

A Model of the Consumption Response to Fiscal Stimulus Payments*

Greg Kaplan

Princeton University, IFS and NBER
gkaplan@princeton.edu

Giovanni L. Violante

New York University, CEPR, and NBER
glv2@nyu.edu

16 FEBRUARY, 2014

Abstract

A wide body of empirical evidence finds that approximately 25 percent of fiscal stimulus payments (e.g., tax rebates) are spent on nondurable household consumption in the quarter that they are received. To interpret this fact, we develop a structural economic model where households can hold two assets: a low-return liquid asset (e.g., cash, checking account) and a high-return illiquid asset that carries a transaction cost (e.g., housing, retirement account). The optimal life-cycle pattern of portfolio choice implies that many households in the model are “wealthy hand-to-mouth”: they hold little or no liquid wealth despite owning sizeable quantities of illiquid assets. Therefore, they display large propensities to consume out of additional transitory income, and small propensities to consume out of news about future income. We document the existence of such households in data from the Survey of Consumer Finances. A version of the model parameterized to the 2001 tax rebate episode yields consumption responses to fiscal stimulus payments that are in line with the evidence, and an order of magnitude larger than in the standard “one-asset” framework. The model’s nonlinearities with respect to the rebate size and the prevailing aggregate economic conditions have implications for policy design.

Keywords: Consumption, fiscal stimulus payments, hand-to-mouth, liquidity.

JEL Classification: D31, D91, E21, H31

*We thank Kurt Mitman for outstanding research assistance. We are grateful to Jonathan Heathcote, Ricardo Lagos, Sydney Ludvigson, and Sam Schulhofer-Wohl for their useful insights, to numerous seminar participants for comments, and to Jonathan Parker and Lubos Pastor for sharing their data. This research is supported by grant no. 1127632 from the National Science Foundation.

1 Introduction

Fiscal stimulus payments, such as transfers to households in the form of tax rebates, are frequently used by governments to alleviate the impact of recessions on households' welfare. This type of fiscal intervention was authorized by the U.S. Congress in the last two downturns of 2001 and 2007-2009.¹ Households received one-off payments that ranged from \$500 to \$1,000, depending on the specific episode. In the aggregate, these fiscal outlays amounted to \$38 billion in 2001 and \$96 billion in 2008, roughly equivalent to 0.4-0.7% of annual GDP.

On the empirical side, substantial progress has been made in measuring the size of household consumption responses to the tax rebate episodes of 2001 and 2008. In both instances the U.S. Treasury scheduled payments based on the last two digits of individual Social Security Numbers, which are effectively random. Johnson, Parker, and Souleles (2006, hereafter JPS) and Parker, Souleles, Johnson, and McLelland (2011, hereafter PSJM) cleverly exploited this randomized timing of the receipt of payments to estimate the effects of the fiscal stimulus on consumption expenditures. Subsequently, Misra and Surico (2013) refined the econometric analysis in these studies. Shapiro and Slemrod (2003a, 2003b, 2009) reinforced this evidence with informative qualitative surveys on how consumers use their rebate.

This collective evidence convincingly concludes that households spend approximately 25 percent of rebates on nondurables in the quarter that they are received. This strong consumption response is measured relative to the control group of households (comparable, because of the randomization) that do not receive the rebate in that same quarter. In the paper we call this magnitude the *rebate coefficient*.²

In spite of this large body of empirical research, there are no quantitative studies of these episodes within dynamic structural models of household behavior. This gap in the literature is troubling because a thorough understanding of the effectiveness of tax rebates as a short-term stimulus for aggregate consumption is paramount for macroeconomists and policy makers.³ Identifying the determinants of how consumers

¹In the context of the latest downturn, Oh and Reis (2011) document that the large fiscal expansion of 2007-2009 consisted primarily of growing social assistance, as opposed to government purchases. Half of this expansion comprised discretionary fiscal stimulus transfers.

²In a regression where the dependent variable is household consumption growth in a given quarter and the right hand side variable is the size of the rebate received in that quarter, possibly zero, the rebate coefficient measures the differential consumption growth of the treatment group -the rebate recipients- relative to the control group of non-recipients.

³Estimates by JPS (2006) feature prominently in the reports prepared by the Congressional Budget

respond to stimulus payments helps in choosing policy options and in assessing whether the same fiscal instrument can be expected to be more or less effective under different macroeconomic conditions.⁴

To develop a structural model that has some hope of matching this micro evidence, one cannot rely on off-the-shelf consumption theory: the rational expectations, life-cycle, buffer-stock model with one risk-free asset (Deaton, 1991; Carroll, 1992, 1997; Rios-Rull, 1995; Huggett, 1996; for a survey, see Heathcote et al., 2009) predicts that the marginal propensity to consume (MPC) out of transitory income fluctuations, such as tax rebates, should be negligible in the aggregate. In this standard one-asset model, the only agents whose consumption would react significantly to receiving a rebate check are those who are constrained. However, under parameterizations where the model’s distribution of net worth is in line with the data, the fraction of constrained households is too small (usually around 10%) to generate a big enough response in the aggregate.⁵

We overcome this challenge by proposing a quantitative framework that speaks to the data on both liquid and illiquid wealth, rather than on net worth alone. To do this, we integrate the classical Baumol-Tobin model of money demand into a partial-equilibrium version of the workhorse incomplete-markets life-cycle economy. In our model, households can store wealth in two types of instruments: a liquid asset, such as cash or bank accounts, and an illiquid asset, such as housing or retirement wealth. Households can also borrow through unsecured credit. The trade-off between the liquid and illiquid asset is that the latter earns an exogenously higher rate of return, but can be accessed only by paying a transaction cost. The model is parameterized to replicate a number of macroeconomic, life-cycle, and cross-sectional targets.

Besides the usual small fraction of poor hand-to-mouth agents with zero net worth, our model features a significant number of what we call wealthy hand-to-mouth households. These are households that hold sizeable amounts of illiquid wealth, yet optimally choose to consume all of their disposable income during a pay-period. Examining as-

Office (CBO, 2009) and the Council of Economic Advisors (CEA, 2010) documenting and forecasting the macroeconomic effects of fiscal stimulus.

⁴JPS (2006, p. 1607) end their empirical analysis of the 2001 tax rebates with: “*without knowing the full structural model underlying these results, we cannot conclude that future tax rebates will necessarily have the same effect.*”. Shapiro and Slemrod (2003a, p. 394) end theirs with “*key parameters such as the propensity to consume are contingent on aggregate conditions in ways that are difficult to anticipate.*”

⁵Even the one-asset model can, under parameterizations where many agents hold close to zero net worth and are very often constrained, predict nontrivial consumption responses. This explains, for example, the sizable MPC out of lump-sum tax cuts reported in some of Heathcote’s (2005) experiments aimed at quantifying deviations from Ricardian neutrality in this class of economies.

set portfolio and income data from the 2001 Survey of Consumer Finances through the lens of our two-asset model reveals that roughly 1/3 of U.S. households fit this profile. Although in our model these households act as if they are constrained, they would not appear constrained from the viewpoint of the one-asset model since they own substantial net worth.

Why would households with sizeable net worth optimally choose to consume all of their randomly fluctuating earnings every period, instead of maintaining a smooth consumption profile? The answer is that such households are better off bearing the welfare loss rather than smoothing shocks because the latter option entails either frequently paying the transaction cost to tap into their illiquid asset, or holding large balances of cash and foregoing the high return on the illiquid asset, or obtaining credit at expensive interest rates. This explanation is reminiscent of calculations by Cochrane (1989) and, more recently, Browning and Crossley (2001) showing that in some contexts the utility loss from setting consumption equal to income, instead of fully optimizing, is second order.⁶

These wealthy hand-to-mouth households are the reason why our model can generate strong average consumption responses to fiscal stimulus payments: such households do not respond to the news of the rebate and have a high MPC when they receive their payment. When we replicate, by simulation, the randomized experiment associated with the tax rebate of 2001 within our structural model, we find rebate coefficients between 11% and 25%, depending on the assumed information structure. Values at the low end of this range are obtained under the assumption that every household is fully aware of the policy one quarter ahead. In this scenario, all the non hand-to-mouth households have already responded to the news when the rebate reaches their pockets, which reduces the effect of the policy at the time of receipt of the checks. Values at the high end correspond to the case where all households are surprised by the payment when they receive it. We set our baseline between these two extremes, where half of households expect the check from the government and half are surprised by it and obtain values near to 15%, i.e., almost two thirds of our preferred estimates of rebate coefficients in the micro data.⁷

⁶The model by Campbell and Hercowitz (2009) also generates wealthy constrained agents endogenously, but through a different mechanism from ours: periodically, households discover they will have a special consumption need T periods ahead (e.g., the education of their kids). This induces them to consume low amounts until they have saved enough for the special consumption need.

⁷In line with this intermediate scenario, for the 2008 episode, Broda and Parker (2012) document that roughly 60% of households learned about the policy in the quarter before Treasury began disbursing payments.

The presence of wealthy hand-to-mouth households is also the crucial source of amplification relative to a plausibly calibrated one-asset model economy where rebate coefficients from model-simulated data are less than 1%. This pronounced magnification works through both the extensive and the intensive margin. First, in our two-asset model there are many more hand-to-mouth consumers, consistent with the SCF data. Second, the wealthy hand-to-mouth display larger MPCs out of tax rebates than their poor counterparts since they have higher wealth (tied up in the illiquid asset) and, therefore, higher desired target consumption.

Several key implications of the model are in agreement with the data. Misra and Surico (2012) estimate the entire empirical distribution of consumption responses for 2001 and document substantial heterogeneity: half of the population displays no response at all and one-fifth display responses over 50%. They also uncover high income households at both ends of the distribution. Our model replicates these two findings for two reasons. First, most of the model agents behave as PIH consumers and have MPCs close to zero, but the wealthy hand-to-mouth have MPCs close to 50%. Second, there are many high-income households among the wealthy hand-to-mouth. Moreover, the model implies a tight negative correlation between the size of the consumption response and the ratio of holdings of liquid wealth to income, as documented, for example, in Souleles (1999) or Broda and Parker (2012). Finally, the model features a marked size-asymmetry in the consumption responses to small and large payments (Hsieh, 2003): large rebates trigger many households to pay the transaction cost and deposit the extra income into the illiquid account, but when they adjust, these households are unconstrained and therefore save the bulk of their rebate.

In a series of experiments, we show how to use the structural model to improve the design of this class of policy instruments. The experiments contain two useful lessons for policy design. First, the aggregate macroeconomic conditions surrounding the policy affect the rebate fraction consumed by households in nontrivial ways. In a mild recession, where income drops are small and short-lived, it is not worthwhile for the wealthy hand-to-mouth households to pay the transaction cost to access some of their illiquid assets (or to use expensive credit) in order to smooth consumption. As a result, liquidity constraints get amplified, and the aggregate consumption response to a fiscal stimulus payment is strong. Conversely, at the outset of a severe recession that induces a large and long-lasting fall in income, many wealthy hand-to-mouth households will choose to borrow or tap into their illiquid account to create a buffer of liquid assets that can be used to smooth consumption. As a result, fewer households are hand-to-mouth

when the rebate is received. Thus, the effect of the stimulus on consumption is lower compared to when the same policy is implemented in a mild downturn.

Second, we compare budget-equivalent policies with various degrees of phasing-out and show that, to achieve the strongest bang for the buck, the rebate should be phased out around median income. A more targeted rebate has smaller effects because its size becomes large enough for the size-asymmetry to kick in, and because it misses many middle class wealthy hand-to-mouth households with high MPCs.

The structural model is also useful to understand when the micro estimates of the rebate coefficients are quantitatively close to what they aim to measure, i.e., the average MPC out of the fiscal stimulus receipt. Recall that identification of the micro estimates comes from the randomized timing of the payments across households. As a result, the consumption response of the treatment group -the group that receives the check in a given week- is measured relative to a control group that is composed of (i) households who are aware of the policy, but will receive the check in a later week, and (ii) households who have already received the payment in a previous week. Thus, the control group's response, which ideally should be unaffected by the policy, is generally a mix of the MPC out of the news about the payment, and the lagged MPC out of the payment. In this paper, we explain that (i) the lag between the date when the policy enters agents' information sets and the date when the transfer enters agents' budget constraints and (ii) the exact specification of the regression, jointly determine whether the empirical estimate is biased. Independently of the regression results, our structural model implies that the average quarterly MPC out of a surprise (anticipated) fiscal stimulus receipt is 20% (6%, respectively), and the MPC out of the news of the payment is 7%.

Our model is related to four strands of literature. A pair of influential papers by Campbell and Mankiw (1989, 1991) show that some key aspects of the comovement of aggregate consumption, income, and interest rates are best viewed as generated not by a single forward-looking type of consumer, but rather by the coexistence of two types: one forward-looking and consuming its permanent income (the saver); the other, highly impatient and following the rule of thumb of spending its current income (the spender).⁸ Our model can be seen as a microfoundation for this spender-saver view because, alongside standard buffer-stock consumers, it endogenously generates hand-to-mouth households. However, most households in this class are patient and own substantial illiquid assets, which critically changes some of the macroeconomic

⁸Recent examples of this model are Gali et al. (2007) and Justiniano et al. (2013).

implications of the model. We return to this point in the Conclusions.

The closest forebears to our framework are Angeletos et al. (2001) and Laibson et al. (2003). These two studies quantitatively compare the lifecycle portfolio allocation properties of two types of consumers: one with quasi-hyperbolic discounting and one with geometric discounting. Relative to the model with standard preferences, with quasi-hyperbolic consumers it is easier to generate both sizable borrowing through unsecured credit (since credit provides funding for instant gratification) and saving predominantly in illiquid assets (since illiquidity protects quasi-hyperbolic agents from future consumption splurges). As a result, the MPC out of predictable income changes can be large.⁹ Our exploration of the two-asset model sheds some new light on its mechanisms and quantitative reach. We demonstrate that, even when this environment is populated by geometric consumers, it can yield large MPCs out of small transitory income changes as long as it features enough wealthy hand-to-mouth households. Hyperbolic discounting magnifies the key economic forces behind the strong (weak) demand for illiquid (liquid) assets, but it is not strictly necessary to obtain a significant amplification relative to the one-asset environment. We explain how to use cross-sectional data on household portfolios to measure such households and, therefore, discipline the model’s parameterization. We apply the framework to quantitatively analyze a relevant policy question that has so far not been addressed through structural modeling.

Although in our model households ride out small shocks, they withdraw from the illiquid account to smooth out large falls in income. This rich adjustment pattern resembles that described by Chetty and Szeidl (2007) in a theoretical model with ex-ante consumption commitments, where the burden of moderate income shocks is only absorbed by fluctuations in the “flexible” consumption good, whereas large shocks also induce ex-post changes in the “commitment” good. Our model, where the illiquid asset (e.g., its housing component) generates a consumption flow, features a similar source of excess sensitivity in nondurable consumption.

Finally, a number of papers embed the Baumol-Tobin insight -the presence of a frictional transaction technology- into portfolio choice models. Prominent recent examples are Alvarez, Atkeson and Kehoe (2002), Alvarez and Lippi (2009), Abel, Eberly and Panageas (2009), and Alvarez, Guiso and Lippi (2012). Although our model is less an-

⁹Another framework that has the ability to generate a large MPC from windfall income is the “rational inattention” model (Reis, 2006). However, without the addition of some form of transaction cost -or a mechanism to generate enough wealthy hand-to-mouth consumers- this framework cannot display small consumption responses to news about future payments, which is a necessary condition to match the size of estimated rebate coefficients.

alytically tractable than most of this literature, it contains a number of additional features, crucial for generating wealthy hand-to-mouth households and empirically plausible rebate coefficients: endogenous non-durable consumption choices, borrowing constraints, uninsurable risk in non-financial income, and a lifecycle saving motive. Some examples of richer frameworks for quantitative analysis exist, but applications are essentially limited to financial issues and monetary policy.¹⁰ Our exercise shows this is also a natural environment to quantitatively analyze fiscal policy.

The rest of the paper proceeds as follows. In Section 2, we describe the 2001 tax rebate episode and present the associated empirical evidence on the estimated consumption responses. In Section 3, we outline our model and in Section 4 we document the presence of wealthy hand-to-mouth consumers in the model and in the data. Section 5 describes our parameterization. Section 6 contains the quantitative analysis of the 2001 tax rebate in the structural model. In Sections 7 and 8, we use the model to perform a number of experiments that are useful to inform the design of policy. Section 9 concludes.

2 Summary and interpretation of the empirical evidence on the 2001 tax rebate

The tax rebate of 2001 was part of a broader tax reform, the Economic Growth and Tax Relief Reconciliation Act (EGTRRA), enacted in May 2001 by the U.S. Congress. The reform included a reduction in the federal personal income tax rate for the lowest bracket (the first \$12,000 of earnings for a married couple filing jointly and the first \$6,000 for singles) from 15% to 10%, effective retroactively to January 2001. In order to make this component of the reform highly visible during calendar year 2001, the Administration paid an advance refund to taxpayers, informally called a tax rebate, for money they would have received from the Treasury only upon filing their tax returns in April 2002. The vast majority of the rebate checks were mailed between the end of July and the end of September 2001, in a sequence based on the last two digits of

¹⁰For example, within incomplete-markets economies, Aiyagari and Gertler (1991) focus on the equity premium; Erosa and Ventura (2002) revisit, quantitatively, the question of welfare effects of inflation; Ragot (2011) studies the joint distribution of money and financial assets. Two recent papers examine whether the existence of two assets featuring different return and liquidity characteristics induces “excess sensitivity” in consumption. In Li (2009), a large MPC out of anticipated income changes is obtained only for calibrations where households hold as little as one-twentieth as much wealth as in the data. Huntley and Michelangeli’s (2011) model focuses exclusively on the distinction between taxable and tax-deferred assets. As a result, the amplification in the MPC is very modest (2-4 pct points) relative to the benchmark one-asset model.

the social security number (SSN). This sequence featured in the news in June. At the same time, the Treasury mailed every taxpayer a letter informing them in which week they would receive their check. The Treasury calculated that checks were sent out to 92 million taxpayers, with almost 80 percent of them paying the maximum amount, (\$600, or 5% of \$12,000, for married couples), corresponding to a total outlay of \$38B, or almost 0.4% of 2001 GDP.

