC and liquid wealth, there are far fewer investigations into the relationship between the MPC and shock size. Fuster et al. (Forthcoming) design and implement a survey intended to explicitly study the extensive- and intensive-margin “size effects” of the consumption response function. They find strong evidence of a positive extensive-margin effect and a negative intensive-margin effect. That is, the consumption response is zero for small shocks and positive for all shocks above some threshold, but then decreases towards zero as the shock size increases. Kueng

Figure 7: Marginal Propensity to Consume for Different Shock Sizes and Wealth Levels



Notes: Marginal propensity to consume out of an income shock ranging from 0% to 100% of monthly income.

(2018) studies payments from the Alaskan Dividend Fund and finds evidence of a negative intensive-margin effect: households for which the payments are smaller fractions of income have larger consumption responses. Fagereng et al. (2021) use Norwegian lottery winnings and find a negative intensive-margin effect even when explicitly controlling for liquidity. A notable exception is Andreolli and Surico (2021), who find a positive intensive-margin effect: on average, the same household reports consuming marginally more of an income shock equal to one year of income than an income shock equal to one month of income.

The BIR model generates a positive extensive-margin effect and a negative intensive-margin effect. To be clear, the former is a mechanical result of the “default choice” of ignoring small shocks, although the threshold at which the extensive margin activates is part of the estimation. In Figure 7, I vary the size of the income shock to between 1% of 100% of monthly income and compare the MPC of a relatively constrained household in panel (a) to a relatively unconstrained household in panel (b). In both cases, the household opts not to smooth very small income shocks and the MPC is zero. As the size of the income shock increases, the optimal planning horizon increases and the MPC decreases.

Constrained households in all three models have large MPCs. In my model, the operative mechanism is bounded intertemporal rationality for small shocks and the financial

constraints channel for large shocks. As such, behavior in the three models is distinct for smaller shocks, but behavior of the BIR household resembles that of the other models for larger shocks. Specifically, the constrained BIR household chooses shorter planning horizons for small shocks, generating MPCs between 0.50 and 1.00. In the standard one- and two-asset models, the constrained household's MPC is between 0.40 and 0.50 for small shocks. As the size of the shock increases, the MPC decreases to between 0.25 and 0.35 in all three models. For larger shocks, even the constrained household in the BIR model is driven by its unmet desire for consumption smoothing, and it therefore spends a larger fraction of the income shock.


Unconstrained households in the BIR model have much larger MPCs than households in either the one- or two-asset models. In those standard models, the household costlessly smooths any income shock and the MPC is roughly 0.10. In contrast, the BIR household opts to partially smooth income shocks, generating a distinct pattern of MPCs. For income shocks up to 40% of monthly income, the benefits of consumption smoothing are dominated by the planning costs and the household selects shorter planning horizons. As the size of the income shock increases, the household is more willing to reoptimize over additional periods but the MPC is still larger than in the one- and two-asset models. Eventually, for a sufficiently large shock, the unconstrained BIR household opts to pay the planning cost and fully smooth the income shock, and the BIR household's behavior resembles that of the other models.

3.3 Implications for the Design of Stimulative Fiscal Policy

The relationship between the marginal propensity to consume and the size of an income shock provides a framework for the design of stimulus programs intended to boost consumption: smaller payments, relative to a household's income, induce less intertemporal smoothing and larger immediate increases in consumption. Panel (a) of Figure 8 plots the distribution of relative Economic Stimulus Payments (divided into monthly income) in 2008. As estimated in Section 2.4.1, households in the left-hand tail of the distribution had larger marginal propensities to consume, despite earning more income and holding more liquid wealth.

In panel (b), I outline an alternative stimulus program that distributes smaller relative payments to households. The program has three tiers that partition the population into low-, medium-, and high-income households. High-income households receive payments equal to a small fraction of income, between 5% and 10%. Medium-income households receive intermediate-sized payments ranging from 11% to 20% of income, while

Figure 8: Comparison Between Fiscal Stimulus Programs



Notes: In panel (a), the distribution of Economic Stimulus Payments relative to monthly income, from the 2008 wave of the Survey of Consumer Expenditures. In panel (b), the distribution of stimulus payments relative to monthly income for a hypothetical stimulus program.

low-income households receive larger payments ranging from 21% to 30% of monthly income.

As noted extensively by Kaplan et al. (2014), the correlation between income and liquid wealth is positive but not very strong. Since high-income households are the most likely to be financially unconstrained, my hypothetical program aims to induce limited consumption smoothing in these households by targeting them with smaller payments. Medium- and low-income households are targeted with slightly larger payments because their thresholds for consumption smoothing are larger. A large payment to a low-income household, which is more likely to be financially constrained, will generate a larger consumption response than for a similarly sized payment to a medium- or high-income household.

In Kaplan and Violante (2014), the pioneering use of a two-asset model to study stimulative fiscal policy also suggests that smaller payments induce higher propensities to consume. However, the mechanisms in the two models are different, ultimately leading

to different policy conclusions. Their ultimate finding is that “the aggregate consumption response is the largest when the policy is phased out around median income” (Kaplan and Violante, 2014, p. 1235), thus excluding half the population from participating in the program. This is because their mechanism works exclusively through the liquidity channel. Despite the presence of high-income households with low liquid wealth and large marginal propensities to consume, lower-income households are more likely to hold low liquid wealth. As such, in a world in which liquid wealth is not observed by the policy-maker, the two-asset model suggests targeting low-income households.