From the point of view of economic theory, the tax rebate of 2001 has three salient characteristics: (i) it is essentially a lump sum, since almost every household received \$300 per adult; (ii) it is anticipated, at least for that part of the population which received the check later and that, presumably, had enough time to learn about the rebate either from the news, from the Treasury letter, or from friends/family who had already collected theirs; and (iii) the timing of receipt of the rebate has the feature of a randomized experiment because the last two digits of a SSN are uncorrelated with any individual characteristics.

Empirical evidence JPS (2006) add a special module of questions to the Consumer Expenditure Survey (CEX) that asks households about the timing and amount of their rebate check. Among the various specifications estimated by JPS (2006) to assess the impact of the rebate on consumption expenditures, we will focus on their baseline:

$$\Delta c_{it} = \sum_s \beta_{0s} \text{month}_s + \beta_1' \mathbf{X}_{i,t-1} + \beta_2 R_{it} + \varepsilon_{it} \quad (1)$$

where Δc_{it} is the change in nondurable expenditures of household i in quarter t , month_s is a time dummy, $\mathbf{X}_{i,t-1}$ is a vector of demographics, and R_{it} is the dollar value of the rebate received by household i in quarter t . The coefficient β_2 , which we label the *rebate coefficient* is the object of interest. Identification of β_2 comes from randomization in the *timing* of the receipt of rebate checks across households. Since the size of the rebate is potentially endogenous, JPS (2006) estimate equation (1) by 2SLS using, as an instrument, an indicator for whether the rebate was received. Their key finding, reproduced in Table 2, is that β_2 is estimated to be 0.37 for nondurable consumption. Since the original estimates of JPS (2006), others have refined this empirical analysis. Hamilton (2008) argues that, since the CEX is notoriously noisy, one should trim the sample to exclude outliers; this procedure leads to smaller rebate coefficients. In Table 2, we report the 2SLS estimate that is obtained by dropping the top and bottom 0.5% and 1.5% of the distribution of nondurable consumption growth from CEX. The rebate coefficient drops to a range of 22 to 24 percent, in line with Hamilton's results. Misra and Surico (2011) use quantile regression techniques to explicitly deal with heterogene-

	Nondurables
JPS 2006, 2SLS ($N = 13,066$)	0.375 (0.136)
Trim top & bottom 0.5%, 2SLS ($N = 12,935$)	0.237 (0.093)
Trim top & bottom 1.5%, 2SLS ($N = 12,679$)	0.219 (0.079)
MS 2011, IVQR ($N = 13,066$)	0.244 (0.057)

Table 1: Estimates of the 2001 rebate coefficient ($\hat{\beta}_2$). Nondurables include food (at home and away), utilities, household operations, public transportation and gas, personal care, alcohol and tobacco, miscellaneous goods, apparel good and services, reading materials, and out-of-pocket health care expenditures. JPS 2006: Johnson, Parker and Souleles (2006); MS 2011: Misra and Surico (2011). 2SLS: Two-Stage Least Squares; IVQR: Instrumental Variable Quantile Regression.

ity in the consumption response across households. Their point estimate is, again, around 0.24. Properly accounting for outliers pushes the rebate coefficient towards the low end of the original JPS estimates and, reassuringly, increases their precision. To facilitate the comparison between model and data, it is useful to focus on one number, and we take 0.25 as our preferred estimate.

Interpretation It is crucial to understand the exact meaning of the rebate coefficient. The estimated coefficient β_2 in equation (1) measures the consumption growth for the treatment group (the rebate recipients at date t) relative to consumption growth of the control group of non-recipients, with the common consumption growth component being subsumed by the time dummies. The control group is composed of those who are already aware of the policy, but will receive the check at a later date, and those who have already received the payment in the past. Thus, the consumption response of the control group, which ideally should be unaffected by the policy is, generally, a mix of the MPC out of the news and the lagged MPC out of the payment. Thus, what exactly does β_2 measure?

To simplify the analysis, we split the population into two groups: early recipients (group A) who receive the check in 2001:Q2 and late recipients (group B) who receive it in 2001:Q3. Let Δc_t^g be consumption growth of group g in quarter t . Then, β_2 is the average of (i) consumption growth of the treatment group in Q2 (group A who receives the check in Q2) net of Q2 consumption growth of the control group (group B who receives the check in Q3) and (ii) consumption growth of the treatment group in Q3 (group B) net of Q3 consumption growth of the control group (group A who received

	Quarter 2 (Q2)		Quarter 3 (Q3)	
	Group A	Group B	Group A	Group B
Surprise for group A	Δc to surprise check	Δc to news	Lagged Δc to surprise check	Δc to anticipated check
Anticipated by all	Δc to anticipated check	0	Lagged Δc to anticipated check	Δc to anticipated check
Surprise for all	Δc to surprise check	0	Lagged Δc to surprise check	Δc to surprise check

Table 2: Economic interpretation of the components of the rebate coefficient β_2 in equation (2) under the three alternative information structures.

the check in Q2), i.e.

$$\beta_2 = \frac{(\Delta c_{Q2}^A - \Delta c_{Q2}^B) + (\Delta c_{Q3}^B - \Delta c_{Q3}^A)}{2}. \quad (2)$$

Consider now three alternative information structures: (i) the policy is announced in 2001:Q1, every consumer becomes aware of it at that date, and thus no consumer is surprised by the check upon receipt; (ii) the policy enters agents' information sets only when the check is actually received, and hence every consumer is surprised by the arrival of the check; (iii) an intermediate structure where the policy enters all agents' information sets after the first batch of checks is sent out (2001:Q2), i.e., group A is surprised, but group B is not. Table 2 describes the economic interpretation of each component Δc_t^g under these three informational assumptions, when β_2 is estimated as in equation (1).

In the case where the policy is fully anticipated by all households, the rebate coefficient β_2 cannot be properly interpreted as an MPC out of the (anticipated) extra income because the consumption growth of the control group A in Q3 incorporates the lagged reaction to the check received in Q2.¹¹ For the same reason, in the case where the policy is a surprise for all, β_2 cannot be interpreted as an MPC out of an unexpected income shock.¹² Interestingly, in both cases one can fully take care of this problem by

¹¹The response of group B in Q2 is the lagged consumption response to the news received in Q1. For unconstrained households it is zero, as they responded already in Q1, and for constrained households is also zero because they have not received the rebate yet.

¹²In this case, one can infer the true MPC out of a surprise check from the consumption response of the earliest recipients.

modifying the specification of equation (1) as

$$\Delta c_{it} = \sum_s \beta_{0s} \text{month}_s + \beta'_1 \mathbf{X}_{i,t-1} + \beta_2 R_{it} + \beta_3 R_{i,t-1} + \varepsilon_{it} \quad (3)$$

because the lag of the rebate variable absorbs the lagged consumption response.¹³ In the intermediate information case, the interpretation of the rebate coefficient is further muddled by the fact that the consumption growth of the control group B in Q2 incorporates the reaction to the news, and thus the addition of the lagged rebate in the regression does not fully resolve the problem.

In spite of these difficulties in mapping directly β_2 to a MPC, we maintain that the rebate coefficient is an informative statistic: only if the true MPC out of the check is sizable and the MPC out of the news is small, can the rebate coefficient be as large as is empirically estimated. The advantage of the structural model is that it enables one to identify all the separate components of equation (2). As a result, it allows one to quantify the current and lagged MPCs out of an income shock, out of an anticipated income change, and out of the news of a future change in income — all magnitudes that are essential for policy analysis.

3 A life-cycle model with liquid and illiquid assets

Our framework integrates the Baumol-Tobin inventory-management model of money demand into an incomplete-markets life-cycle economy. We first describe the full model; next, we use a series of examples to highlight the economic mechanisms at work.

3.1 Model description

Demographics The stationary economy is populated by a continuum of households, indexed by i . Age is indexed by $j = 1, 2, \dots, J$. Households retire at age J^w and retirement lasts for J^r periods.

Preferences Households have an Epstein-Zin-Weil objective function defined recursively by

$$V_{ij} = \left[(1 - \beta) \left(c_{ij}^\phi s_{ij}^{1-\phi} \right)^{1-\sigma} + \beta \{ \mathbb{E}_j [V_{i,j+1}^{1-\gamma}] \}^{\frac{1-\sigma}{1-\gamma}} \right]^{\frac{1}{1-\sigma}} \quad (4)$$

¹³In JPS and PSJM, the baseline specification is equation (1). This augmented specification with one or more lags is used by the authors to calculate the cumulative effect of the rebate over several months.

where $c_{ij} \geq 0$ is consumption of nondurables and $s_{ij} \geq 0$ is the service flow from housing for household i at age j . The parameter β is the discount factor, ϕ measures the weight of nondurables relative to housing services in period-utility, γ regulates risk aversion, and $1/\sigma$ is the elasticity of intertemporal substitution.¹⁴

Idiosyncratic earnings In any period during the working years, household labor earnings (in logs) are given by

$$\log y_{ij} = \chi_j + \alpha_i + z_{ij}, \quad (5)$$

where χ_j is a deterministic age profile common across all households, α_i is a household-specific fixed effect, and z_{ij} is a stochastic idiosyncratic component that obeys the conditional c.d.f. $\Gamma^z(z_{j+1}, z_j)$.

Assets Households can hold a liquid asset m_{ij} and an illiquid asset a_{ij} . The illiquid asset pays a gross financial return $1/q^a$, whereas positive balances of the liquid asset pay $1/q^m$. When the household wants to make deposits into, or withdrawals from, the illiquid account, it must pay a transaction cost κ .¹⁵ The trade-off between these two savings instruments is that the illiquid asset earns a higher return, in the form of capital gain and consumption flow, but its adjustments are subject to the transaction cost. Households start their working lives with an exogenously given quantity of each asset.

Illiquid assets are restricted to be always non-negative, $a_{ij} \geq 0$. Because of the prevalence of housing among commonly held illiquid assets (see Section 5), we let the stock of illiquid assets a_{ij} yield a utility flow with proportionality parameter $\zeta > 0$. Households are also free to purchase or rent out housing services $h_{ij} \geq -\zeta a_{ij}$ on the market.¹⁶ As a result, $s_{ij} = \zeta a_{ij} + h_{ij}$.

We allow borrowing in the liquid asset to reflect the availability of unsecured credit up to an ad-hoc limit, $\underline{m}_{j+1}(y_{ij})$ expressed as a function of current labor earnings. The

¹⁴Piazzesi et al. (2007) offer both (i) microevidence from CEX on the variation of housing expenditure share across different household types, and (ii) time-series evidence on the relationship between the aggregate expenditure share and the relative price of housing services. Both dimensions of the data suggest an elasticity of substitution between nondurable and housing consumption very close to one, which is the Cobb-Douglas case that we adopt in our preference specification.

¹⁵It is straightforward to allow for a utility cost or a time cost proportional to labor income rather than a monetary cost of adjustment. We have experimented with both types of costs and obtained similar results in both cases. See Kaplan and Violante (2011).

¹⁶This assumption adds realism to the model. Technically, it is useful because, with our Cobb-Douglas period utility specification, housing services are an essential consumption good and, without a rental market, even the poorest households would be forced to pay the transaction cost in order to deposit into the illiquid account to start enjoying a minimum amount of housing services.

interest rate on borrowing is denoted by $1/\bar{q}^m$ and we define the function $q^m(m_{i,j+1})$ to encompass both the case $m_{i,j+1} \geq 0$ and $m_{i,j+1} < 0$.

Financial returns to the liquid and illiquid assets, as well as the borrowing rate, are exogenous. Two reasons dictate the choice of abstracting from the equilibrium determination of returns. First, the total outlays from the 2001 rebate amounted to less than 0.1% of aggregate net worth, surely not enough to move asset prices significantly. Second, 83% of aggregate wealth is held by the top quintile of the distribution (Diaz-Gimenez et al., 2011, Table 6), and the portfolio allocation of such households is unlikely to be affected by the receipt of a \$500 check from the government.¹⁷

Government Government expenditures G are not valued by households. Retirees receive social security benefits $p(\chi_{J^w}, \alpha_i, z_{iJ^w})$ where the arguments proxy for average gross lifetime earnings. The government levies proportional taxes on consumption expenditures (τ^c) and on asset income (τ^a, τ^m), a payroll tax $\tau^{ss}(y_{ij})$ with an earnings cap, and a progressive tax on labor income $\tau^y(y_{ij})$. There is no deduction for interest paid on unsecured borrowing. We denote the combined income tax liability function as $\mathcal{T}(y_{ij}, a_{ij}, m_{ij})$. For retirees, the same tax function applies with y_{ij} taking the value $p(\cdot)$. Finally, we let the government issue one-period debt B at price q^g .

Household problem We use a recursive formulation of the problem. Let $\mathbf{s}_j = (m_j, a_j, \alpha, z_j)$ be the vector of individual states at age j . The value function of a household at age j is $V_j(\mathbf{s}_j) = \max\{V_j^0(\mathbf{s}_j), V_j^1(\mathbf{s}_j)\}$, where $V_j^0(\mathbf{s}_j)$ and $V_j^1(\mathbf{s}_j)$ are the value functions conditional on not adjusting and adjusting (i.e., depositing into or withdrawing from) the illiquid account, respectively. This decision takes place at the beginning of the period, after receiving the current endowment shock, but before consuming.¹⁸

Consider a household of age j . If the household chooses not to adjust its illiquid assets because $V_j^0(\mathbf{s}_j) \geq V_j^1(\mathbf{s}_j)$, it solves the dynamic problem:

¹⁷In simulations, the aggregate stock of illiquid wealth increases by only 0.14% during the first year of the transition, an amount hardly large enough to have an impact on the rate of return.

¹⁸Because of this timing, after the earnings shock the household can always choose to pay the transaction cost, access the illiquid account, and use all its resources to finance consumption. Hence, our model *does not* feature a cash-in-advance (CIA) constraint. See Jovanovic (1982) for an exhaustive discussion of the difference between models with transaction costs and models with CIA constraints.

$$\begin{aligned}
V_j^0(\mathbf{s}_j) &= \max_{c_j, h_j, m_{j+1}} \left[(1 - \beta) \left(c_j^\phi s_j^{1-\phi} \right)^{1-\sigma} + \beta \{ \mathbb{E}_j [V_{j+1}^{1-\gamma}] \}^{\frac{1-\sigma}{1-\gamma}} \right]^{\frac{1}{1-\sigma}} \\
&\text{subject to} \quad : \\
(1 + \tau^c) (c_j + h_j) + q^m (m_{j+1}) m_{j+1} &= y_j + m_j - \mathcal{T}(y_j, a_j, m_j) \\
s_j &= h_j + \zeta a_j \\
q^a a_{j+1} &= a_j \\
c_j \geq 0, \quad h_j \geq -\zeta a_j, \quad m_{j+1} \geq -\underline{m}_{j+1}(y_j) \\
y_j &= \begin{cases} \exp(\chi_j + \alpha + z_j) & \text{if } j \leq J^w \\ p(\chi_{J^w}, \alpha, z_{J^w}) & \text{otherwise} \end{cases}
\end{aligned} \tag{6}$$

where z_j evolves according to the conditional c.d.f. Γ_j^z .

The household that adjusts its holding of illiquid assets because $V_j^0(\mathbf{s}_j) < V_j^1(\mathbf{s}_j)$ solves:

$$\begin{aligned}
V_j^1(\mathbf{s}_j) &= \max_{c_j, h_j, m_{j+1}, a_{j+1}} \left[(1 - \beta) \left(c_j^\phi s_j^{1-\phi} \right)^{1-\sigma} + \beta \{ \mathbb{E}_j [V_{j+1}^{1-\gamma}] \}^{\frac{1-\sigma}{1-\gamma}} \right]^{\frac{1}{1-\sigma}} \\
&\text{subject to} \quad : \\
(1 + \tau^c) (c_j + h_j) + q^m (m_{j+1}) m_{j+1} + q^a a_{j+1} &= y_j + m_j + a_j - \kappa - \mathcal{T}(y_j, a_j, m_j) \\
s_j &= h_j + \zeta a_j \\
c_j \geq 0, \quad h_j \geq -\zeta a_j, \quad m_{j+1} \geq -\underline{m}_{j+1}(y_j), \quad a_{j+1} \geq 0 \\
y_j &= \begin{cases} \exp(\chi_j + \alpha + z_j) & \text{if } j \leq J^w \\ p(\chi_{J^w}, \alpha, z_{J^w}) & \text{otherwise} \end{cases}
\end{aligned} \tag{7}$$

The problem for the retired household of age $j > J^w$ is analogous, with pension benefits $p(\cdot)$ in place of earnings y_j . Appendix E describes the computational algorithm used to solve problems 6 and 7.

Balanced budget The government always respects its intertemporal budget constraint

$$G + \sum_{j=J^w+1}^J \int p(y_{J^w}) d\mu_j + \left(\frac{1}{q^g} - 1 \right) B = \tau^c \sum_{j=1}^J \int c_j d\mu_j + \sum_{j=1}^J \int \mathcal{T}(y_j, a_j, m_j) d\mu_j \tag{8}$$

where μ_j is the distribution of households of age j over the individual state vector \mathbf{s}_j .

4 Hand-to-mouth households in model and data

In this section we first illustrate, by means of numerical examples, how hand-to-mouth behavior arises endogenously in our model, even when agents hold positive illiquid wealth. Next, we measure hand-to-mouth households in the Survey of Consumer Finances.

4.1 Behavior in the model: the “wealthy hand-to-mouth”

For ease of exposition, we focus on a stylized version of the model with time-separable preferences ($\gamma = \sigma$), without service flow from illiquid assets ($\phi = 1, \zeta = 0$), with logarithmic period utility, deterministic labor income ($z_j = 0$), and no taxes ($\mathcal{T}(\cdot) = \tau^c = 0$). Moreover, we assume that $\bar{q}^m < q^a < q^m$. The second inequality states that the illiquid asset has a higher return and the first one ensures that households do not borrow to deposit into the illiquid account.

Two Euler equations Consumption and portfolio decisions are characterized by a short-run Euler equation (EE-SR) that corresponds to borrowing or saving in the liquid asset, and a long-run Euler equation that corresponds to (dis)saving in the illiquid asset (EE-LR). In periods where the working household does not adjust:

$$u'(c_j) = \frac{\beta}{q^m(m_{j+1})} u'(c_{j+1}). \quad (\text{EE-SR})$$

The slope of her consumption path is governed by $\beta/q^m(m_{j+1})$. For plausible parameterizations, when the household is in debt ($m_{j+1} < 0$) this slope is above one: the consumption path is increasing as the household saves his way out of expensive borrowing. When the household is saving ($m_{j+1} > 0$) this slope is below one: consumption declines over time because of impatience and the low real return on cash. There are two kinks in the budget constraints where equation (EE-SR) does not hold: $m_{j+1} = -\underline{m}_{j+1}(y_j)$, the debt limit, and $m_{j+1} = 0$, because of the wedge between the return on liquid saving and the interest on unsecured credit ($\bar{q}^m < q^m$). Households on the kinks are hand to mouth, i.e., consume all their income.

During the working life, an agent will eventually want to save to finance consumption in retirement by making deposits into the illiquid account. Given the fixed cost of adjusting, households accumulate liquid funds and choose infrequent dates at which to add some or all of their liquid holdings to the illiquid asset (the “cake-baking” problem). Across two such adjustment dates N periods apart, consumption dynamics

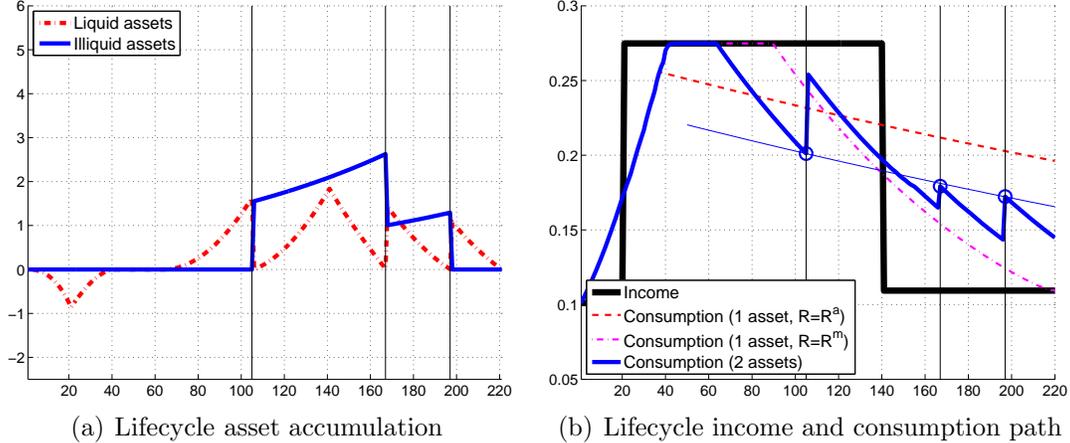


Figure 1: Example of lifecycle of a poor hand-to-mouth agent in the model

are dictated by

$$u'(c_j) = \left(\frac{\beta}{q^a}\right)^N u'(c_{j+N}). \quad (\text{EE-LR})$$

Since $\beta/q^a > \beta/q^m$, consumption grows more (or falls less) across adjustment dates, than between adjustments.

During retirement, the household faces a cake-eating problem, where optimal decisions closely resemble those in Romer (1986). Consumption in excess of pension income is financed by making periodic withdrawals from the illiquid account. Between each withdrawal, the household runs down its liquid holdings and consumption falls according to (EE-SR). The withdrawals are timed to coincide with the period where cash is exhausted. Equation (EE-LR) holds across withdrawals.

Poor hand-to-mouth behavior Figure 1 shows consumption and wealth dynamics in an example where an agent starts her working life with zero wealth, receives an increasing endowment while working, and a constant endowment when retired. To make this example as stark as possible, we impose a very large transaction cost. Panel (a) shows that, because of the increasing earnings profile, the agent in this example chooses first to borrow to smooth consumption, and then starts saving for retirement. She adjusts her illiquid account at only three points in time: one deposit while working, after repaying her debt, and two withdrawals in retirement. After its inception, the value of the illiquid account grows at rate $1/q^a$.¹⁹

¹⁹Over the working life, the household piles up liquid funds in anticipation of her deposit into the liquid account, but also to smooth consumption across her transition into retirement. As we show in Appendix C.4, this pattern of accumulation of liquid wealth around retirement survives in the richer model with heterogeneity and uncertainty and is also distinctly visible in the micro data.

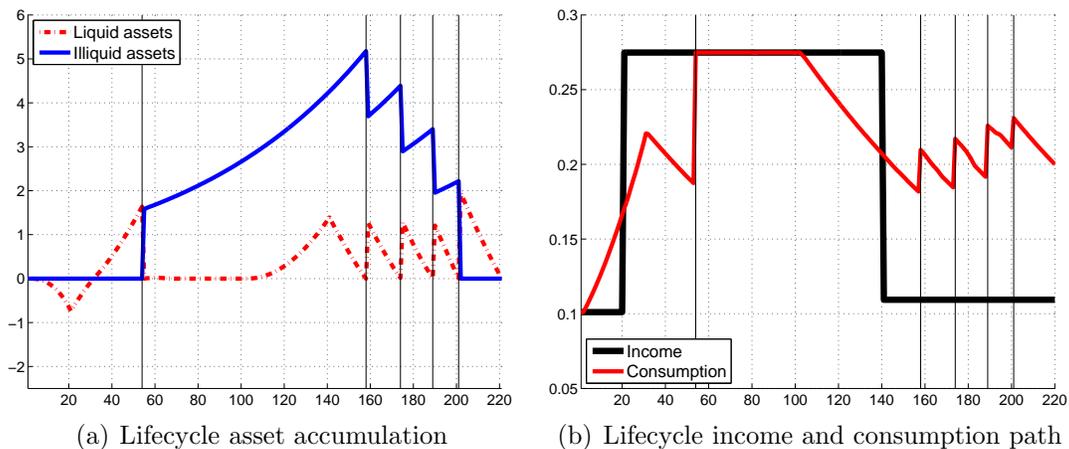


Figure 2: Example of lifecycle of a wealthy hand-to-mouth agent in the model

Panel (b) shows her associated earnings and consumption paths. In the same panel, we have also plotted the paths for consumption arising in the two versions of the corresponding one-asset model: one with the short-run interest rate $1/q^m$ (m_{j+1}), and one with the long-run rate $1/q^a$. The sawed pattern for consumption that arises in the two-asset model is a combination of the short-run and long-run behavior: between adjustment dates the consumption path is parallel to the path in the one-asset model with the low return; while across adjustment dates, the slope is parallel to consumption in the one-asset model with the high return. Finally note that, after repayments of her debts, this agent is *poor* hand-to-mouth. In other words, she keeps zero net worth and consumes all her income for a phase of her life, before starting to save.

Wealthy hand-to-mouth behavior Figure 2 illustrates how the model can feature households with positive net worth who consume their income every period: the *wealthy* hand-to-mouth agents. The parameterization is the same as in Figure 1, except for a higher return on the illiquid asset. This higher return leads to stronger overall wealth accumulation, but rather than increasing the number of deposits during its working life, the household changes the timing of its single deposit: the deposit into the illiquid account is now made earlier in life in order to take advantage of the high return for a longer period (compare the left panels across Figures 1 and 2). Thus, the household optimally chooses to hold zero liquid assets in the middle of the working life, after her deposit, while the illiquid asset holdings are positive and are growing in value. Intuitively, since her net worth is large, this household would like to consume more than her earnings flow, but the transaction cost and the high interest rate on unsecured borrowing dissuade her from doing so. This is a household that, upon receiving the rebate, will consume a large part of it and, upon the news of the rebate,

will not increase her expenditures.

Why would households choose to consume all of their earnings and deviate from the optimal consumption path imposed by the short-run Euler equation (EE-SR), even for long periods of time? The answer is that households are better off taking this welfare loss because avoiding it entails either (i) paying the transaction cost more often to withdraw cash in order to consume more than income; or (ii) holding larger balances of liquid wealth and hence foregoing the high return on the illiquid asset (and, therefore, the associated higher level of long-run consumption); or (iii) using expensive unsecured credit to finance expenditures.²⁰ We note that this logic is reminiscent of Cochrane’s (1989) insight that the utility loss from setting consumption equal to income is second-order in a representative agent model with reasonable risk aversion and income volatility. Browning and Crossley (2001) report similar calculations in the context of a life-cycle one-asset model of consumption and saving.

4.2 The SCF data

We begin with some descriptive statistics about household portfolios in the Survey of Consumer Finances (SCF). We then explain how we exploit these data to estimate the proportion of hand-to-mouth households in the US.

Households’ portfolio data Our data source is the 2001 wave of the SCF, a triennial cross-sectional survey of the assets and debts of US households. For comparability with the CEX sample in JPS (2006), we exclude the top 5% of households by net worth. Average (median) labor income for the working-age population is \$52,745 (\$41,000), a number close to the one reported by JPS (2006, Table 1).²¹ Our definition of liquid assets comprises: cash, money market (MM), checking, savings and call accounts as well as directly held mutual funds (MF), stocks, bonds, and T-Bills net of revolving debt on credit card balances. In Appendix B.1 we describe our identification of revolving debt and our cash imputation procedure, needed because the SCF does not record household cash holdings.²²

²⁰While we have focused our examples on poor and wealthy hand-to-mouth behavior at the kink for zero liquid wealth, there is a second type of hand-to-mouth behavior when agents borrow up to the credit limit. This limit is the second kink in the budget constraint. In this case, option (iii) is obviously not feasible. In Appendix A, we illustrate an example of wealthy hand-to-mouth behavior at the credit limit.

²¹In our definition of household labor income, we include unemployment and disability insurance, TANF, and child benefits.

²²Briefly, our cash imputation uses data from the Survey of Consumer Payment Choice administered by the Federal Reserve Bank of Boston. To calculate revolving unsecured debt, we use a combination of different SCF questions. This strategy, which is common in the literature (see Telyukova, 2011),

	Median (\$2001)	Mean (\$2001)	Fraction Positive	Return (%)
Earnings plus benefits (age 22-59)	41,000	52,745	–	–
Net worth	62,442	150,411	0.90	1.7
Net liquid wealth	2,629	31,001	0.77	-1.5
Cash, checking, saving, MM accounts	2,858	12,642	0.92	-2.2
Directly held MF, stocks, bonds, T-Bills	0	19,920	0.29	1.7
Revolving credit card debt	0	1,575	0.41	–
Net illiquid wealth	54,600	119,409	0.93	2.3
Housing net of mortgages	31,000	72,592	0.68	2.0
Retirement accounts	950	34,455	0.53	3.5
Life insurance	0	7,740	0.27	0.1
Certificates of deposit	0	3,807	0.14	0.9
Saving bonds	0	815	0.17	0.1

Table 3: Household Portfolio Composition. Authors’ calculations based on the 2001 Survey of Consumer Finances (SCF). The return reported in the last column is the real after-tax risk-adjusted return. MM: money market; MF: mutual funds. See Appendix B.1 for additional details.

Our baseline measure of illiquid assets includes housing net of mortgages and home equity loans, retirement accounts (e.g., IRA, 401K), life insurance policies, CDs, and saving bonds. Table 3 reports some descriptive statistics.

As expected, the bulk of household wealth is held in illiquid assets, notably housing and retirement accounts. For example, the median of the liquid and illiquid asset distributions are \$2,629 and \$54,600, respectively. Moreover, over their working life households save disproportionately through illiquid wealth and keep holdings of liquid wealth fairly stable: median illiquid assets grow by around \$100,000 from age 30 to retirement, whereas median liquid wealth increases by \$5,000 or less.

Measurement of hand-to-mouth households In the model, we define a household to be hand to mouth (hereafter, HtM) if she chooses to be at one of the kinks of her budget constraint, either zero liquid wealth or the credit limit. Such a household will have a high marginal propensity to consume out of an extra dollar of windfall income. How can we identify these HtM households in the SCF data?

To measure HtM households at the zero kink for liquid wealth, we start from the avoids including purchases made through credit cards in between regular payments as debt.

observation that, since these households do not borrow and do not save through liquid assets, they do not carry any liquid wealth *across* pay periods. If we observed liquid balances at the end of the period in the data, we could easily identify these HtM agents, but the SCF reports only the *average* liquid balance during the last month. Average balances are positive for all households (HtM and not) because labor income is paid as liquid assets and because of a mismatch in the timing of consumption and earnings *within* a pay period. Then, a strict criterion to identify these HtM agents in the data is to count those households in the SCF whose average balance of liquid wealth is equal to or less than half their earnings per pay-period. The “half” presumes resources being consumed at a constant rate).²³ Symmetrically, we measure HtM agents at the credit limit as those SCF households with negative holdings of liquid wealth that are lower than half their pay-period earnings minus their self-reported total credit limit.

Any sample split based on income and liquid wealth is bound to contain both type I and type II classification error (see, e.g., Jappelli, 1990). Nevertheless, our estimate is likely to be a lower bound because, while all *non* HtM households would always hold average liquid balances above half their earnings, some HtM households at the zero kink may fall in this latter group as well.²⁴

The examples in Section 4.1 show that there are two types of HtM agents. There are poor HtM agents without any illiquid assets, and wealthy HtM agents who have positive balances of illiquid wealth. In the SCF we identify wealthy HtM agents as those households who satisfy the HtM requirements listed above and, at the same time, hold illiquid assets.

Appendix B.2 contains more details on this measurement. There, we also perform a robustness analysis with respect to the frequency of the pay-period (weekly, bi-weekly, monthly), the definition of liquid wealth (whether it only includes cash and bank accounts or also directly held stocks and bonds), and the definition of illiquid wealth (whether it also includes vehicles), and the definition of wealthy HtM (whether the HtM household holds at least \$3,000 in its illiquid account, which is the median amount of liquid wealth).

²³Alvarez and Lippi (2009) suggest this calculation as a test of the liquidity management model.

²⁴If the household starts the period with some savings in addition to earnings and ends the period with some savings, its average balance would be above half earnings. If its initial balance equals only earnings for that period and it ends the period with positive savings, the average balance would also be above half earnings. Neither of these households are HtM. However, if a household starts the period with positive savings in addition to earnings and ends the period with zero liquid savings, its average liquid balance would be above half earnings, but she is a HtM household in that period.

Our estimates imply that between 17.5% and 35% of households are HtM in the US. Among these, between 40 and 80 percent are wealthy HtM, depending mainly on the pay frequency and on whether one expands the notion of illiquid wealth by including vehicles. This group of wealthy HtM households, which represents a sizeable fraction of the population (between 7% and 26%), is only visible through the lens of the two-asset model. From the distorted point of view of the standard one asset model, these are households with positive net worth, and are hence unconstrained. It is useful to compare these estimates with those that one would obtain when HtM agents are measured in terms of net worth.²⁵ We compute that between 4% and 14% of US households are HtM in terms of net worth, depending largely on whether vehicles are considered part of wealth.

Because of the lower bound nature of our estimator, in the model we target a total fraction of HtM households on the high end of the range, around 1/3 of the population. This target is also consistent with three additional pieces of survey evidence. First, the SCF asks households whether “in the past year their spending exceeded their income, but did not spend on a new house, a new vehicle, or on any investment.” Almost 36% of households fall into this category. Second, Lusardi, Schneider, and Tufano (2011) document that around 1/3 of US households would “certainly be unable to cope with a financial emergency that required them to come up with \$2,000 in the next month.” The authors also report that, among those giving that answer, a high proportion of individuals are at middle class levels of income. Similarly, Broda and Parker (2012) document, from the AC Nielsen Homescan database, that 40% of households report that they do not have “at least two months of income available in cash, bank accounts, or easily accessible funds”.

5 Calibration

Demographics and initial asset positions Decisions in the model take place at a quarterly frequency. Households begin their active economic life at age 22 ($j = 1$) and retire at age 60 ($J^w = 152$). The retirement phase lasts for 20 years ($J^r = 80$). We use observed wealth portfolios of SCF households aged 20 to 24 to calibrate the age $j = 0$ asset positions in the model. Our procedure also targets the observed correlation between initial earnings, liquid wealth, and illiquid wealth.²⁶

²⁵We define HtM households in terms of net worth in the same way. A household is HtM (in terms of net worth) if it has (i) positive net worth below half its earnings per pay-period, or (ii) negative net worth lower than half its earnings minus its credit limit.

²⁶See Appendix C.1 for details.

Preferences We calibrate the discount factor β to replicate median illiquid wealth as a fraction of average income in the SCF.²⁷ The annualized value of β is 0.941, and hence our results are not driven by an implausibly low discount factor that makes households highly impatient. We set the coefficient of relative risk aversion γ to 4 and the elasticity of intertemporal substitution ($1/\sigma$) to 1.5.²⁸ Finally, we set $\phi = 0.85$ to match the ratio of expenditures on housing services to total consumption expenditures in the National Income and Product Account, which is around 15 percent on average over the period 1960-2009. In Section 6 we discuss the robustness of our results to this parameterization of preferences.

Appendix C.2 explains in detail how we compute the service flow from housing which maps into the parameter ζ . In short, we account for the fact that owning housing wealth has both costs (maintenance, insurance, property taxes, and mortgage interests) and benefits (imputed rental value of the space and tax deductibility of mortgage interests and property taxes). From this, we arrive at a conservative estimate for ζ of 1 percent per quarter. Since the median ratio of gross housing wealth to net illiquid assets in the SCF is around one, we apply ζ to the entire stock a_j .