In my model and alternative stimulus plan, low-income households are targeted because of the same liquidity channel as in the two-asset model. Medium- and high-income households are targeted with smaller payments because of their bounded intertemporal rationality; even if these households do not hold low liquid wealth, the small payments will induce little consumption smoothing and a large marginal propensity to consume. Instead of phasing out the policy completely around median income, the model suggests a more gradual decrease in relative payment size that can extend to more of the high-income population.

4 Conclusions

I develop a model of consumption behavior in which households form consumption and savings plans over stochastic fluctuations in income but reoptimize in response to unanticipated windfall income shocks. My estimated model produces results that are consistent with two motivating facts: the large consumption response out of income shocks for unconstrained households and the negative relationship between the size of the consumption response and the size of the income shock.

I label the households in my model as displaying bounded intertemporal rationality because although they are fully rational, their ability to make plans for intertemporal substitution is bounded by the presence of planning costs. Absent these costs, my model collapses to the standard one-asset model with full consumption smoothing for unconstrained households. Financially constrained households immediately spend positive income shocks because of an unmet desire to smooth consumption, while even unconstrained households have high marginal propensities to consume because they opt to only partially smooth income shocks. For both types of households, the larger the shock, the stronger the incentive for consumption smoothing and the smaller the marginal propensity to consume.

My contribution to the literature is a partial equilibrium framework that focuses on

income shocks and produces plausible marginal propensities to consume along the entire distribution of wealth. More realistic consumption responses at the micro level generate more realistic aggregate marginal propensities to consume. This allows for macroeconomic models that can better understand the dynamic and distributional effects of shocks and can aid in designing policies to maximize aggregate welfare. Future work in this area will extend in two directions: first, it will expand the framework to analyze other shocks, such as to interest rates or borrowing limits; and second, it will embed bounded intertemporal rationality into a broader framework to fully examine the effects of fiscal and monetary policy in general equilibrium.

References

- Aguiar, Mark, Mark Bils, and Corina Boar**, “Who Are the Hand-to-Mouth?,” Working Paper 2020.
- Ameriks, John, Andrew Caplin, and John Leahy**, “Wealth Accumulation and the Propensity to Plan,” *The Quarterly Journal of Economics*, Aug. 2003, 118 (3), 1007–1047.
- Andreolli, Michele and Paolo Surico**, “LESS is More: Consumer Spending and the Size of Economic Stimulus Payments,” *CEPR Discussion Paper No. DP15918*, Mar 2021.
- Arkes, Hal R., Cynthia A. Joyner, Mark V. Pezzo, Jane-Gradwohl Nash, Karen Siegel-Jacobs, and Eric Stone**, “The Psychology of Windfall Gains,” *Organizational Behavior and Human Decision Processes*, September 1994, 59, 331–347.
- Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub**, “Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models,” Working Paper 2020.
- , **Matthew Rognlie, and Ludwig Straub**, “The Intertemporal Keynesian Cross,” Working Paper 2018.
- Baugh, Brian, Itzhak Ben-David, Hoonsuk Park, and Jonathan A. Parker**, “Asymmetric Consumption Smoothing,” *American Economic Review*, Jan. 2021, 111 (1), 192–230.
- Beatty, Timothy K. M., Laura Blow, Thomas F. Crossley, and Cormac O’Dea**, “Cash by any other name? Evidence on labeling from the UK Winter Fuel Payment,” *Journal of Public Economics*, Oct. 2014, 118, 86–96.
- Blundell, Richard, Luigi Pistaferri, and Ian Preston**, “Consumption Inequality and Partial Insurance,” *American Economic Review*, 2008, 98 (5), 1887–1921.
- Bornstein, Gideon and Sasha Indarte**, “The Impact of Social Insurance on Household Debt,” Working Paper. 2020.
- Boutros, Michael**, “Household Finances and Fiscal Stimulus in 2008,” Technical Report, Working Paper 2019.
- , “The Prevalence and Relevance of Credit Card Debt,” Working Paper. Feb. 2020.
- Browning, Martin and M. Dolores Collado**, “The Response of Expenditures to Anticipated Income Changes: Panel Data Estimates,” *American Economic Review*, Jun. 2001, 91 (3), 681–692.