²⁷In the literature on quantitative macroeconomic models with heterogeneous households and incomplete markets, there are two approaches to calibrating the discount factor. The first is to match median wealth (e.g., Carroll, 1992, 1997). The second is to match average wealth (e.g., Aiyagari, 1994; Rios-Rull, 1995; Krusell and Smith, 1998). There is a trade off in this choice. Matching median wealth allows one to reproduce the wealth distribution more closely for the vast majority of households, with the exception of the upper tail that holds a large portion of total assets. Matching average (and aggregate) wealth allows one to fully incorporate equilibrium effects on prices at the cost of overstating wealth holdings and understating the MPC for a large fraction of households (due to the concavity of the consumption function, see Carroll and Kimball, 1996). We choose the former approach because for the question at hand, a plausible distribution of MPCs across the population is far more important than aggregate price effects.

²⁸We have chosen a value of the intertemporal elasticity of substitution above one based on theoretical and empirical grounds. Two recent promising approaches to account for asset pricing facts – the long-run risk hypothesis and the rare disasters model – point towards a high willingness to substitute intertemporally. Bansal and Yaron (2004) show that to replicate the estimated consumption volatility effects on price-dividend ratios, one needs an elasticity above one. In the context of the rare disasters literature, Barro (2009) makes the analogous observation that an intertemporal elasticity below one has the counterfactual implication that a rise in the probability (or the size) of a disaster increases asset prices. The literature examining the empirical magnitude of this elasticity based on aggregate time series leads to a wide range of estimates. As discussed at length in Bansal et al. (2012, section 4.6), low estimates are typically obtained by estimating the elasticity as the slope coefficient from a regression of consumption growth on the real interest rate. This traditional approach can lead to severely downward biased estimates because of attenuation bias (when the real rate is measured with error) or endogeneity bias (when omitted variables are correlated with the real rate or when consumption volatility is time-varying). To deal with endogeneity, Gruber (2006) uses cross-individual differences in after-tax real interest rates that derive from arguably exogenous differences in capital income tax rates and estimates an elasticity around 2. In general, when a GMM approach is used instead of the regression approach (with a larger set of moment restrictions including, for example, other asset market data), the values for this elasticity are well above one (Hansen et al. 2007).

Earnings heterogeneity From the Panel Study of Income Dynamics (PSID), we construct a sample of households with 22-59 year-old heads in 1969-1996, following the same selection criteria as in Heathcote, Perri, and Violante (2010). We use a fourth-order polynomial in age to extract the common life-cycle earnings profile χ_j . Since the residual variance from this regression rises almost linearly with age, we model z_{ij} as a unit root process with quarterly variance of the innovation equal to 0.003 to match the total increase over the age range we consider. The variance of the individual fixed effect (α_i) is set to 0.18 to reproduce the dispersion of initial earnings at age 22.

Asset returns We measure financial returns on liquid and illiquid wealth in four steps. First, we compute returns on each individual asset class over the period 1960-2009. Second, we perform a risk-adjustment on each of these returns that acknowledges the fact that in our model there is no aggregate uncertainty. Third, we apply these risk-adjusted returns and the corresponding capital income tax rates to each individual household portfolio in the SCF, and compute the average return on liquid and illiquid wealth, (and net worth for the one asset version of our model) in the population. The average risk-adjusted after-tax real returns we obtain are -1.48% for liquid wealth, 2.29% for illiquid wealth, and 1.67% for net worth (see Table 3). Appendix C.3 reports details of these calculations.

Credit limit and borrowing rate The SCF asks households to report their total credit limit. The median ratio of credit limit to quarterly labor income for households aged 22 to 59 is 74% . For working-age households, we therefore specify the function $\underline{m}_{j+1}(y_j)$ as $\underline{m} \cdot y_j$, with $\underline{m} = 0.74$. For retirees the borrowing limit is set to zero.

The interest rate on unsecured debt $1/\bar{q}^m$ is set so that the model reproduces the fraction of borrowers in the data. In the SCF, one could define borrowers in two ways: (i) as households with negative *net* liquid wealth, or (ii) as households with credit card debt, independent of their balances on checking accounts, saving accounts, etc. Around 17% of working-age households are borrowers according to (i) and 37% according to (ii). The second definition is more conventional, but the first one is the exact counterpart of borrowers in the model, since the model only speaks to net holdings of liquid wealth.²⁹ We target a fraction of borrowers in the middle of this range. At a nominal borrowing rate of 10% (or 6% real), 26% of agents have $m_{j+1} < 0$ in the model. The implied wedge between the unsecured borrowing cost and real after-tax return on liquid assets

²⁹The model is not designed to tackle the so-called “credit card puzzle” (i.e., households who have positive balances of liquid wealth and credit card debt at the same point in time). Telyukova (2011) documents the extent of this puzzle in the data and proposes a solution based on the the existence of certain “cash” good expenditures whose size is unpredictable.

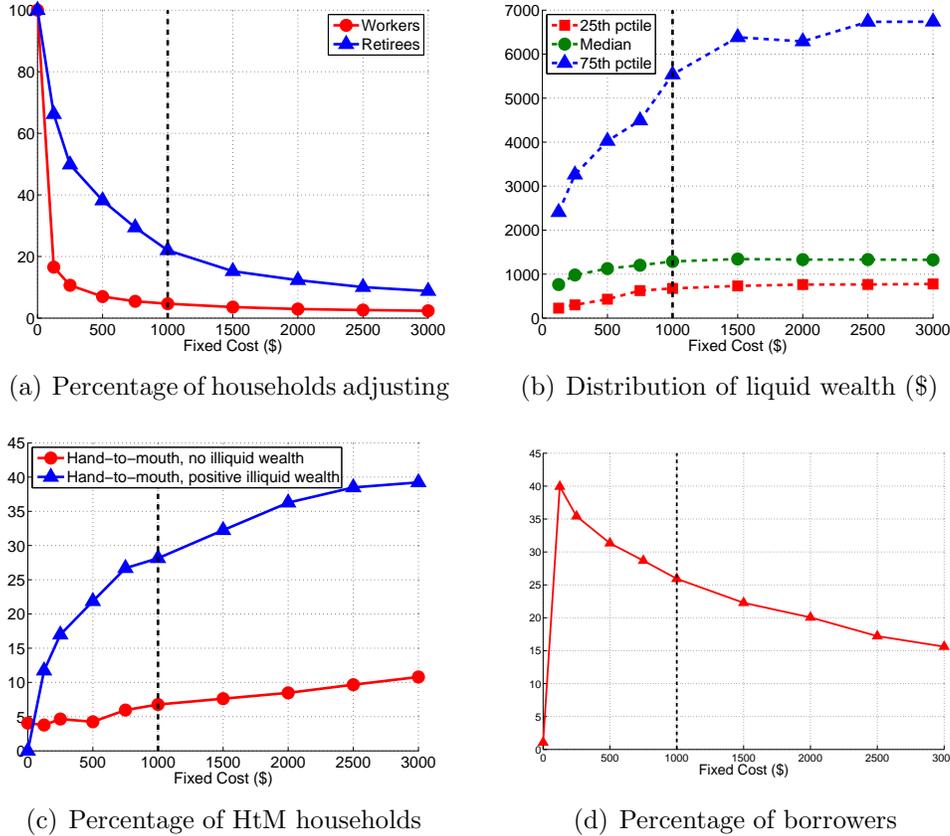


Figure 3: Features of two-asset model, by transaction cost

(6% + 1.5% = 7.5%) is in line with estimates by Davis et al. (2006) who report wedges between 6.5% and 8.5% for the period 1991-2001.

Transaction cost Because of the lack of systematic evidence on transaction costs, we set the value of κ to match the proportion of hand-to-mouth households in the data. For a value of $\kappa = \$1,000$, the model implies that roughly 1/3 of agents in the model are (poor and wealthy) hand-to-mouth, consistently with the estimates presented in Section 4.2. We note that this value of κ corresponds to 0.9% of the stock of illiquid assets, on average, for adjusting households.³⁰

Figure 3 displays some features of the model as a function of κ . For each value of $\kappa > 0$, we re-calibrate β to match median holdings of illiquid wealth. Panel (a) shows that the fraction of households adjusting -accessing the illiquid account to withdraw or deposit-

³⁰Transaction costs for housing are commonly estimated around 5% of the asset value (e.g., OECD, 2011). Alvarez, Guiso, and Lippi (2012, Table 5) report transaction costs on durables of the order of 1%. Individual retirement accounts are subject to set-up costs and penalties for early distributions (typically, 10% of the amount withdrawn). In light of these estimates, our value of κ appears reasonable.

falls with the size of the transaction cost κ . As illustrated in the simulations of Section 3, retirees adjust more often than working-age households because they finance their consumption largely by withdrawing from the illiquid account. At $\kappa = \$1,000$, 4.5% of workers and 21% of retirees adjust each quarter. Holdings of liquid wealth increase with the transaction cost (panel (b)), because when κ is larger households deposit into or withdraw from the illiquid account less often and carry larger balances of liquid assets. However, even for large transaction costs, median liquid wealth remains small. Liquid balances are more sensitive to κ at the upper end of the distribution since, in that range, transaction costs have more of an impact on the optimal frequency of adjustment. Panel (c) plots the fraction of hand-to-mouth consumers in the model and divides them into those who also have zero illiquid wealth and those with positive illiquid wealth. The size of both groups is increasing in κ . At $\kappa = \$1,000$, the split between poor and wealthy hand-to-mouth, roughly 1/5 and 4/5, is in line with the data presented in Section 4.2. Panel (d) shows how the fraction of borrowers in the model declines with κ . This result is the mirror image of our findings of panel (b): as κ grows, households hold larger liquid balances and respond to negative shocks by dissaving rather than by taking up debt.

Taxes and social security benefits The consumption expenditure tax τ^c is set to 7.2% (McDaniel, 2007). We specify the tax function $\mathcal{T}(y_j, a_j, m_j)$ as a sum of four components: (i) a progressive tax on labor income $\tau^y(y_j)$ modelled as a smooth approximation to the estimates in Kiefer et al. (2002, Table 5) who report effective tax rates on wage income for ten income brackets in the year 2000; (ii) a payroll tax $\tau^{ss}(y_j)$ set to 12.4% up to an earnings cap of 0.5 times average annual earnings, in order to reproduce the Old-Age, Survivors, and Disability Insurance (OASDI) tax rates in 2000; (iii) a tax of 23.2% on income from liquid assets (τ^m), and (iv) a tax of 7.9% on income from illiquid assets (τ^a).³¹ The implied tax rate on capital income from net worth is 10.4%. To compute social security benefits, our proxies for individual average lifetime earnings $y_{iJw} = \exp(\chi_{Jw} + \alpha_i + z_{iJw})$ are run through a formula based on replacement rates and bend points as in the actual system in the year 2000.

Calibration of one-asset model For the one-asset model: (i) we set β to reproduce median net worth; (ii) the interest rate is the average after-tax real return on

³¹Kiefer et al. (2002, Table 5) also report the effective tax schedule on interests and dividends, and on long term capital gains, by ten income brackets for the year 2000. We apply these tax schedules to each household portfolio in our SCF sample, and take the average to compute τ^m and τ^a . We follow the same strategy to compute the tax on capital income from net worth and obtain 10.4%. See Appendix C.3 for more details.

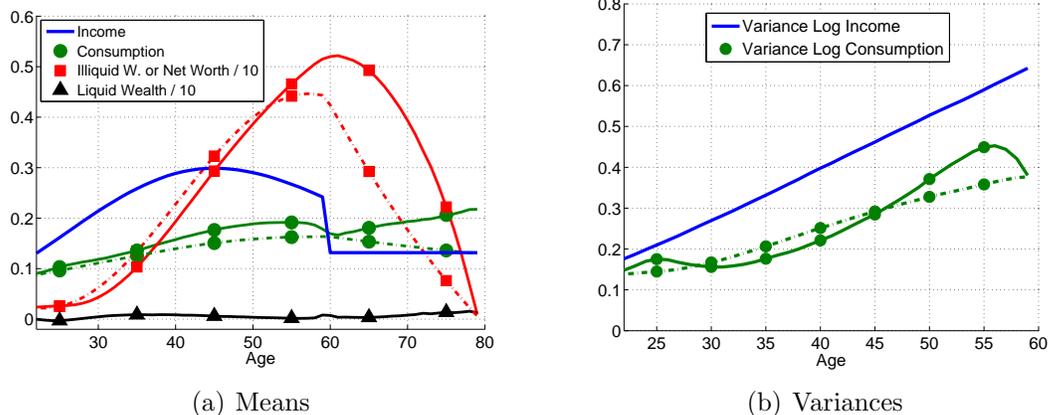


Figure 4: Left panel: Mean income, nondurable consumption, net worth (in the one-asset model), and liquid and illiquid wealth (in the two-asset model). Right panel: Variance of log income and nondurable consumption. Dashed lines: one-asset model. Solid lines: two-asset model.

net worth in the SCF data (see Table 3); (iii) the parameter ζ , which measures the consumption flow from housing, is applied to the entire stock of net worth; and (iv) the credit limit remains at 74% of quarterly household income.

Life-cycle profile Figure 4 compares the life-cycle means and variances of labor income, nondurable consumption, and wealth across the one-asset and two-asset models. Panel (a) shows that the path of average consumption is very similar in the two models, except during the retirement phase. In the two-asset model, because of the high rate of return on the illiquid asset, the long-run Euler equation (EE-LR) dictates that consumption should grow across withdrawals, which induces an upward trend in consumption (see, e.g., Figure 2(b)). Both models produce a hump shape in net worth/illiquid wealth.³² Panel (b) of Figure 4 shows that consumption inequality from middle age to retirement grows somewhat faster in the two-asset economy. In that phase of the life-cycle, most of a household’s wealth is held in the illiquid asset, which is seldom used for consumption smoothing. Overall, both models reproduce the key features of the data reasonably well (see Heathcote et al., 2010; Kaplan, 2012), and would be hardly distinguishable based on life cycle data on income, consumption, and net worth, given the noise present in typical cross-sectional household surveys.

³²The two-asset model has a slightly higher average wealth-to-income ratio, but the same median wealth-to-income ratio by calibration.

6 The tax rebate experiment

We now reproduce the 2001 tax rebate episode within our economic model.

Experiment design The economy is in a steady state in 2001:Q1. The rebate checks are randomly sent out to half the eligible population in 2001:Q2 (group A), and to the other half in 2001:Q3 (group B). The size of the rebate is set to \$500 based on JPS (2006) who report that the average rebate check was \$480 per household. We assume that the news/check reaches households before making their consumption/saving and adjustment decisions for that quarter. The government finances the rebate program by increasing debt, and after ten years it permanently increases the payroll tax to gradually repay the accumulated debt (plus interest).³³

Building on our discussion of Section 2, one could take different views about the timing of when the rebate enters households' information sets. At one extreme, households become fully aware of the rebate when the Bill is discussed in Congress and enacted. This scenario implies that the news arrives in 2001:Q1 and the check is thus fully anticipated by both groups. At the other extreme, households become aware of the rebate only after receiving their own check and thus both groups of households treat the rebate as a surprise. An intermediate view is that all households learn about the rebate in 2001:Q2, when the first batch of Treasury checks is received. Under this timing, the check is a surprise for group A, but it is fully anticipated by group B since they receive the check in 2001:Q3.

What information structure is the closest approximation to reality? Survey data are typically not rich enough to identify when the rebate enters households' information sets. An important exception is a recent paper by Broda and Parker (2012) which studies the consumption response to the fiscal stimulus payment of 2008. The authors conduct a survey of roughly 60,000 households in the Nielsen Consumer Panel and, among other questions, ask when the surveyed household learned about the rebate. They document that 60% of households knew about the policy the quarter before payments began to be disbursed.³⁴ Moreover, they show that even those households who learned in advance did not have a significant spending response before receipt of their payment. The first finding offers support for the intermediate informational

³³We have experimented with other lengths of time before the tax rate is increased to repay the rebate outlays, and with a case where the rebate is entirely financed by expenditure cuts. These choices have no quantitative bearing on the results.

³⁴The Bill was passed by Congress in February, and payments begun in late April. 60% of households responded they learned in February or March.

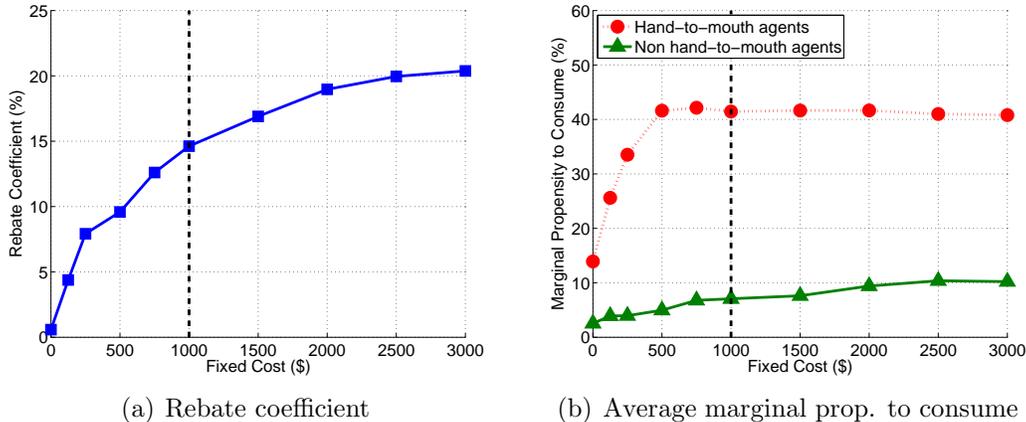


Figure 5: Rebate coefficient and marginal propensity to consume, by transaction cost

assumption; the second for the view that the policy is, effectively, a surprise for all households. We choose the intermediate timing as our baseline and explore the other two alternative timing assumptions later in this section.

We start by studying an economy where the tax rebate occurs in isolation. In Section 7, we incorporate two features of 2001’s macroeconomic environment: the broader income tax reform and the recession.

Baseline results Figure 5(a) displays the rebate coefficient in the model for a range of the transaction cost between zero and \$3,000. The rebate coefficient is computed through regression (1) run on simulated panel data, exactly as in JPS (2006). The rebate coefficient grows steadily from 0.6% at $\kappa = 0$ (the one-asset model) to 20% at $\kappa = \$3,000$. For $\kappa = \$1,000$, the calibrated value of the transaction cost, the model generates a rebate coefficient of 15% or nearly 2/3 of the empirical estimate. Figure 5(b) shows the model’s MPC out of the unanticipated fiscal stimulus payment (i.e., the consumption response of group A in 2001:Q2) for two types of households: those who are hand-to-mouth and those who are not. Note how the average MPC is over 40% for the HtM, while for the non HtM it is only 7%. Therefore, the vast majority of households in the model behave as predicted by the PIH and have small MPCs. The high rebate coefficient is entirely driven by HtM households. Such households have significant MPCs out of the rebate check (when they are in the treatment group) and do not respond to the news of the check (when they are in the control group).

Figure 5(a) also displays the powerful amplification mechanism intrinsic in the two-asset model: the rebate coefficient is 14 percentage points larger than its one-asset model counterpart ($\kappa = 0$). This amplification works through both an extensive and

	Δc_{Q2}^A	Δc_{Q2}^B	Δc_{Q3}^A	Δc_{Q3}^B	β_2
Baseline	0.20	0.06	-0.09	0.07	0.15
Anticipated by all	0.07	0.00	-0.08	0.07	0.11
Surprise for all	0.20	0.00	-0.09	0.20	0.25

Table 4: Breakdown of the model’s rebate coefficient into different component for the three different informational assumptions. $\Delta c_{Q_t}^g$ denotes consumption growth of group $g \in \{A, B\}$ at quarter $t \in \{2, 3\}$. The last column is the rebate coefficient (β_2) computed as $[(\Delta c_{Q2}^A - \Delta c_{Q2}^B) + (\Delta c_{Q3}^B - \Delta c_{Q3}^A)] / 2$

an intensive margin. First, the two-asset model features a much larger fraction of HtM consumers, many of which hold sizeable quantities of illiquid assets.³⁵ Second, even among HtM agents, the wealthy HtM have larger MPCs out of tax rebates than the poor HtM (44% vs 34%) since they have higher wealth (tied in the illiquid asset) and, therefore, higher desired target consumption.

Anatomy of the rebate coefficient Using the expression in equation (2), we now decompose the rebate coefficient into the four components described in Table 2. The term Δc_{Q2}^A (consumption growth of group A in Q2) is the average MPC out of the unexpected \$500 check. Table 4 shows that this component equals 20% (an average of the MPCs of HtM and non HtM agents plotted in Figure 5(b)). The term Δc_{Q2}^B is the MPC out of the news (that a \$500 check will be received next quarter) and equals 6%. The term Δc_{Q3}^A is the lagged consumption growth of group A. This term is negative (−9%) since consumption of group A peaks in Q2 upon receiving the check, after which it declines steeply following the slope dictated by the short-run Euler equation. Finally, the term Δc_{Q3}^B , which equals 7%, is a combination of a large response of the HtM agents in group B net of the consumption drop of the unconstrained agents in group B who already responded to the news in Q2. Averaging out the four components, we obtain (modulo the rounding) the estimated value of the rebate coefficient, 15%.

From this decomposition, we learn three key numbers for policy analysis. In our model, the average quarterly MPC out of a small income shocks is 20%. The average MPC out of an anticipated (one quarter ahead) income change is 6%; and the average MPC

³⁵The fraction of HtM households in the one-asset model ($\kappa = 0$) is 7%, and hence within the range of the estimates obtained from the SCF 2001 (see Table B1). Since β is set to match median net worth, and all other parameters are disciplined directly by the data, the fraction of HtM agents is not an explicit target in the one-asset model. If, instead of targeting median net worth, we set β to reproduce 15% of HtM agents (the upper bound of our estimates), the implied rebate coefficient increases to 2.5%. In conclusion, there is essentially no scope for the one asset model to generate large rebate coefficients, while remaining consistent with SCF data on the distribution of net worth.

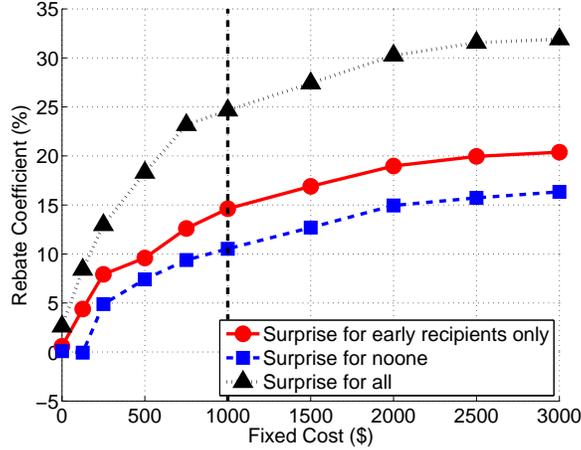


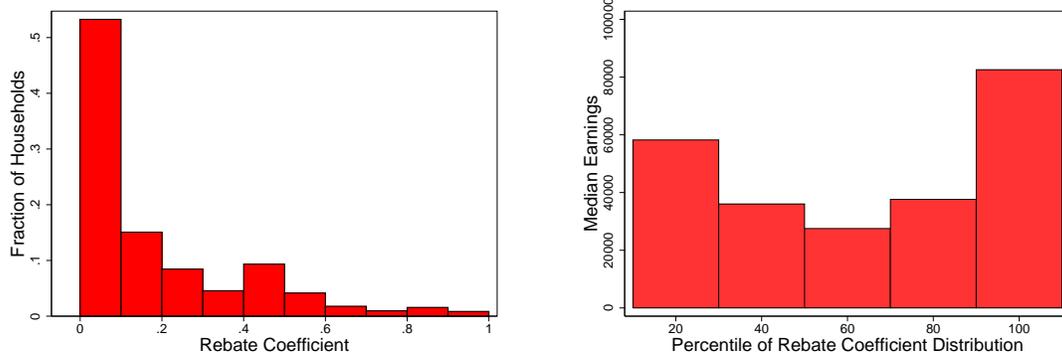
Figure 6: Rebate coefficients under alternative assumptions on timing of arrival of news

out of the announcement (the news) of a future income change is 7%. It is clear that, since the estimated rebate coefficient mixes these three objects, one has to be cautious when directly using its empirical estimate for policy analysis.

Alternative information structures In Figure 6, we report the model’s rebate coefficient under alternative assumptions about when the news of the rebate enters households’ information sets. When the rebate is anticipated by all households (the news arrives in Q1, i.e., one quarter ahead of the check for the first group and two quarters ahead for the second group), the estimated rebate coefficient drops by 4 percentage points compared to the baseline. Non HtM households (2/3 of the population) increase their consumption upon arrival of the news and not when they receive the check either one or two quarters later. However, the rebate coefficient remains of a sizable magnitude, around 11% for $\kappa = \$1,000$, and, most importantly, the amplification with respect to the one-asset model (where the rebate coefficient is now 0.1%) is still very large. The reason is that liquidity constrained households are those responsible for the amplification mechanism in the two-asset model and learning about the policy ahead of time does not affect their behavior.

When the policy is a surprise for all (i.e., households learn about the policy only upon receiving their check), the rebate coefficient increases significantly relative to the baseline. At $\kappa = \$1,000$, the model-implied rebate coefficient reaches 25%, the same magnitude as its empirical counterpart. Under this information structure, the control group who receives the check in Q3 cannot respond to the news in Q2, like it does in the baseline. The absence of this anticipation effect raises the model’s rebate coefficient.

This analysis reinforces our point that the rebate coefficient is not a MPC. The rebate



(a) Distribution of rebate coefficients in the population

(b) Median earnings by quintile of rebate coefficient distribution

Figure 7: Heterogeneity in rebate coefficients in the model ($\kappa = \$1,000$)

coefficient varies between 11% and 25%, depending on how households process information, but as is clear from Table 4, the MPC out of the unexpected fiscal stimulus payment is always 20%. Therefore, the rebate coefficient may underestimate or overestimate the true MPC. Only a structural model can help disentangle the true MPC from the empirical results of regressions such as (1).

Heterogeneity The stark dichotomy in the MPC of HtM and non HtM agents documented in Figure 5(b) suggests that our model features a large amount of heterogeneity in consumption responses to fiscal stimulus payments across households. Figure 7(a) plots the distribution of rebate coefficients in the model: almost half of households in the model have consumption responses close to zero, 15% spend more than half the rebate in the quarter they receive it, and the remaining third are in between. Misra and Surico (2013) apply quantile regression techniques to the JPS (2006) data to estimate the empirical cross-sectional distribution of consumption responses to the 2001 rebate. Their results line up remarkably well with the model predictions. They estimate that between 40% and 50% of U.S. households have responses that are statistically indistinguishable from zero; another 20% of households have rebate coefficients that are significantly above one half; and the remaining households fall somewhere in between.

Misra and Surico (2013, Figure 5) also document that high income households are disproportionately concentrated in the two tails of the distribution of consumption responses, a finding that rationalizes two former results in the literature. JPS (2006) report that, when splitting the population into three income groups, differences in rebate coefficient across groups are not statistically significant. Similarly, Shapiro and Slemrod (2003a, 2003b) find no evidence of higher spending rates among low income

households. Figure 7(b) shows that our model can replicate the bimodality of the income distribution by size of the rebate coefficient. The reason why there are high earnings households at both ends of the distribution in the model, is that some of them are unconstrained (those at the bottom end) and some are wealthy HtM (those at the top end). In particular, because the rebate is a lump sum, among wealthy HtM agents the income-richest have the highest MPCs.³⁶

Correlation with liquid wealth The model predicts that households carrying low levels of liquid wealth *across pay periods*, i.e., the HtM households, should have strong consumption responses. Although it is not possible to construct an analogous measure in the data, an imperfect proxy can be obtained by grouping households based on *liquid wealth-to-income ratios*. This is because for a HtM household, the quantity of liquid assets that are held for within pay-period expenditures is, on average, half its income. Broda and Parker (2012) split households in two groups and find very strong (and statistically significant) evidence that households with a low ratio of liquid assets to income spend at least twice as much as the average household, precisely as predicted by our model. Souleles (1999) studies the consumption response to anticipated tax refunds (whose median size is around \$560). When the sample is split between low and high liquid wealth to income ratio households, the former are found to have statistically significant larger responses to the refund (Souleles, 1999, Table 4).³⁷

Size asymmetry Figure 8 shows how, in our baseline economy, the rebate coefficient declines with the size of the rebate. With a \$1,000 transaction cost, the rebate coefficient drops by over a factor of two (from 15% to 6%) as the size of the stimulus payment increases from \$500 to \$2,000. A large enough rebate loosens the liquidity constraint, and even constrained households find it optimal to save a portion of their

³⁶A further validation of our mechanism comes from another finding in Misra and Surico (2013): in contrast to the high income households at the bottom of the distribution, those at the top tend to have high mortgage debt. They therefore do own illiquid wealth in the form of housing, and their large interest payments mean that they are likely to be wealthy HtM households.

³⁷JPS (2006) estimate rebate coefficients for sub-groups of households with different amounts of liquid assets and they do find stronger responses for the group with less than \$1,000 in liquid wealth. These effects are imprecisely estimated, though, for three reasons. First, the sample becomes very small when divided into sub-groups. Second, the asset data in the CEX must be viewed with extreme caution, due to the large amount of item non-response. JPS (2006) have data on liquid wealth for less than half of the sample, and hence it is likely that respondents are a highly selected group. Third, households hold liquid wealth both to finance consumption expenditures *within* pay periods, and to save *across* pay periods. Therefore, even hand-to-mouth households will be observed to hold positive, and possibly large, quantities of liquid wealth if they are sampled at a point in time between pay dates, as done in the CEX. Therefore, empirically, the relationship between rebate coefficients and the *level* of liquid wealth can be statistically weak. As explained, the liquid wealth-to-income ratio may be more informative.

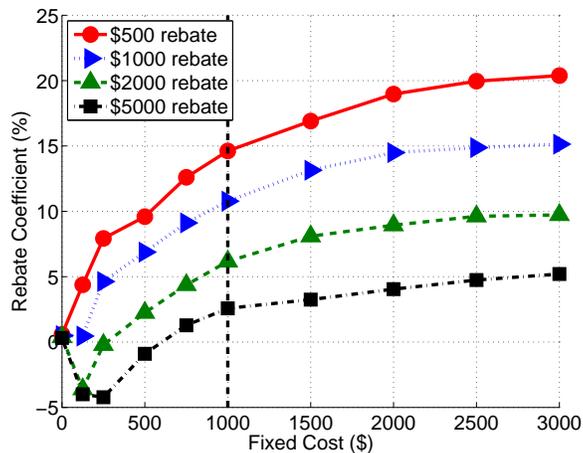


Figure 8: Rebate coefficients by stimulus payment size

payment. Moreover, for rebates that are sufficiently large relative to the transaction cost, many working households will choose to pay the transaction cost and make a deposit upon receipt of the rebate. But adjusting households are unconstrained, so they save a large portion of the rebate, as in the one asset model. Figure 8 also shows how estimated rebate coefficients (but not the MPC) may become negative when the stimulus payment is large relative to the transaction cost. In this case, many working households choose to make a deposit into the illiquid account upon receipt of the payment. As a result, these households consume even less than the control group during that period. The finding that the rebate coefficient falls with the size of the payment is mirrored by the behavior of the true MPC out of a surprise payment: as the check grows from \$500 to \$5,000, this MPC drops from 20 to 3 percent.

Our mechanism’s size asymmetry feature is consistent with two well-known empirical findings. Hsieh (2003) shows that the same CEX consumers who “overreact” to small income tax refunds respond very weakly to much larger payments (around \$2,000 per household) received from the Alaskan Permanent Fund. Browning and Collado (2001) document similar evidence from Spanish survey data: workers who receive anticipated double-payment bonuses (hence, again, large amounts) in the months of June and December do not alter their consumption growth significantly in those months. Our interpretation of these findings is that although households spend substantial portions of small anticipated income changes, they predominantly save large ones, since only large enough payments trigger an adjustment.

Robustness Appendix D contains an extensive sensitivity analysis with respect to preference parameters (risk aversion and IES), access to credit (borrowing costs and

limits), desirability of the illiquid asset (financial return and consumption flow), and size of idiosyncratic risk. One of the main findings is the role played by the IES. Households who are more willing to substitute consumption intertemporally are more likely to save heavily in the illiquid asset during working-age (and thus to be wealthy HtM) in order to enjoy higher consumption at retirement. Quantitatively, the effects are substantial: doubling the IES from 1 to 2 more than doubles the rebate coefficient.

7 Role of aggregate economic conditions

We now incorporate two features of 2001’s macroeconomic environment into the analysis: the broader income tax reform and the recession. These additional experiments also highlight that our model features a strong aggregate state-dependence of the consumption response to fiscal stimulus payments: same-size rebates distributed under different economic conditions can have different effects.

2001 Tax reform The 2001 rebate was part of a broader tax reform which, beyond decreasing the lowest rate, also reduced all other marginal rates by 3% or more. We construct the sequence of effective tax schedules implied by the reform based on Kiefer et al. (2002).³⁸ These changes were phased in gradually over the five years 2002:Q1-2006:Q1 and planned to “sunset” in 2011. A tax reform is defined as a sequence of income tax schedules $\{\mathcal{T}_t\}_{t=t^*}^{t^{**}}$ which is announced, jointly with the rebate, in 2001:Q2. Date t^* , the first quarter of the change in the tax code, is 2002:Q1. Date t^{**} , the last quarter of the change in the tax code, is 2011:Q1 when the tax reform sunsets, as originally legislated. The tax cut is deficit-financed for ten years, after which the payroll tax is increased permanently (by roughly 0.2%) to gradually reduce the debt to its pre-reform level.³⁹

Figure 9(a) shows the consumption responses to the tax rebate when the baseline economy is augmented with the tax reform. The fall in future tax liabilities leads to a rise in the desired level of lifetime consumption which, in turn, triggers two offsetting forces. On the one hand, households who are already borrowing sizeable amounts may reach their credit limit, which tends to increase the number of HtM households in the economy. On the other hand, HtM households at the zero kink may start borrowing

³⁸Kiefer et al. (2002, Table 5) report the pre and post reform income tax rates, and describe the timing of the reduction in the various brackets.

³⁹Instead of sunseting as originally planned, subsequent legislation further extended the tax cuts. An alternative scenario, where the tax cuts expire later yields almost identical results. Similarly, when the tax cuts are funded by lower expenditures, the model’s rebate coefficient is unchanged.