- Campbell, Jeffrey R. and Zvi Hercowitz**, “Liquidity Constraints of the Middle Class,” *American Economic Journal: Economic Policy*, Aug. 2019, 11 (3), 130–155.
- Campbell, John Y.**, “Household Finance,” *Journal of Public Economics*, Aug. 2006, 61 (4), 1553–1604.
- Chetty, Raj and Adam Szeidl**, “Consumption Commitments and Risk Preferences,” *The Quarterly Journal of Economics*, May 2007, 122 (2), 831–877.
- Cochrane, John H.**, “The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives,” *The American Economic Review*, Jun. 1989, 79 (3), 319–337.
- DellaVigna, Stefano**, “Structural Behavioral Economics,” in B. Douglas Bernheim, Stefano DellaVigna, and David Laibson, eds., *B. Douglas Bernheim, Stefano DellaVigna, and David Laibson, eds., Vol. 1 of Handbook of Behavioral Economics: Applications and Foundations 1*, North-Holland, 2018, pp. 613–723.
- Epley, Nicholas, Dennis Mak, and Lorraine Chen Idson**, “Bonus or Rebate? The Impact of Income Framing on Spending and Saving,” *Journal of Behavioral Decision Making*, January 2006, 19, 1–15.
- Ergin, Haluk and Todd Sarver**, “A Unique Costly Contemplation Representation,” *Econometrica*, July 2010, 78 (4), 1285–1339.
- Fagereng, Andreas, Martin B. Holm, and Gisle J. Natvik**, “MPC Heterogeneity and Household Balance Sheets,” *American Economic Journal: Macroeconomics*, 2021, 13 (3), 1–54.
- Fogel, Suzanne**, “Income Source Effects,” Working Paper 1999.
- Forster, Jens and Laura Dannenberg**, “GLOMOsys: A systems account of global versus local processing,” *Psychological Inquiry*, 2010, 21 (3), 175–197.
- Fuchs-Schuendeln, Nicola and Tarek Alexander Hassan**, *Natural Experiments in Macroeconomics*, Vol. 2 of *Handbook of Macroeconomics*, North-Holland,
- Fuster, Andreas, Greg Kaplan, and Basit Zafar**, “What Would You Do with \$500? Spending Responses to Gains, Losses, News and Loans,” *Review of Economic Studies*, Forthcoming.
- Gabaix, Xavier**, *Behavioral Inattention*, Vol. 2 of *Handbook of Behavioral Economics*, Elsevier,

- , **David Laibson, Guillermo Moloche, and Stephen Weinberg**, “Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model,” *American Economic Review*, Sept. 2006, 96 (4), 1043–1068.
- Ganong, Peter, Damon Jones, Pascal Noel, Diana Farrell, Fiona Greig, and Chris Wheat**, “Wealth, Race, and Consumption Smoothing of Typical Income Shocks,” Technical Report, Working Paper 2020.
- Gelman, Michael**, “What Drives Heterogeneity in the Marginal Propensity to Consume? Temporary Shocks vs Persistent Characteristics,” *Journal of Monetary Economics*, 2021, 117, 521–542.
- , “The Self-Constrained Hand-to-Mouth,” *The Review of Economics and Statistics*, Forthcoming.
- Grossman, Sanford and Guy Laroque**, “Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods,” *Econometrica*, 1990, 58 (1), 25–51.
- Harris, Christopher and David Laibson**, “Dynamic Choices of Hyperbolic Consumers,” *Econometrica*, 2001, 69 (4), 935–957.
- Heathcote, Jonathan, Fabrizio Perri, and Giovanni L. Violante**, “Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967–2006,” *Review of Economic Dynamics*, Jan. 2010, 13 (1), 15–51.
- Hsieh, Chang-Tai**, “Do Consumers React to Anticipated Income Changes? Evidence from the Alaska Permanent Fund,” *American Economic Review*, Mar. 2003, 93 (1), 397–405.
- Huang, Kevin X. D. and Frank Caliendo**, “Rationalizing Multiple Consumption-Saving Puzzles in a Unified Framework,” *Frontiers of Economics in China*, 2011, 6 (3), 359–388.
- Ilut, Cosmin and Rosen Valchev**, “Economic Agents as Imperfect Problem Solvers,” Technical Report, Working Paper 2019.
- Jappelli, Tullio**, “Who is Credit Constrained in the U. S. Economy?,” *The Quarterly Journal of Economics*, Feb. 1990, 105 (1), 219–234.
- **and Luigi Pistaferri**, “Fiscal Policy and MPC Heterogeneity,” *American Economic Journal: Macroeconomics*, October 2014, 6 (4), 107–136.
- Kaplan, Greg and Giovanni L. Violante**, “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica*, July 2014, 82 (4), 1199–1239.

- , **Ben Moll**, and **Gianluca Violante**, “Monetary Policy According to HANK,” *American Economic Review*, 2018, 3 (108), 697–743.
- , **Giovanni L. Violante**, and **Justin Weidner**, “The Wealthy Hand-to-Mouth,” *Brookings Papers on Economic Activity*, 2014, 1, 77–138.
- Kimball, Miles S.**, “Precautionary Saving in the Small and in the Large,” *Econometrica*, Jan. 1990, 58 (1), 53–73.
- Krueger, Dirk**, **Kurt Mitman**, and **Fabrizio Perri**, *Macroeconomics and Heterogeneity, Including Inequality*, Vol. 2 of *Handbook of Macroeconomics*, Elsevier,
- Kueng, Lorenz**, “Excess Sensitivity of High-Income Consumers,” *The Quarterly Journal of Economics*, November 2018, 133 (4), 1693–1751.
- Laibson, David**, “Golden Eggs and Hyperbolic Discounting,” *Quarterly Journal of Economics*, 1997, 112 (2), 443–477.
- , **Peter Maxted**, and **Benjamin Moll**, “Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy,” Working Paper 2021.
- Lewis, Daniel**, **Davide Melcangi**, and **Laura Pilossoph**, “Latent Heterogeneity in the Marginal Propensity to Consume,” Working Paper 2021.
- Lian, Chen**, “Mistakes in Future Consumption, High MPCs Now,” Working Paper 2021.
- Mankiw, Gregory N.**, “The Savers-Spenders Theory of Fiscal Policy,” *American Economic Review*, May 2000, 90 (2), 120–125.
- McDowall, Robert A.**, “Consumption Behavior Across the Distribution of Liquid Assets,” Working Paper 2020.
- Misra, Kanishka** and **Paolo Surico**, “Consumption, Income Changes, and Heterogeneity: Evidence from Two Fiscal Stimulus Programs,” *American Economic Journal: Macroeconomics*, October 2014, 6 (4), 84–106.
- Moon, Hyungsik** **Roger** and **Frank Schorfheide**, “Estimation with Overidentifying Inequality Moment Conditions,” *Journal of Econometrics*, Dec. 2009, 153 (2), 136–154.
- Navon, David**, “Forest before trees: The precedence of global features in visual perception,” *Cognitive Psychology*, July 1977, 9 (3), 353–383.