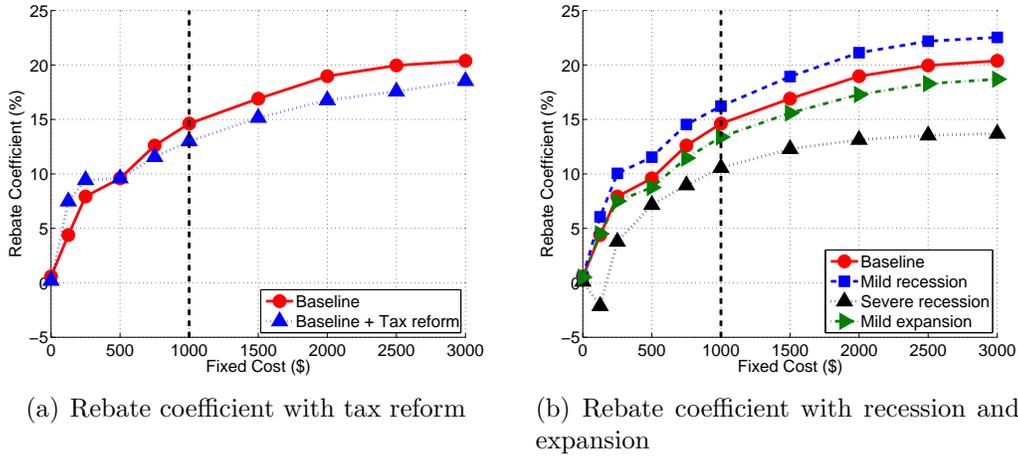


Figure 9: Effect of tax reform and aggregate economic conditions on rebate coefficient

and, once off the kink, they have low MPCs out of the rebate. For low transaction costs, when there are already lots of households borrowing (see panel (d) of Figure 3), the first channel dominates, and the rebate coefficient is slightly higher than in the baseline. However, for higher transaction costs, the second channel appears to be stronger. At $\kappa = \$1,000$, one year after the tax reform the fraction of households using credit is twice the initial one. Overall, the fraction of HtM agents is much lower and, as a consequence, the rebate coefficient drops by roughly 2 points.⁴⁰

2001 Recession To model the downturn of 2001, we assume that at the onset of 2001:Q2 households become aware that they are entering a recession. At this time they learn that their labor income will fall evenly for the next three quarters, generating a cumulative drop of 3%, and will then fully recover at a constant rate over the following eight quarters.⁴¹ Figure 9(b) shows that the occurrence of a mild recession, such as the 2001 episode, increases the number of hand-to-mouth households in the economy and adds nearly 2 percentage points to the rebate coefficient.

State dependence Figure 9(b) also shows that the consumption response to the rebate is highly dependent on the aggregate economic conditions. For example, when the rebate is distributed during a mild expansion (of the same size of the mild recession

⁴⁰To further understand the importance of credit for these effects, we simulated an economy without borrowing ($\underline{m}_{j+1} = 0$). Here, the tax reform increases the rebate coefficient by 7-8 percentage points relative to the baseline experiment. The reason is that the announced tax cuts exacerbate liquidity constraints, and the government transfer enables HtM households to start consuming immediately out of the additional future disposable income.

⁴¹The NBER dates the 2001 recession as starting in March 2001 and ending in November 2001. The magnitude of the downturn and the duration of its recovery are calibrated from HP-filtered quarterly GDP (NIPA Table 1.1.6).

of 2001, with the sign reversed, and of the same duration), the consumption response is more muted in the model. Since most episodes of fiscal stimulus payments occur in recessions, it is difficult, empirically, to isolate the role of aggregate economic conditions on the size of the consumption response. A unique piece of evidence is offered by JPS (2009) who examine the impact of the child tax credit of 2003, a period of sustained growth. Their point estimates of the contemporaneous response of consumption for the 2003 episode are about half of those estimated for 2001 in similar specifications (although not statistically different). This leads these authors to conjecture “a more potent response to such payments in recessions, when liquidity constraints are more likely to bind, than during times of more typical economic growth.” Our model offers a mechanism why this force may be at work and quantifies its significance.

The state dependence is, however, quite complex. A central, and novel, implication of our model is that the aggregate consumption response to a stimulus payment can decrease with the severity of the recession. Recall that wealthy HtM behavior is optimal to the extent that the welfare gain from smoothing consumption (by tapping into the illiquid account) is small enough relative to the transaction cost and the foregone return. The size and expected duration of the income drop caused by the recession affects this trade-off. A sufficiently sharp recession leads many wealthy HtM households to pay the transaction cost and withdraw from their illiquid account in order to avoid an abrupt dip in consumption. Similarly, the poor HtM at the zero liquid wealth kink start using credit heavily to sustain their consumption. As a result, many households who were HtM before the recession become effectively unconstrained at the time of the rebate, and their consumption response to the transfer can be quite low. In Figure 9(b), we report the results of a rebate handed out during a severe downturn (5 times deeper than the mild recession examined before). Two quarters into the downturn, the fraction of households who have used credit or have withdrawn from their their illiquid account since the start of the sharp recession is almost twice as large as in the mild recession case, and the rebate coefficient is 6 percentage points lower.⁴²

Aggregate impact of the policy When we run the tax rebate experiment within an environment which combines both the tax reform and the recession, the economic forces discussed in this section tend to balance out, and the rebate coefficient falls only slightly (by roughly half a percentage point) relative to the baseline.⁴³

⁴²A similar drop, from 19% in the mild recession to 13% in the severe recession, is observed in the true MPC out of a surprise payment.

⁴³Combining the tax reform and recession leads to minor changes in rebate coefficients also for the other two information structures. In the case where the policy is anticipated by all, the rebate

	2001:Q1	2001:Q2	2001:Q3	2001:Q4
Baseline	0	0.13	0.22	0.28
Anticipated by all	0.06	0.19	0.26	0.30
Surprise for all	0	0.10	0.25	0.32

Table 5: Cumulative aggregate impact of the policy measured as the fraction of the total rebate outlays spent on nondurable consumption within the year 2001, in the model with both tax reform and recession.

Within this macroeconomic environment, we exploit our structural model to quantify the impact of the 2001 fiscal stimulus payments on aggregate nondurable consumption expenditures. Table 5 summarizes the results. We find that, in the model, households spend around 30% of the total rebate outlays (\$38B) by the fourth quarter of 2001, independently of the assumed information structure. However, the time path of expenditures during 2001 is obviously affected by the exact timing of when households are assumed to become aware of the policy.

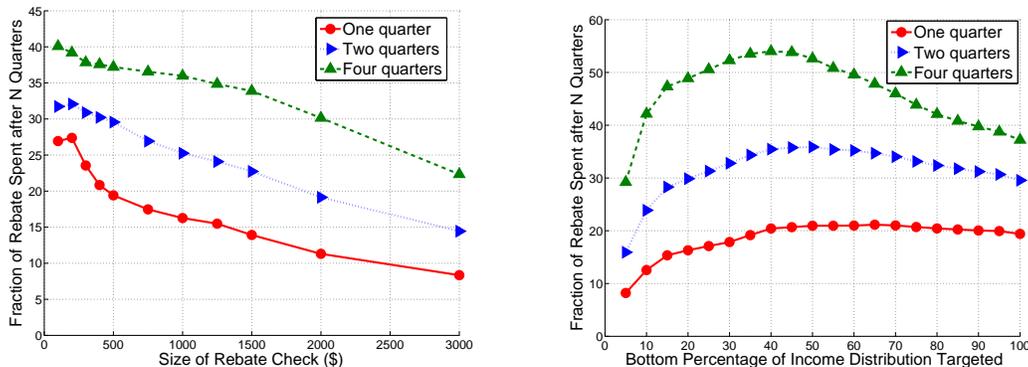
8 Implications for stimulus policy design

The main lesson from our model is that the sizable estimated response of aggregate consumption to fiscal stimulus payments is largely attributable to the behavior of HtM households, many of which are wealthy HtM. This conclusion has implications for policy design. A government that aims at stimulating consumption expenditures in the short-run (the declared objective of such policies) should recognize that (i) increasing the magnitude of the stimulus will not raise household expenditures proportionately, and (ii) targeting, whenever possible, the group of wealthy HtM households in the population will yield stronger effects. In this section, we illustrate these prescriptions in more detail by running two policy experiments.⁴⁴

Stimulus size In the first experiment, we compute the fraction of the rebate spent at different short-run horizons (1, 2, and 4 quarters) for transfers of different sizes, starting at \$100 up to \$2,500 per household. Larger fiscal stimulus checks clearly induce larger household expenditures, but as explained in Section 6, our model displays a strong size-dependence due to the infrequent adjustment of the illiquid asset: larger

coefficient increases by 1.5 percentage points, and in the case where the rebate is a surprise for all it increases by 3 percentage points.

⁴⁴To keep the policy experiments simple, we assume that (i) the policy is a surprise for all households, and (ii) all the rebates are paid at the same time. All our qualitative results are robust to using the baseline (i) information structure and (ii) staggering of payments.



(a) Experiment on the size of the stimulus payment (b) Experiment on the income-targeting of the stimulus payment

Figure 10: Alternative designs of fiscal stimulus policies: implications for aggregate consumption

payments trigger anticipated deposits into the illiquid account, a feature that tends to dampen the short-run consumption response. Figure 10(a) shows that this mechanism is quantitatively significant: increasing the magnitude of the government transfer from \$500 to \$2,000 per household reduces the fraction of the rebate spent by over 10 percentage points at all horizons.

Stimulus targeting In the second experiment, we consider a series of policies with different targeting based on household income that are budget-equivalent to our baseline experiment. For example, when targeted to the bottom half of the income distribution the rebate, is twice as large (\$1,000) as when it is paid to the entire population (\$500). Figure 10(b) plots the percentage of the total outlays (the same in each simulation) spent at different horizons. All the curves are hump shaped. Targeting income-poorer households makes it more likely to reach the HtM agents, but there are two countervailing forces. First, the wealthy HtM are not the income poorest, so an excessively narrow targeting may miss many agents with high MPCs. Second, as the policy targets fewer agents the size of the payments increases, which leads some households to save a large fraction of their transfer into the illiquid asset instead of consuming it.

The implications for policy design are quite stark. A steep phasing out is required for the policy to reach its highest “bang for the buck”: at all horizons, the aggregate consumption response is the largest when the policy is phased out around median income.⁴⁵

⁴⁵Consistently with our findings, Broda and Parker (2012) estimate significantly higher consumption

9 Concluding remarks

By integrating the Baumol-Tobin model with the standard incomplete-markets life-cycle framework, one can provide a theoretical foundation, and a quantitative validation, for the observation that the MPC out of small temporary income changes is large – an empirical finding that is substantiated by quasi-experimental evidence. Going forward, our analysis can be expanded in several directions.

More immediately, the model can be used to analyze the fiscal stimulus payments of 2008. This episode is of particular interest because both PJSM (2011) and Broda and Parker (2012) measure responses in nondurable expenditures around half of the size of the 2001 estimates. The 2008 episode differs from the one studied in this paper in four ways: (i) its magnitude was roughly twice as large; (ii) eligibility phased out quickly starting at \$75,000 of gross individual income; (iii) the 2008 recession was much deeper than its 2001 counterpart; and, (iv) the 2008 episode was not part of any broader tax reform. As explained here, each of these factors matters for households' consumption responses, and only a quantitative analysis that contains all of these ingredients can shed light on what accounted for the more modest effects of the 2008 stimulus program.

Taking a broader view, the framework used in this paper can be seen as the second generation of the spender-saver model of Campbell and Mankiw (1989, 1991). Compared to its original formulation, where the measure of spenders is exogenous and entirely composed of impatient wealth-poor households, here the fraction of hand-to-mouth agents is endogenously determined and mostly composed of patient individuals who own assets tied up in illiquid instruments. This distinction changes some of the key macroeconomic implications of this model. For example, one well-known problem of the model with exogenous spenders is that the volatility of aggregate consumption is too high relative to the data. But in the time-series for aggregate income, there are large and small innovations. While the consumption response of the wealthy hand-to-mouth agents and that of the impatient spender are similar with respect to small shocks, large shocks induce the former type of agents to adjust their portfolio and, as a result, better smooth the change in income.

In a similar vein, major fiscal or monetary policy interventions that influence the relative return between liquid funds and illiquid assets (large public debt expansions or changes in the federal fund rate) will affect the endogenous fraction of wealthy hand-to-mouth consumers in the second-generation models, thereby complicating the analysis

responses to the 2008 stimulus payments for households with income below the median.

of the impact of policy on the macroeconomy.

As just exemplified, some applications of the model cannot abstract from general equilibrium effects on prices. Given the high-frequency OLG structure, solving a version of our two-asset model with aggregate shocks and asset returns determined endogenously is not numerically feasible (see Krueger and Kubler, 2004). To make progress in these directions, one could develop an infinite-horizon version of our economy with a stochastic transition between work and retirement. To close the model, one would interpret the illiquid asset as productive capital with a return equal to its marginal product, and the return on the liquid asset could be pinned down by a monetary policy rule.

References

- ABEL, A., J. EBERLY, AND S. PANAGEAS (2009): “Optimal Inattention to the Stock Market with Information Costs and Transactions Costs,” *NBER Working Paper*, 15010.
- AIYAGARI, S. (1994): “Uninsured Idiosyncratic Risk and Aggregate Saving,” *The Quarterly Journal of Economics*, 109(3), 659–684.
- AIYAGARI, S., AND M. GERTLER (1991): “Asset Returns with Transactions Costs and Uninsured Individual Risk,” *Journal of Monetary Economics*, 27(3), 311–331.
- ALVAREZ, F., A. ATKESON, AND P. KEHOE (2002): “Money, Interest Rates, and Exchange Rates with Endogenously Segmented Markets,” *Journal of Political Economy*, 110(1), 73–112.
- ALVAREZ, F., L. GUISO, AND F. LIPPI (2010): “Durable Consumption and Asset Management with Transaction and Observation Costs,” *American Economic Review*, 102(5), 2272–2300.
- ALVAREZ, F., AND F. LIPPI (2009): “Financial Innovation and the Transactions Demand for Cash,” *Econometrica*, 77(2), 363–402.
- ANGELETOS, G., D. LAIBSON, A. REPETTO, J. TOBACMAN, AND S. WEINBERG (2001): “The Hyperbolic Consumption Model: Calibration, Simulation, and Empirical Evaluation,” *Journal of Economic Perspectives*, 15(3), 47–68.
- BANSAL, R., D. KIKU, AND A. YARON (2012): “An Empirical Evaluation of the Long-Run Risks Model for Asset Prices,” *Critical Finance Review*, 1, 183–221.
- BANSAL, R., AND A. YARON (2004): “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles,” *The Journal of Finance*, 59(4), 1481–1509.
- BARRO, R. J. (2009): “Rare Disasters, Asset Prices, and Welfare Costs,” *American Economic Review*, 99(1), 243–264.
- BRODA, C., AND J. PARKER (2012): “The Economic Stimulus Payments of 2008 and the Aggregate Demand for Consumption,” *Mimeo*.
- BROWNING, M., AND M. COLLADO (2001): “The Response of Expenditures to Anticipated Income Changes: Panel Data Estimates,” *American Economic Review*, 91(3), 681–692.
- BROWNING, M., AND T. CROSSLEY (2001): “The Life-Cycle Model of Consumption and Saving,” *Journal of Economic Perspectives*, 15(3), 3–22.
- CAMPBELL, J., AND N. MANKIW (1989): “Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence,” *NBER Macroeconomics Annual*, 4, 185–216.
- (1991): “The Response of Consumption to Income: A Cross-Country Investigation,” *European Economic Review*, 35(4), 723–756.
- CAMPBELL, J. R., AND Z. HERCOWITZ (2009): “Liquidity Constraints of the Middle Class,” *Federal Reserve Bank of Chicago Working Paper*, 09-20.
- CARROLL, C., AND M. KIMBALL (1996): “On the Concavity of the Consumption Function,” *Econometrica*, 64(4), 981–992.
- CARROLL, C. D. (1992): “The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence,” *Brookings Papers on Economic Activity*, pp. 61–156.

- (1997): “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis,” *Quarterly Journal of Economics*, 112(1), 1–55.
- CHETTY, R., AND A. SZEIDL (2007): “Consumption Commitments and Risk Preferences,” *Quarterly Journal of Economics*, 122(2), 831–877.
- COCHRANE, J. (1989): “The Sensitivity of Tests of the Intertemporal Allocation of Consumption to Near-Rational Alternatives,” *American Economic Review*, 79(3), 319–337.
- CONGRESSIONAL BUDGET OFFICE (2009): “Did the 2008 Tax Rebates Stimulate Short-Term Growth?,” *Economic and Budget Issue Brief*.
- COUNCIL OF ECONOMIC ADVISORS (2010): “The Economic Impact of the American Recovery and Reinvestment Act (Third Quarterly Report),” *Executive Office of the President*.
- DAVID LAIBSON AND ANDREA REPETTO AND JEREMY TOBACMAN (2003): “A Debt Puzzle,” in *Knowledge, Information, and Expectations in Modern Economics: In Honor of Edmund S. Phelps*, ed. by Philippe Aghion and Roman Frydman and Joseph Stiglitz and Michael Woodford, pp. 228–266. Princeton University Press.
- DAVIS, S. J., F. KUBLER, AND P. WILLEN (2006): “Borrowing Costs and the Demand for Equity over the Life Cycle,” *Review of Economics and Statistics*, 88(2), 348–362.
- DEATON, A. (1991): “Saving and Liquidity Constraints,” *Econometrica*, 59(5), 1221–1248.
- DIAZ-GIMENEZ, J., A. GLOVER, AND J.-V. RIOS-RULL (2011): “Facts on the Distributions of Earnings, Income, and Wealth in the United States: 2007 Update,” *Federal Reserve Bank of Minneapolis Quarterly Review*, 31(1), 2–35.
- EROSA, A., AND G. VENTURA (2002): “On Inflation as a Regressive Consumption Tax,” *Journal of Monetary Economics*, 49(4), 761–795.
- GALÍ, J., J. D. LÓPEZ-SALIDO, AND J. VALLÉS (2007): “Understanding the Effects of Government Spending on Consumption,” *Journal of the European Economic Association*, 5(1), 227–270.
- GRUBER, J. (2006): “A Tax-Based Estimate of the Elasticity of Intertemporal Substitution,” *NBER Working Paper Series*, 11945.
- HAMILTON, D. (2008): “A Reexamination of Johnson, Parker, and Souleles 2001 Tax Rebate Estimate,” *Mimeo*.
- HANSEN, L. P., J. HEATON, J. LEE, AND N. ROUSSANOV (2007): “Intertemporal Substitution and Risk Aversion,” *Handbook of Econometrics*, 6, 3967–4056.
- HEATHCOTE, J. (2005): “Fiscal Policy with Heterogeneous Agents and Incomplete Markets,” *Review of Economic Studies*, 72(1), 161–188.
- HEATHCOTE, J., F. PERRI, AND G. VIOLANTE (2010): “Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States 1967-2006,” *Review of Economic Dynamics*, 13(1), 15–51.
- HEATHCOTE, J., K. STORESLETTEN, AND G. VIOLANTE (2009): “Quantitative Macroeconomics with Heterogeneous Households,” *Annual Reviews in Economics*, 1(1), 319–354.
- HSIEH, C. (2003): “Do Consumers React to Anticipated Income Changes? Evidence from the Alaska Permanent Fund,” *American Economic Review*, 93(1), 397–405.