- Olafsson, Arna and Michaela Pagel**, “The Liquid Hand-to-Mouth: Evidence from Personal Finance Management Software,” *The Review of Financial Studies*, November 2018, 31 (11), 4398–4446.
- Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland**, “Consumer Spending and the Economic Stimulus Payments of 2008,” *American Economic Review*, October 2013, 103 (6), 2530–2553.
- Reis, Ricardo**, “Inattentive Consumers,” *Journal of Monetary Economics*, 2006, 53 (8), 1761–1800.
- Rudebusch, Glenn D. and Eric T. Swanson**, “The Bond Premium in a DSGE Model with Long-Run Real and Nominal Risks,” *American Economic Journal: Macroeconomics*, Jan. 2012, 4 (1), 105–43.
- Shapiro, Matthew D. and Joel Slemrod**, “Did the 2008 Tax Rebates Stimulate Spending?,” *American Economic Review: Papers & Proceedings*, 2009, 99 (2), 374–379.
- Spears, Dean**, “Cognitive Limits, Apparent Impatience, and Monthly Consumption Cycles: Theory and Evidence from the South African Pension,” Technical Report, Working Paper 2012.
- Thaler, Richard H.**, “Anomalies: Saving, Fungibility, and Mental Accounts,” *Journal of Economic Perspectives*, Winter 1990, 4 (1), 193–205.

Appendices

A Proofs

Let $V_t(\Delta, k; w)$ denote the lifetime value from smoothing shock Δ over k periods for a household at time t with long-term wealth w . Formally:

$$V_t(\Delta, k; w) = \max_{\{\tilde{c}_\tau^\Delta\}_{\tau=t}^{t+k-1}} \sum_{\tau=t}^T \beta^{\tau-t} u(c_\tau + \tilde{c}_\tau^\Delta),$$

with the entire shock spent over the k planning periods, i.e.,

1. over the first k periods,

$$\sum_{\tau=t}^{t+k-1} \frac{\tilde{c}_\tau^\Delta}{(1+r)^{t-\tau}} = \Delta,$$

2. for the remaining periods, $\tau > t + k - 1$,

$$\tilde{c}_\tau^\Delta = 0.$$

In this equation, total consumption is expressed using the definition of “excess consumption” from Section 1. Total consumption is the sum of the initial lifecycle consumption plan, c_τ , and the marginal consumption out of the income shock, \tilde{c}_τ^Δ . As discussed, the initial consumption and savings plan continue to evolve according to the policy functions from the lifecycle optimization and are unaffected by the windfall-induced reoptimization. For clarity and to focus on the reoptimization plans, I suppress the notation indicating that in each period, c_τ depends on wealth, w , in period τ .

Let $c_\tau^\Delta(\Delta, k; w)$ denote the optimal consumption out of the windfall shock at time τ for a household at time t that has wealth w , faces shock Δ , and reoptimizes over k periods. Note that for every period τ , this function is defined over the initial state variable w , as these will be the focus of the proofs.

Then the expression above can be rewritten as:

$$V_t(\Delta, k; w) = \sum_{\tau=t}^{t+k-1} \beta^{\tau-t} u(c_\tau + c_\tau^\Delta(\Delta, k; w)) + \sum_{\tau=t+k}^T \beta^{\tau-t} u(c_\tau),$$

Lemma 1. *Fix wealth, w . Then:*

$$\frac{\partial V_t(\Delta, k+1; w)}{\partial \Delta} > \frac{\partial V_t(\Delta, k; w)}{\partial \Delta}.$$

Proof. With respect to Δ :

$$\frac{\partial V_t(\Delta, k; w)}{\partial \Delta} = \sum_{\tau=t}^{t+k-1} \beta^{\tau-t} u'(c_\tau + c_\tau^\Delta(\Delta, k; w)) \cdot \frac{\partial c_\tau^\Delta(\Delta, k; w)}{\partial \Delta} > 0.$$

From before, the entire shock is spent inside the planning horizon, and the derivative of this constraint is given by:

$$\sum_{\tau=t}^{t+k-1} \frac{1}{(1+r)^{t-\tau}} \frac{\partial c_\tau^\Delta(\Delta, k; w)}{\partial \Delta} = 1.$$

Without loss, assume that $\beta = (1+r)^{-1}$. The sum above is weighted in each period by the derivative of the short-run consumption function, with weights summing to unity.

When k increases, there is an additional term in the summation and the weights continue summing to unity. By construction, the short-term construction function is decreasing in the number of planning horizons. The value entering the marginal utility function decreases in every term, and since marginal utility is decreasing, each term is larger. Thus the sum is over more terms, and each term is increasing, so the total summation is larger. \square

Lemma 2. Fix Δ . Then:

$$\frac{\partial V_t(\Delta, k+1; w)}{\partial w} > \frac{\partial V_t(\Delta, k; w)}{\partial w}.$$

Proof. The derivative with respect to w :

$$\begin{aligned} \frac{\partial V_t(\Delta, k; w)}{\partial w} &= \sum_{\tau=t}^T \beta^{\tau-t} u'(c_\tau + c_\tau^\Delta(\Delta, k; w)) \\ &\quad \times \left(\frac{\partial c_\tau}{\partial w} + \frac{\partial c_\tau^\Delta(\Delta, k; w)}{\partial w} \right) \end{aligned}$$

Note that from $t+k$ to T , all of the c_τ^Δ terms are zero. When w changes, the marginal value is how utility changes with consumption, $u'(\cdot)$, multiplied by how consumption changes.

Consider the difference between the left- and right-hand side expressions in the inequality. The terms $\tau > t+k$ are equal and net to zero. In $\tau = t+k$, short-run consumption is zero for planning horizon k but positive for planning horizon $k+1$, and the terms in $\tau \in \{t, t+1, \dots, t+k+1\}$ differ since the short-term consumption function is different for the two planning horizons.

As in Lemma 1, the multiplicative term in brackets can be normalized into a weighted

average composing of the marginal utility functions. From the constraint for short-term consumption, we know that:

$$\sum_{\tau=t}^{t+k-1} \frac{1}{(1+r)^{t-\tau}} \frac{\partial c_{\tau}^{\Delta}(\Delta, k; w)}{\partial w} = 0.$$

This states that when long-term wealth changes, the total change in short-run consumption does not change, since the income shock does not change. As such, for either planning horizon, the sum of the weights does not change, but is re-arranged across the different terms.

Consider the terms of the summation above which differ:

$$\begin{aligned} & \sum_{\tau=t}^{t+k} \beta^{\tau-t} u'(c_{\tau} + c_{\tau}^{\Delta}(\Delta, k+1; w)) \times \left(\frac{\partial c_{\tau}}{\partial w} + \frac{\partial c_{\tau}^{\Delta}(\Delta, k; w)}{\partial w} \right) \\ & > \sum_{\tau=t}^{t+k-1} \beta^{\tau-t} u'(c_{\tau} + c_{\tau}^{\Delta}(\Delta, k; w)) \times \left(\frac{\partial c_{\tau}}{\partial w} + \frac{\partial c_{\tau}^{\Delta}(\Delta, k; w)}{\partial w} \right) + \beta^{t+k} u'(c_{t+k} + 0) \times \left(\frac{\partial c_{t+k}}{\partial w} + 0 \right) \end{aligned}$$

The terms $\tau > t+k$ do not appear because they are equal. In $\tau \in \{t, t+1, \dots, t+k-1\}$, the terms differ since the short-term consumption function is different for the two planning horizons. In the first line with the longer planning horizon, the summation is to $t+k$, whereas in the second line representing the shorter planning horizon, the summation is to $t+k-1$. I include the term for $t+k$ in the second line for the sake of comparison. In $\tau = t+k$, short-term consumption (and its derivative) are zero for the shorter planning horizon.

Recognizing that these two expressions can be expressed this way as an equal number of terms, then it immediately follows from convexity of the marginal utility function that the first expression is strictly greater than the second.¹⁸

□

A.1 Proof of Proposition 1

Consider three planning horizons of decreasing length, $k_2 > k_1 > k_0$. Let $k_1 \equiv k^*(\Delta)$ denote the optimal planning horizon for the smaller income shock, Δ . I establish the weak inequality in two steps.

¹⁸Convexity of the marginal utility function follows from the presence of incomplete markets and occasionally binding borrowing constraints in the stochastic case and is directly assumed (i.e., prudence) in the deterministic case.

First, I prove that for Δ' , the planning horizon k_1 dominates any $k_0 < k_1$. Given the optimality of k_1 for Δ , we have that

$$V_t(\Delta, k_1; w) - \phi_{k_1} > V_t(\Delta, k_0; w) - \phi_{k_0},$$

and, re-arranging, that

$$V_t(\Delta, k_1; w) - V_t(\Delta, k_0; w) > \phi_{k_1} - \phi_{k_0}.$$

This expression states that the marginal value from increasing the planning horizon is more than offset by the marginal increase in planning costs.

Since $k_2 > k_1$, by Lemma 1,

$$V_t(\Delta', k_1; w) - V_t(\Delta', k_0; w) > V_t(\Delta, k_1; w) - V_t(\Delta, k_0; w).$$

The marginal value from increasing planning horizons is larger for Δ' than it is for Δ . Combining with the above and re-arranging,

$$\begin{aligned} V_t(\Delta', k_1; w) - V_t(\Delta', k_0; w) &> \phi_{k_1} - \phi_{k_0} \\ V_t(\Delta', k_1; w) - \phi_{k_1} &> V_t(\Delta', k_0; w) - \phi_{k_0}, \end{aligned}$$

establishing that for Δ' , k_1 is preferred over k_0 . Intuitively, if increasing planning horizons from k_0 to k_1 is preferred for the smaller shock, then this is also preferred for the larger shock given that the slope of the value function with respect to planning horizons is increasing in the income shock.