- HUGGETT, M. (1996): “Wealth Distribution in Life-Cycle Economies,” *Journal of Monetary Economics*, 38(3), 469–494.
- HUNTLEY, J., AND V. MICHELANGELI (2011): “Can Tax Rebates Stimulate Spending in a Life-Cycle Model?,” *Mimeo*.
- JOHNSON, D., J. PARKER, AND N. SOULELES (2006): “Household Expenditure and the Income Tax Rebates of 2001,” *American Economic Review*, 96(5), 1589–1610.
- (2009): “The Response of Consumer Spending to Rebates During an Expansion: Evidence from the 2003 Child Tax Credit,” *Mimeo*.
- JOVANOVIC, B. (1982): “Inflation and Welfare in the Steady State,” *Journal of Political Economy*, 90(3), 561–577.
- JUSTINIANO, A., G. E. PRIMICERI, AND A. TAMBALOTTI (2013): “Household Leveraging and Deleveraging,” *NBER Working Papers*, 18941.
- KAPLAN, G., AND G. L. VIOLANTE (2011): “A Model of the Consumption Response to Fiscal Stimulus Payments,” *NBER Working Papers*, 17338.
- KIEFER, D., R. CARROLL, J. HOLTZBLATT, A. LERMAN, J. MCCUBBIN, D. RICHARDSON, AND J. TEMPALSKI (2002): “The Economic Growth and Tax Relief Reconciliation Act of 2001: Overview and Assessment of Effects on Taxpayers,” *National Tax Journal*, 55(1), 89–118.
- KRUEGER, D., AND F. KUBLER (2004): “Computing Equilibrium in OLG Models with Stochastic Production,” *Journal of Economic Dynamics and Control*, 28(7), 1411–1436.
- KRUSELL, P., AND A. SMITH, JR. (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, 106(5), 867–896.
- LI, G. (2009): “Transaction Costs and Consumption,” *Journal of Economic Dynamics and Control*, 33(6), 1263–1277.
- LUSARDI, A., D. SCHNEIDER, AND P. TUFANO (2011): “Financially Fragile Households: Evidence and Implications,” *Brookings Papers on Economic Activity*, 42(1), 83–150.
- MCDANIEL, C. (2007): “Average Tax Rates on Consumption, Investment, Labor and Capital in the OECD: 1950–2003,” *Mimeo*.
- MISRA, K., AND P. SURICO (2013): “Consumption, Income Changes and Heterogeneity: Evidence From Two Fiscal Stimulus Programmes,” *CEPR Discussion Paper*, 9530.
- OH, H., AND R. REIS (2011): “Targeted Transfers and the Fiscal Response to the Great Recession,” *NBER Working Paper*, 16775.
- ORGANISATION FOR ECONOMIC CO-OPERATION AND DEVELOPMENT (2011): “Housing and the Economy: Policies for Renovation,” *Economic Policy Reforms 2011: Going for Growth*.
- PARKER, J., N. SOULELES, D. JOHNSON, AND R. MCCLELLAND (2011): “Consumer Spending and the Economic Stimulus Payments of 2008,” *NBER Working Paper*, 16684.
- PIAZZESI, M., M. SCHNEIDER, AND S. TUZEL (2007): “Housing, Consumption and Asset Pricing,” *Journal of Financial Economics*, 83(3), 531–569.
- RAGOT, X. (2011): “The Case for a Financial Approach to Money Demand,” *Mimeo*.

- REIS, R. (2006): “Inattentive Consumers,” *Journal of Monetary Economics*, 53(8), 1761–1800.
- RÍOS-RULL, J.V. (1995): “Models with Heterogeneous Agents,” in *Frontiers of Business Cycles Research*, ed. by T. F. Cooley, pp. 98–125. Princeton University Press.
- ROMER, D. (1986): “A Simple General Equilibrium Version of the Baumol-Tobin Model,” *Quarterly Journal of Economics*, 101(4), 663–685.
- SHAPIRO, M., AND J. SLEMROD (2003a): “Consumer Response to Tax Rebates,” *American Economic Review*, 93(1), 381–396.
- (2003b): “Did the 2001 Tax Rebate Stimulate Spending? Evidence from Taxpayer Surveys,” *Tax Policy and the Economy*, 17, 83–109.
- (2009): “Did the 2008 Tax Rebates Stimulate Spending?,” *American Economic Review*, 99(2), 374–379.
- SOULELES, N. (1999): “The Response of Household Consumption to Income Tax Refunds,” *American Economic Review*, 89(4), 947–958.
- TELYUKOVA, I. (2013): “Household Need for Liquidity and the Credit Card Debt Puzzle,” *Review of Economic Studies*, 80(3), 1148–1177.

APPENDIX NOT FOR PUBLICATION

This Appendix is organized as follows. Section A illustrates another case of wealthy hand to mouth behavior which completes the discussion of Section 4.1 in the main text. Section B contains more details on our definitions of liquid and illiquid wealth from the SCF and on the measurement of hand-to-mouth households. Section C describes certain steps of the model’s calibration omitted from the main text. Section D contains a robustness analysis on the baseline rebate experiment. The numerical computation of the model and the simulation of the policy experiments are delineated in Section E.

A Wealthy hand-to-mouth behavior in the model

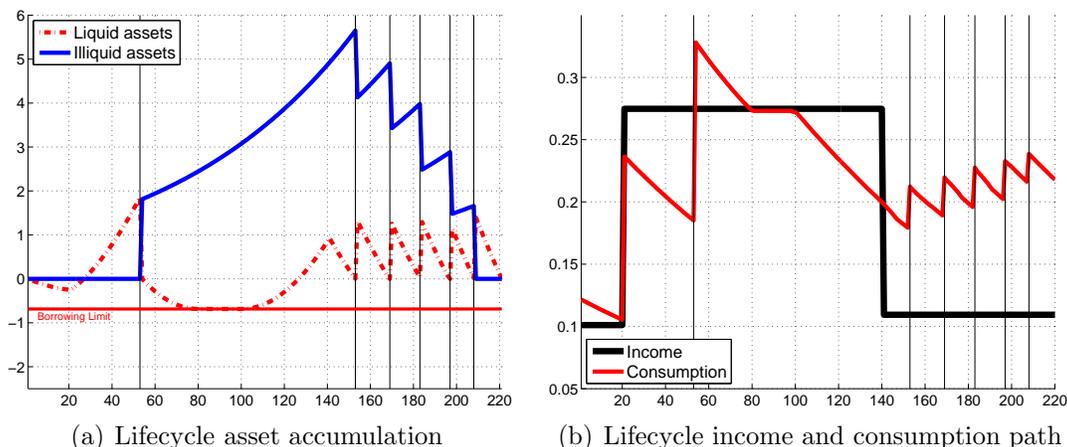


Figure A1: Example of lifecycle of a “wealthy hand-to-mouth” agent in the two-asset model where hand-to-mouth behavior occurs when the agent hits the credit limit.

Figure A1 illustrates how the model can feature households with positive illiquid assets who, at the same time, use credit up to the limit. This is another type of wealthy hand-to-mouth (HtM) behavior, in addition to the one described in the main text (the latter being more prevalent in the data and in the model simulations). In Figure 2, the HTM behavior arises because the agent is at the zero kink for liquid wealth, whereas here it is at the borrowing limit.

After the first deposit into the illiquid account, households would like to increase their consumption to a target level that reflects the higher rate of return earned on their savings. In Figure 2, borrowing costs were prohibitive for the household, and after the deposit the household was immediately constrained. The key difference in the parameterization between the example in this Appendix and the example in the main text is that credit is much cheaper here. As a result, the household starts borrowing

to finance consumption after its deposit (see panel (b) in Figure A1), and it quickly reaches the credit limit. At that point, it stays at the limit for several periods, and consumes all of its earnings, net of the interest payment on debt. During this phase of the lifecycle, upon receiving the rebate check, it will consume a large part of the check and, upon receiving the news of the rebate, it will not increase her expenditures.

As retirement gets closer, the lifecycle saving motive starts kicking in, and it begins repaying its debt and accumulating liquid wealth.

B SCF data and measurement of hand-to-mouth households

B.1 Estimation of cash holdings and credit card debt

Cash imputation The Survey of Consumer Finances (SCF) does not record cash holdings of households. To impute cash holdings to our measure of liquid assets, we make use of the Survey of Consumer Payment Choice, administered by the Federal Reserve Bank of Boston, for 2008 (the earliest survey year). This survey reports that median cash holdings on person and property was \$69 (Foster et al., 2011, Table 9). Median wealth in checking, saving, money market, and call accounts in the SCF 2001 is \$2,858. We therefore increase proportionately all individual household holdings of these assets by a factor of $1 + (69 \times 2) / 2,858 = 1.05$, where the 2 multiplying the median individual holdings of cash accounts for the fact that there are two adults in most households.

Unsecured debt As for the calculation of revolving credit card debt, the SCF asks the following questions about credit card balances: (i) “How often do you pay your credit card balance in full?” Possible answers are: (a) Always or almost always; (b) Sometimes; or (c) Almost never. (ii) “After the last payment, roughly what was the balance still owed on these accounts?” From the first question, we identify households with revolving debt as those who respond (b) Sometimes or (c) Almost Never. We then use the answer to the second question, for these households only, to compute statistics about credit card debt. This strategy (common in the literature, e.g., see Telyukova, 2011) avoids including, as debt, purchases made through credit cards in between regular payments.

B.2 Measurement of hand-to-mouth households

Based on the discussion of Section 4 in the paper, we use the following definitions of hand-to-mouth (HtM) households. Let m_i be the average balance of liquid assets over the past month for household i , and a_i be the stock of illiquid assets, as reported by the SCF. Let y_i be monthly labor income (annual labor income from the SCF divided by 12). Finally, let \underline{m}_i be household's i reported credit limit in the survey.

Household i is HtM if either

$$0 \leq m_i \leq \frac{y_i}{2 \cdot f} \quad (\text{B1})$$

or

$$m_i < 0 \text{ and } m_i \leq \frac{y_i}{2 \cdot f} - \underline{m}_i \quad (\text{B2})$$

where f is the frequency of pay. For monthly frequency $f = 1$, for biweekly $f = 2$, and for weekly $f = 4$. Since the frequency of pay is not available from the SCF, we do all our calculations under three alternative assumptions: weekly, bi-weekly, and monthly frequency.

Household i is *wealthy* HtM if either

$$0 \leq m_i \leq \frac{y_i}{2 \cdot f} \text{ and } a_i > 0. \quad (\text{B3})$$

or

$$m_i < 0 \text{ and } m_i \leq \frac{y_i}{2 \cdot f} - \underline{m}_i \text{ and } a_i > 0 \quad (\text{B4})$$

Poor HtM households are all the residual HtM households who are not wealthy HtM, i.e., those who have $a_i = 0$.

Table B1 row (i) reports the calculation with the baseline definition of liquid and illiquid wealth described in the main text.

We also offer a robustness analysis on these measures. First, we use a stricter definition of liquid wealth that only includes cash, checking, saving, money markets and call accounts (and therefore excludes directly held mutual funds, stocks, bonds and T-Bills which are, arguably, less liquid). Second, we define wealthy HtM only those with illiquid wealth above a positive threshold. As threshold, we choose \$3,000, which is roughly the median amount of liquid wealth held by the U.S. populations. Third, we use a broader definition of illiquid wealth that also includes vehicles (excluded from the baseline definition of illiquid assets). Note that over 80% of households in the SCF own a vehicle and, for many households, this is the major component of their non-liquid wealth. While the first modification increases the total number of HtM agents, the second and third ones only affect the split between “poor” and “wealthy” HtM, but

	Total HtM Households		
	Week	Bi-week	Month
(i) Baseline	0.175	0.225	0.311
(ii) Strict liquid definition	0.188	0.245	0.350
(iii) Strict illiquid definition	0.175	0.225	0.311
(iv) Vehicles	0.175	0.225	0.311

	Wealthy HtM Households		
	Week	Bi-week	Month
(i) Baseline	0.089	0.124	0.190
(ii) Strict liquid definition	0.100	0.142	0.226
(iii) Strict illiquid definition	0.071	0.102	0.165
(iv) Vehicles	0.134	0.179	0.260

	Total HtM in Net Worth		
	Week	Bi-week	Month
(i) Baseline	0.094	0.113	0.139
(ii) Strict liquid definition	0.094	0.113	0.139
(iii) Strict illiquid definition	0.094	0.113	0.139
(iv) Vehicles	0.044	0.050	0.059

Table B1: Estimates of hand-to-mouth (HtM) households. Entries are fraction of the total population. The labels Week, Bi-week and Month refer to the assumptions on the frequency of pay. Row (i) reports the calculation with the baseline definition of liquid and illiquid wealth; (ii) uses a stricter definition of liquid wealth which excludes directed held mutual funds, stocks, and bonds; (iii) defines wealthy HtM only those HtM households with at least \$3,000 in illiquid assets; (iv) adds to illiquid wealth the net value of vehicles.

not the total fraction of HtM households.

As reported in the main body of the paper, our analysis leads us to conclude that between 17.5% and 35% of US households are HtM. This is a conservative estimate (for reasons explained in the main text). Moreover, we estimate that between 40% and 80% of these households are wealthy HtM, depending mainly on the pay frequency and on whether one expands the notion of illiquid wealth to include vehicles.

Finally, for comparison, we also compute the fraction of HtM households in terms of net worth. We apply the definition in (B2) and (B1), with the only difference that in those definitions we use net worth instead of liquid wealth. The bottom part of Table B1 show that the fraction of agents HtM in terms of net worth never exceeds 14%, and is as low as 4-5% when including vehicles as wealth.

C Calibration

C.1 Initial asset positions

We divide households in the SCF into 21 groups based on their earnings and calculate (i) the fraction with zero holdings, and (ii) the median liquid and illiquid wealth in each group, conditional on positive holdings. When we simulate life-cycles in the model, we create the same groups based on the initial earnings draw. Within each group, we initialize a fraction of agents with zero assets, and the rest with the corresponding median holdings of liquid and illiquid wealth. For example, in the median initial earnings group, the fraction of households with zero liquid (illiquid) wealth is 14% (55%). For those with positive holdings, median liquid wealth is \$2,300, and median illiquid wealth is \$7,700.

C.2 Service flow from housing

To calculate the service flow from housing (the parameter ζ in the model), we start from the following relationship holding at any given date t :

$$\zeta_t = r_t^h - m_t^h - n_t^h - (1 - \tau_t^{ded}) (\tau_t^{prop} + i_t^{mort}) \quad (C1)$$

where (as for the left hand side variable) every variable on the right hand side is expressed as a fraction of a unit of housing stock. Specifically, r_t^h is the rental value of a unit of housing, m_t^h are maintenance and repair expenditures, n_t^h are home-owner insurance expenditures, τ_t^{prop} are property taxes, and i_t^{mort} are mortgage interests payments. The formula accounts for the fact that these latter two items are tax deductible at the (average) marginal tax rate τ^{ded} . This formula reflects that owning housing wealth has both costs (maintenance, insurance, property taxes, and mortgage interests) and benefits (imputed rental value of the space and tax deductibility of mortgage interests and property taxes).

We omit from this calculation housing price appreciation net of physical depreciation because this component is included in the calculation of the financial return on total illiquid wealth described in Section C.3. We now explain how we measure all the ingredients in equation (C1). Our final value for ζ is computed as an average of ζ_t over the period 1960-2009, the same period used to compute asset returns in Section C.3.

Our starting point is the total value of residential housing from the Flow of Funds (Table B100). Residential housing can be tenant-occupied or owner-occupied. NIPA

Table 2.5.5 (line 20) reports rents from tenant-occupied housing. For owner-occupied housing, the National Income and Product Accounts (NIPA) use a “rental equivalence approach” stating that the housing services produced by a owner-occupied unit are deemed to be equal in value to the rentals that would be paid on the market for accommodations of the same size, quality, and type. NIPA Table 2.5.5 (line 21) reports these “imputed” rents. Computing total rents over the total value of the residential housing stock over the sample period yields $r^h = 7.9\%$.

We set maintenance and repair expenditures m^h at 1 percent of the stock (an upper bound, see below). The Federal Reserve Board (<http://www.federalreserve.gov/pubs/settlement/default.htm>) estimates the cost of home-owner insurance n^h at 0.35 percent per year. Poterba and Sinai (2008) report an average annual property tax τ^{prop} of 1 percent.

To compute mortgage interest payments i^{mort} as a fraction of the value of the housing stock, we proceed as follows. As a measure of mortgage interest rates, we use the 30-year interest rate on conventional mortgages (series MORTG from the Federal Reserve Bank of St. Louis Federal Reserve Economic Data – “FRED”) which averages 8.3 percent over this period. To calculate the average loan-value ratio, we divide the total outstanding stock of home mortgages from the Flow of Funds (series HMLBSHNO from FRED) by the total value of residential housing from the Flow of Funds (the same series used above), which gives an average value of 0.36 over this period. By multiplying, year by year, the interest rate by the loan-value ratio, we obtain an estimate of mortgage interest payments per unit of housing owned of 2.9 percent.

Finally, Barro and Redlick (2011) report that the average marginal Federal tax rate τ^{ded} over this period was 23.8 percent.