Second, I prove that for Δ' , the planning horizon $k_2 > k_1$ may be optimal. This is the case when

$$V_t(\Delta', k_2; w) - \phi_{k_2} > V_t(\Delta', k_1; w) - \phi_{k_1},$$

which holds if the marginal value from increasing the planning horizon is larger than the marginal cost,

$$V_t(\Delta', k_2; w) - V_t(\Delta', k_1; w) > \phi_{k_2} - \phi_{k_1}.$$

This expression may obtain given the structure of the value function or planning costs.

A.2 Proof of Proposition 2

This proof proceeds similarly to the proof for Proposition 1. Given income y , consider two levels of wealth at time t such that $w' > w$ and three planning horizons of decreasing

length, $k_2 > k_1 > k_0$. Let $k_1 \equiv k_t^*(w)$ denote the optimal planning horizon for the smaller level of initial wealth, w . I establish the weak inequality in two steps.

First, I prove that for w' , the planning horizon k_1 dominates any $k_0 < k_1$. Given the optimality of k_1 for w , we have that

$$V_t(\Delta, k_1; w) - \phi_{k_1} > V_t(\Delta, k_0; w) - \phi_{k_0},$$

and, re-arranging, that

$$V_t(\Delta, k_1; w) - V_t(\Delta, k_0; w) > \phi_{k_1} - \phi_{k_0}.$$

As above, this expression states the the marginal value of increasing the planning horizon is more than offset by the marginal increase in planning costs. Since $k_2 > k_1$, by Lemma 2,

$$V_t(\Delta, k_1; w', y) - V_t(\Delta, k_0; w', y) > V_t(\Delta, k_1; w, y) - V_t(\Delta, k_0; w, y).$$

The marginal value from increasing planning horizons is larger for w' than it is for w . Combining with the above and re-arranging,

$$\begin{aligned} V_t(\Delta, k_1; w', y) - V_t(\Delta, k_0; w', y) &> \phi_{k_1} - \phi_{k_0} \\ V_t(\Delta, k_1; w', y) - \phi_{k_1} &> V_t(\Delta, k_0; w', y) - \phi_{k_0} \end{aligned}$$

establishing that for w' , k_1 is preferred over k_0 . Intuitively, if increasing planning horizons from k_0 to k_1 is preferred for the lower level of wealth, then this is also preferred for the larger level of wealth given that the slope of the value function with respect to planning horizons is increasing in wealth.

Second, I prove that for w' , the planning horizon $k_2 > k_1$ may be optimal. This is the case when

$$V_t(\Delta, k_2; w) - \phi_{k_2} > V_t(\Delta, k_1; w) - \phi_{k_1},$$

which holds if the marginal value from increasing the planning horizon is larger than the marginal cost,

$$V_t(\Delta, k_2; w) - V_t(\Delta, k_1; w) > \phi_{k_2} - \phi_{k_1}.$$

This expression may obtain given the structure of the value function or planning costs.

B Additional Regressions Using Survey of Consumer Expenditures

Table 6: Spending Response of Consumption to ESP by Income Group

Panel (a): Terciles by Income						
	(a) Non-Durables		(b) Durables		(c) Total	
	Estimate	Implied MPC	Estimate	Implied MPC	Estimate	Implied MPC
ESP (Base: Tercile 1)		0.251 (0.161)		0.914* (0.486)		1.165** (0.520)
ESP \times Tercile 2	-0.104 (0.119)	0.147 (0.122)	-0.518 (0.366)	0.396 (0.409)	-0.622 (0.390)	0.543 (0.439)
ESP \times Tercile 3	-0.006* (0.130)	0.245** (0.116)	-0.373 (0.379)	0.541 (0.379)	-0.378 (0.406)	0.787* (0.406)
Observations		8,592		8,592		8,592
R^2		0.017		0.005		0.008

Panel (b): Terciles by $\frac{1}{\text{Income}}$						
	(a) Non-Durables		(b) Durables		(c) Total	
	Estimate	Implied MPC	Estimate	Implied MPC	Estimate	Implied MPC
ESP (Base: Tercile 1)		0.248** (0.116)		0.566 (0.375)		0.813** (0.403)
ESP \times Tercile 2	-0.107 (0.108)	0.141 (0.122)	-0.207 (0.350)	0.359 (0.414)	-0.314 (0.377)	0.500 (0.444)
ESP \times Tercile 3	0.003 (0.129)	0.251 (0.161)	0.347 (0.376)	0.913* (0.486)	0.350 (0.403)	1.163** (0.520)
Observations		8,592		8,592		8,592
R^2		0.017		0.005		0.008

Notes: Standard errors in parentheses. *, **, *** denote significance at the 0.10, 0.05, and 0.01 levels under the assumption of a single test. Estimated using two stages least squares and instrumenting for ESP amount with an indicator for ESP receipt. See Parker et al. (2013) for more details.

C ESPs to Wealthy Hand-to-Mouth Households

Kaplan and Violante (2014) introduce a new class of constrained households that they call the wealthy hand-to-mouth. They define households as hand-to-mouth using liquid wealth, and define them as poor or wealth using illiquid wealth. Traditionally, empirical

analysis in this literature focuses only on total net worth and thus only the poor hand-to-mouth. Instead, Kaplan and Violante show that a large fraction of households that have low liquid wealth but high net worth behave similarly to households with low liquid wealth and low net worth. These households are defined as the wealthy hand-to-mouth.

A key insight to their analysis is that the ratio of liquid wealth to income is the relevant statistic, as opposed to the level of liquid wealth:

$$LWI = \frac{\text{liquid wealth}}{\text{periodic income}}.$$

For example, a household that earns \$1000 per month and carries \$5000 in liquid wealth has $LWI = 5$, whereas a household that earns \$10,000 per month and carries \$5000 in liquid wealth has $LWI = 0.5$.

Kaplan and Violante deem households hand-to-mouth if their LWI ratios fall within one of two intervals. First, if their liquid wealth to income ratio is between 0 and 1, the household is hand-to-mouth because they keep less than one month of income on hand. Second, allowing for a credit limit up to one month of income, households whose liquid wealth to income is less than -1 are also hand-to-mouth.

Table 7: CEX Regressions with H2M Indicator

	Estimate	Implied MPC
ESP (Base: Tercile 1)		0.186 (0.163)
ESP \times Hand-to-Mouth	0.045 (0.152)	0.231 (0.143)
Observations		3,446
R^2		0.024

In the CEX, I calculate LWI and remove extreme outliers ($LWI > 10$). I then calculate hand-to-mouth status using the two criteria above. To estimate the differential MPC for hand-to-mouth consumers, I estimate the baseline regression and interact the hand-to-mouth indicator with ESP payment. The estimated coefficients are in Table 7. Although none of the estimated coefficients are statistically significant, the patterns are consistent with existing evidence. The MPC for hand-to-mouth households is 0.231, which is 0.045 percentage points or almost 25% larger than the MPC for non-hand-to-mouth households, 0.186. Again, however, the MPC for non-hand-to-mouth households is much larger than the prediction of near-zero MPCs in standard models.

Table 8: CEX Regressions with LWI Terciles

	Estimate	Implied MPC
ESP (Base: Tercile 1)		0.334* (0.173)
ESP \times Tercile 2	-0.263 (0.178)	0.071 (0.154)
ESP \times Tercile 3	-0.082 (0.191)	0.252 (0.180)
Observations		3,446
R^2		0.024

As a robustness check, I also separate households into terciles based on liquid wealth to income. In Kaplan et al. (2014), the authors check robustness by changing the criteria used to define LWI (i.e. pay periods, credit limits, etc.), which essentially changes the intervals that define hand-to-mouth status. Dividing households by LWI serves the same purpose. I estimate the baseline regression and interact LWI tercile with ESP payment and the estimated coefficients are in Table 8. Again consistent with existing evidence, I find that households in the lowest tercile are those with the highest MPC. However, the estimated relationship between LWI tercile and MPC is U-shaped, similar to previous findings using this data for the relationship between MPC and wealth. Households in the middle tercile have an MPC of only 0.071, while households in the high tercile have an MPC of 0.252. Overall, it is hard to infer too much from this U-shaped pattern, but it remains the case that the MPC for unconstrained households is too high relative to what standard models would predict.

D Theoretical Construction of the Marginal Propensity to Consume

In standard consumption-savings problems, the marginal propensity to consume out of a temporary income shock is a partial derivative of the consumption function. If the temporary income shock is represented as a distinct state, then the derivative is taken with respect to that state. For example, consider a standard one-asset model in which a household forms state-contingent plans over wealth, a , autoregressive permanent income, ν , and perfectly transient temporary income shocks, ϵ . The marginal propensity to consume

out of temporary income shocks is given by

$$\text{MPC}(a, \nu, \epsilon) = \frac{\partial c(a, \nu, \epsilon)}{\partial \epsilon}.$$

In practice, the statespace in standard models can be reduced by one dimension since the temporary income shock is equivalent to wealth. To see this, note that when the household's problem is written recursively, the consumption policy function is given by

$$c(a, \nu, \epsilon) = \arg \max_c u(c) + \beta E[V(a', \nu', \epsilon') | \nu],$$

with $a' = (1+r)(\nu + \epsilon + a - c(a, \nu, \epsilon))$. Via the budget constraint, a change in ϵ is equivalent to a change in a . Economically, the perfectly transient income shock is equivalent to the household beginning the period with a different level of wealth. Importantly, a change in the temporary income shock, ϵ , is not equivalent to a change in permanent income, ν , because the latter is autoregressive and enters the conditional expectation.

When the statespace is reduced to wealth and permanent income, the marginal propensity to consume out of a temporary income shock, ϵ , can be written as

$$\lim_{\epsilon \rightarrow 0} \frac{c(a + \epsilon, \nu) - c(a, \nu)}{\epsilon},$$

which is the partial derivative of the consumption function with respect to wealth.

E Discounting of Future Periods

The finite planning horizon model developed in this paper is isomorphic to the standard recursive consumption-saving model with a specific discount rate structure. In this section, I consider a household that lives for T periods and faces an income shock at time $t = 1$, and I compare the discount rate structure of exponential or quasi-hyperbolic discounting with the finite planning horizon model.

Suppose that, for the same reasons as in the finite planning horizon model, the household in the standard model must reoptimize consumption and savings plans to account for the unexpected change in income at time $t = 1$. The first two rows of Table 9 show how future periods are discounted with exponential and quasi-hyperbolic discounting.

In the standard model, discount rates are a geometric series with base β . With quasi-hyperbolic discounting, the household discounts between time $t + 1$ and $t + 2$ using $\delta\beta$, but

Table 9: Discounting Factors in Consumption-Saving Models

	Periods After $t = 1$								
	1	2	...	$k - 1$	k	$k + 1$...	$T - 1$	T
Standard Discounting	β	β^2	...	β^{k-1}	β^k	β^{k+1}	...	β^{T-1}	β^T
Quasi-Hyperbolic Discounting	β	$\delta\beta^2$...	$\delta\beta^{k-1}$	$\delta\beta^k$	$\delta\beta^{k+1}$...	$\delta\beta^{T-1}$	$\delta\beta^T$
Finite k -period Plan	β	β^2	...	β^{k-1}	β^k	0	...	0	0

then discounts any two further future periods, i.e., $t + 3$ and $t + 4$,¹⁹ using only β . This generates present bias since the household discounts the immediate future more than the distant future.

The third row of Table 9 shows discount rates in the finite horizon model. Faced with an income shock at time $t = 1$, the household determines an optimal k -period planning horizon over which to reoptimize. Within the planning horizon, the household uses standard exponential discounting. Beyond the planning horizon, the household behaves as-if the income shock had never occurred and uses its existing long-term consumption and savings plans. As a result, from the perspective of its reoptimization, it is as-if the household completely disregards all periods beyond k , i.e., discounts them with rate zero.

F Constructing Liquid Wealth in the CEX


In the 2008 wave of the Survey of Consumer Expenditures (CEX), approximately half of households did not respond to the optional question on liquid assets, and there is no corresponding question on liquid debt. As such, in this section, I describe the method by which I merge the CEX and Survey of Consumer Finances (SCF), which contains detailed data on both liquid assets and debt.

The SCF collects responses from two groups: a random sample of US households and an additional selected sample of high-wealth households (who, generally, have higher income). In order to make the distributions of household income in the SCF and CEX comparable, I drop the top 6% of observations (see, for example, Heathcote, Perri and Violante (2010) for a broader discussion on comparing household income and wealth across

¹⁹More generally, $t + s$ and $t + s + 1$ for any $s \geq 2$. I also deviate from the standard notation in quasi-hyperbolic discounting in order to maintain comparability with the notation in standard discounting. Specifically, I use β for the exponential/geometric discounting and δ for the additional first-period discounting, as opposed to the opposite notation usually employed in this literature (see, for example, Laibson (1997)).

surveys). In Figure 9, I plot monthly income by percentile in each survey. I evaluated dropping between the top 1 and 10% of households in the SCF and found that 6% minimized the mean squared error between each line in the figure. Notably, household income in each percentile is remarkably similar until around the 95th percentile.

Figure 9: Monthly Income by Percentile in CEX and SCF

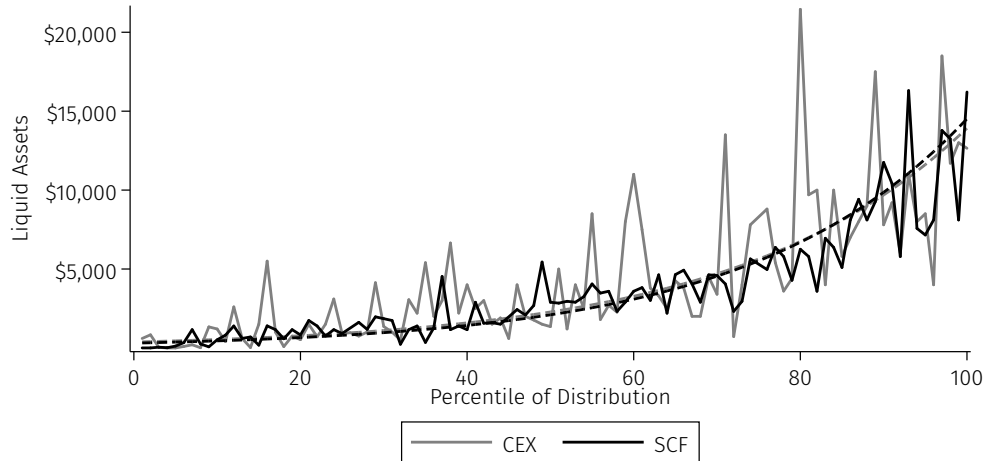


Notes: Top 6% of income distribution in SCF truncated in order to approximate income distribution in CEX. See text for details.

Figure 10 verifies that liquid assets in each distribution are similar. The solid lines represent the median value of liquid assets in each survey. Unsurprisingly, liquid assets in the CEX display much more variability. In both surveys, liquid assets are increasing in income; this is made clearer by the dashed lines, which are estimated logarithmic regressions of the income percentile on median liquid assets. The lines are almost perfectly overlaid, demonstrating the similarity between liquid assets in the raw CEX and truncated SCF distributions.


Finally, Figure 11 plots median liquid assets and liquid wealth by income percentile in the SCF. The difference between the two is liquid debt, which has been documented to be increasing in household income (see, for example, Boutros (2020) and Bornstein and Indarte (2020)). For all households, there is a notable difference between liquid assets and wealth. Altogether, these values of liquid wealth from the SCF are matched to households in the CEX by percentile of monthly income.

Figure 10: Liquid Assets by Percentile in CEX and SCF



Notes: Top 6% of income distribution in SCF truncated in order to approximate income distribution in CEX. See text for details.

Figure 11: Liquid Assets and Liquid Wealth by Percentile in SCF



Notes: Top 6% of income distribution in SCF truncated in order to approximate income distribution in CEX. See text for details.