Combining all these components into (C1), and averaging over the sample period, we obtain an estimate of ζ of 4.2 percent per year. This estimate is a lower bound for various reasons.

First, if one repeats the calculation for r^h only on the stock of owner-occupied housing by using the value of residential housing wealth at current cost (i.e., market value) for owner-occupied housing from NIPA Table 5.1 (line 11) together with the imputed rents from owner-occupied housing from NIPA Table 2.5.5, one obtains a higher value for r^h , 8.6% instead of 7.9%, a result which confirms that the conventional wisdom that the stock of owner-occupied housing is, on average, of better quality.

Second, the Census reports estimates of “maintenance and repair” expenditures for both owner-occupied housing and for all residential properties (<http://www.census.gov/>

construction/c50/c50index.html). These estimates are considerably below our baseline of 1 percent per year. Using the Census estimates for m_t^h , we obtain values of ζ which are 0.8-0.9 percentage point higher.

Third, property taxes can be thought of as the price to pay to gain access to certain local services (notably, public schooling). As a result, they are not entirely a cost, as they imply a utility flow as well. Adding back 50 percent of property taxes in the calculation increases ζ by 0.9 percentage points.

Fourth, the service flow originates from the housing stock, whereas in the model a is the net value of illiquid assets. These two values differ because 1) housing is a leveraged claim, and 2) housing is only one asset class (albeit the largest) among illiquid wealth. From the SCF 2001, the median and mean gross housing wealth to net illiquid wealth ratios are, respectively, 1 and 1.6. By applying ζ to a we implicitly use a ratio of one.

To conclude, we choose a value of 1 percent per quarter for ζ and the calculations reported in this section lead us to think that this may be a conservative estimate.

C.3 Returns on liquid and illiquid assets

Risk adjustment Since in the model we abstract from aggregate risk, we perform a “risk-adjustment” on the returns of all our asset classes.

In the data, assets have different returns because of the risk properties of their dividend stream and because of their liquidity value. In our model, the only source of return differentials is liquidity summarized (arguably, in reduced form) by the existence of transaction costs.

We outline two approaches to identify the portion of the return associated with the liquidity properties of the asset in question. The residual approach uses a minimum amount of asset pricing theory to filter out from the observed return the component due to aggregate risk and identifies the one due to liquidity residually. The direct approach uses existing estimates of liquidity premia from the literature.

C.3.1 Residual Approach

The Euler equation for an asset i at date t can be written as

$$1 = E_t [MRS_{t+1} (1 + r_{t+1}^i) (1 - \ell_{t+1}^i)] \quad (C2)$$

where MRS_{t+1} is the marginal rate of substitution of the asset holder, r_{t+1}^i is the return of the asset (price appreciation cum dividend), and $\ell_{t+1}^i \geq 0$ is an additional component of the return that captures the “liquidity value” of asset i (highest for $\ell_{t+1}^i = 0$). For example, Lagos (2010, equation 1) derives the Euler equation (C2) from a model with search frictions where some assets, beyond paying a stream of dividends, are better than others as a medium of exchange for the final consumption good in a decentralized frictional market. There, ℓ_{t+1}^i is a function of the model primitives (e.g., the lower the probability for the holder of asset i to meet a buyer in the frictional market, the higher is ℓ_{t+1}^i).

For an asset which is safe, yields no dividends, and has perfect liquidity, the Euler equation (C2) implies

$$1 = E_t [MRS_{t+1}]. \quad (C3)$$

Abstracting from second order terms, $(1 + r_{t+1}^i) (1 - \ell_{t+1}^i) \simeq 1 + r_{t+1}^i - \ell_{t+1}^i$, rearranging (C2) and using (C3), one can obtain the following reformulation for (the unconditional version of) that Euler equation:

$$E(r^i) = -cov(MRS, r^i) + cov(MRS, \ell^i) + E(\ell) \quad (C4)$$

which yields an intuitive expression for the average return of the asset. The first term in the RHS of (C4) encodes the classical risk premium due to the comovement between the return of the asset and the marginal rate of substitution of the asset holder. The second and third terms capture the additional components of the return associated with the liquidity value. An asset with low liquidity properties ($E(\ell)$ large) and liquidity value that is negatively correlated with the marginal rate of substitution (positive correlation between ℓ and MRS) must command a high financial return to be held by risk-averse households. See Lagos (2010, equation 20), for a reinterpretation of the Euler equation (C2) exactly along these lines.

In this context, risk adjusting the return r^i means eliminating the first covariance component $cov(MRS, r^i)$ from the return in (C4). This covariance-component, however, is model-specific since the MRS depends on preferences and market structure. Our model cannot be used for such calculation since it has no aggregate uncertainty. We therefore propose two empirical strategies to perform this risk adjustment.

First, a plausible assumption, which allows making a risk-adjustment without taking a stand on the MRS , is

$$var(r^i) > -cov(r^i, MRS) \quad (C5)$$

Under this inequality, one can subtract from the expected return the observed variance of the return and obtain a lower bound for the component of the return which is

associated to liquidity, i.e., for the risk-adjusted return.

A second plausible upper bound for the term $-cov(r^a, MRS)$ can be constructed using the insight that, empirically and theoretically, aggregate income volatility exceeds the volatility of the aggregate component of consumption. From NIPA Table 2.1 (series: Compensation of Employees plus $0.66\times$ series Proprietor’s Income) and from the St. Louis FRED database (series: Civilian Employment), we compute labor income per worker and estimate a stochastic process for the residuals of this series around a deterministic linear trend. These residuals are well approximated as an AR1 with autoregressive coefficient of 0.95 and annualized variance of the innovation equal to 0.003. Next, we use our Epstein-Zin-Weil preference specification parameterized as in our calibration (i.e., with risk aversion equal to 4, IES equal to 1.5 and discount factor equal to 0.941) to compute the implied volatility of the MRS, when the consumption process equals the labor income process. See Chen, Favilukis, and Ludvigson (2013, equation 5) for the analytical expression of the MRS with Epstein-Zin-Weil preferences.

Let \overline{MRS} denote this alternative time series proxy for the MRS . We find that $std(\overline{MRS}) = 0.044$. Since this is, arguably, an upper bound for the volatility of the MRS in the data, we can write the inequality

$$-cov(r^a, MRS) < std(r^a) \cdot std(MRS) < std(r^a) \cdot std(\overline{MRS}). \quad (C6)$$

and use the last (measurable from the data) term in this inequality for the risk adjustment. In what follows, we refer to the first strategy based on inequality (C5) as risk-adjustment strategy S1 and to the second strategy based on inequality (C6) as strategy S2.

Nominal returns We apply this methodology to all individual asset classes we consider within the liquid and illiquid wealth groups. All our calculations refer to the period 1960-2006. We perform this calculation in nominal terms first, since we are interested in after-tax returns and taxes apply to nominal returns. Then, we make an adjustment for inflation. We set the annual inflation rate to 4% (the average over this period was 4.1%).

Recall that our definition of liquid assets comprises: cash, money market, checking, savings and call accounts plus directly held mutual funds, stocks, bonds, and T-Bills. Our baseline measure of illiquid assets includes net housing worth, retirement accounts, life insurance policies, CDs, and saving bonds.

We set the nominal return on cash and all non-interest bearing accounts to zero. We set the return on savings accounts, T-Bills, savings bonds, and life insurance (assuming

actuarially fair contracts) to the interest rate on 3-month T-Bills (Federal Reserve Board, FRB hereafter, database). Over the period 1960-2006, we obtain an average nominal return on 3-month T-Bills of 5.33% (SD 2.76%) with an implied risk-adjusted return of 5.25% under strategy S1 and 5.21% under strategy S2.

For CDs (for which data are available only starting from 1964 in the FRB database) we compute a return of 6.29% (SD 3.13%) corresponding to a risk-adjusted return of 6.2% under both strategies.

For equities, we use Center for Research in Security Prices (CRSP) value-weighted returns, assuming dividends are reinvested, and obtain an annualized nominal return of 11.1% (SD 17.89%), with an implied risk-adjusted nominal return of 7.9% under strategy S1 and 10.3% under strategy S2. Note that our risk-adjustment S1 closes half of the gap between equity and bond returns. This is a generous adjustment, in light of the fact that Lagos (2010) concludes that 90% of the equity premium is liquidity driven (and hence the risk adjustment would only account for 10 percent of the gap, similarly to what obtained from our risk-adjustment strategy S2).

The SCF reports the equity share for directly held mutual funds, stocks and bonds and for retirement accounts, which allows us to apply separate returns on the equity and safe components of each saving instrument. An important feature of retirement accounts is the employer's matching rate. Over 70% of households in our sample with positive balance on their retirement account have employer-run retirement plans. The literature on this topic finds that, typically, employers match 50% of employees' contributions up to 6% of earnings, but the vast majority of employees do not contribute above this threshold (e.g., Papke and Poterba, 1995). As a result, we raise the return on retirement accounts by a factor of 1.35.

To compute the rate of return on housing (appreciation net of physical depreciation), we follow two alternative methods. The first method replicates the calculation in Favilukis, Ludvigson, and Van Nieuwerburgh (2010). We measure housing wealth for the household sector from the Flow of Funds (Table B100) and construct an index measuring the growth in residential housing wealth. We then subtract population growth in order to correct for the growth in housing quantity. We obtain an average annual nominal return of 6.6% (SD 7.3%) implying a risk-adjusted nominal return of 6% under the risk-adjustment strategy S1 and 6.2% under strategy S2.

Second, we use the calculations of Piazzesi and Schneider (2007) who list different estimates for the real return on housing over the postwar period. Their Tables B1 and B2 report both means and standard deviation, and hence we can calculate risk-adjusted returns. We find that their estimates range between 1.7 and 2.7 percent per

	Nominal		
	Mean	SD	Risk-Adjusted
Cash, checking accounts	0.00	0.00	0.00
3-month T-bills	5.33	2.76	5.25
Saving acc./bonds, Life ins.	5.33	2.76	5.25
3-month CDs (1964-2009)	6.29	3.13	6.20
Stocks	11.06	17.89	7.86
Housing	6.56	7.30	6.03
	Nominal	Mean	Real After-Tax
	Risk-Adjusted	Tax Rate	Risk-Adjusted
Liquid Wealth	3.30	23.19	-1.48
Illiquid Wealth	6.84	7.86	2.29
Net Worth	6.30	10.37	1.67

Table C1: Summary of calculations for returns of various asset classes (1960-2009). Risk adjustment based on strategy S1.

year in real terms under both risk-adjustment strategies, and hence in line with the 6% nominal obtained from the first approach, given our assumed inflation rate of 4%.

Finally, we note that both risk-adjustment strategies lead to very similar results, except for the case of stocks where the first strategy S1 leads to much lower risk-adjusted returns.

C.3.2 Direct Approach

We take the view that the entire return on saving bonds, 3-month T-Bills and on 3-month CDs is due to their imperfect liquidity (relative, say, to cash or bank accounts), and hence we do not perform any risk-adjustment. The calculations based on the residual approach outlined above suggest the adjustment would be rather trivial anyway.

The most widely cited recent paper on the measurement of liquidity risk for equities is Pastor and Stambaugh (JPE, 2003, PS thereafter). PS study whether liquidity (measured as the temporary effect of order flows on stock prices) is a relevant factor in explaining the cross-section of stock returns, over and above the standard Fama-French factors. Their answer is quite striking: the authors rank stocks by decile of sensitivity to their measure of aggregate liquidity risk and show that liquidity accounts for an excess return of 7.5% between the top and the bottom decile, and roughly 3.5% between the median and the bottom decile over the period 1966-1999.

If we assume that stocks in the bottom decile of the PS classification (the most liquid) are akin to T-Bills in their liquidity properties, and that the median stock is representative of the equity portfolio held by our agents, then we obtain a risk-adjusted nominal return for stocks of $5.33 + 3.5 = 8.83\%$ under this strategy (that we call S3).

Since we are not aware of an equivalent calculation in the literature for housing, we proceed as follows. Over the period 1966-1999, the illiquidity premium computed by PS represents $3.5/6.9=51\%$ of the excess return for stocks. It is reasonable to think therefore that, since housing is less liquid than the median stock, a larger portion of the excess return of housing (1.23%) stems from its illiquid nature. If we assume that this portion is $2/3$, we obtain a risk-adjusted nominal return for housing of $5.33 + 1.23 * 0.66 = 6.14\%$.

Overall, strategy S3 yields a return differential between total illiquid and liquid wealth in between that obtained with strategy S1 and that obtained with strategy S2.

C.3.3 Calculation of real after-tax returns on liquid and illiquid assets

In light of these results, we proceed with our calculations using the first, more conservative, strategy for risk-adjustment, S1. To complete our calculations we need estimates for (i) tax rates and (ii) inflation.

Capital income tax rates Kiefer et al. (2002, Table 5) report the effective tax schedule on interests and dividends, and on long term capital gains by ten income brackets in 2000. We apply the interests and dividend tax rates on all asset returns with two exceptions. First, we apply the capital gain tax rate on the return to retirement accounts. Second, we follow Poterba and Sinai (2008) and set the effective tax rate on housing returns to zero. They write that “since 1997, married (single) homeowners have been able to realize \$500,000 (\$250,000) of capital gains tax-free after a holding period of two years. Relatively few accruing housing capital gains are likely to face taxation under this regime.”

Real after-tax returns We apply these nominal returns (by asset type) and these tax rates (by asset type and household income bracket) to each household portfolio in the SCF and compute average risk-adjusted after-tax nominal returns in the population for liquid wealth, illiquid wealth, and net worth. Finally, we subtract 4% inflation to each rate of return, and obtain risk-adjusted after-tax real returns of -1.48% for liquid wealth, 2.29% for illiquid wealth, and 1.67% for net worth. Table C1 summarizes these calculations.

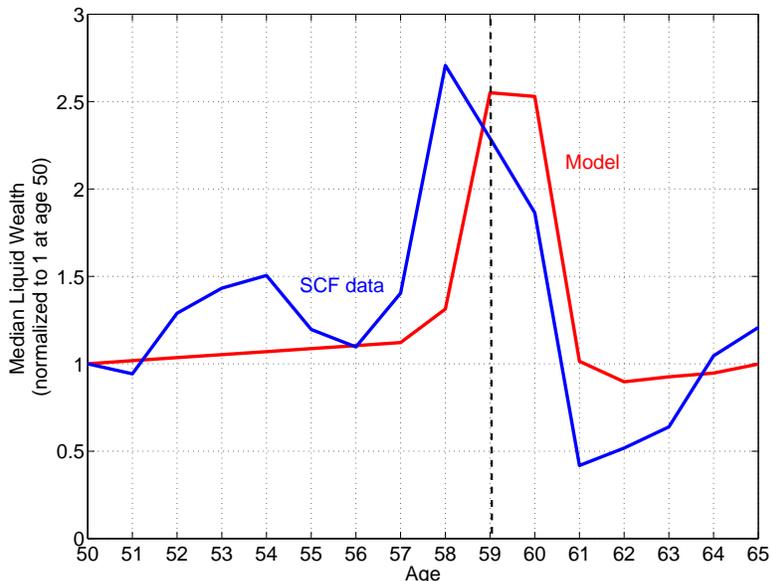


Figure C1: Pattern of median liquid wealth around retirement in the model (where retirement age is 59 for all households) and in the SCF data. SCF data are 3-year moving averages. Model is yearly averages of quarterly values.

C.4 Dynamics of liquid wealth around retirement

Figure C1 zooms on the age range 50-65 to display the hump in median liquid wealth around retirement in the model and in the SCF data. In the model, households accumulate liquid wealth in anticipation of retirement to smooth the drop in income. The micro data do display a similar pattern. Unsurprisingly, in the data the hump is smoother since not every individual retires at the same age.

D Robustness

Table D1 summarizes our sensitivity analysis with respect to preference parameters (risk aversion and IES), access to credit (borrowing costs and limits), desirability of the illiquid asset (financial return and consumption flow), and size of the idiosyncratic risk.⁴⁶ The analysis is done for all three information structures, and for both the one-asset and the two-asset models. For every parameterization, we recalibrate the discount factor β to match median illiquid wealth, or net worth, (as a fraction of average income).

⁴⁶The table does not report sensitivity with respect to the transaction cost κ because it can be easily inferred from the figures in the paper.

Preferences Increasing the coefficient of relative risk aversion from 2 to 6 raises the rebate coefficient because households hold more illiquid wealth as a precautionary saving instrument in case they are hit by large shocks. As a result, the calibrated discount factor needed to match the median illiquid wealth-income ratio is lower. Higher impatience increases the MPC of all agents.

As we mention in the main text, the IES plays a powerful role. Households who are more willing to substitute consumption intertemporally are more likely to save heavily in the illiquid asset, and to be wealthy hand-to-mouth, during working-age to enjoy higher consumption at retirement. Moreover, those households who learn about the rebate in advance are less likely to use costly credit to start spending the check earlier, and would rather wait one extra quarter to consume it. Indeed, with higher IES there are more hand-to-mouth agents and fewer agents using credit in the economy. Both forces push up the rebate coefficient.

Credit Lowering and increasing the borrowing cost, relative to the baseline, increases the rebate coefficient. Cheap credit creates an arbitrage opportunity: many households borrow up to the limit to invest into the illiquid asset, and end up wealthy hand-to-mouth at the credit limit (recall the example in Appendix A). When credit is very expensive, few households ever borrow and there are many more hand-to-mouth households at the zero kink for liquid wealth.

Table D1 shows that our credit limit is not too binding. Doubling the limit has no impact on the rebate coefficient. Tightening the limit down to zero has similar effects to prohibitively increasing borrowing costs.

Desirability of the illiquid asset Raising the return wedge and the housing-service flow makes the illiquid asset more desirable and induces more households to be wealthy hand-to-mouth which, in turn, increases the rebate coefficient.

Idiosyncratic earnings risk Making the individual earnings process more volatile has similar effects to raising risk aversion. It pushes households in the model to hold more illiquid wealth as a precautionary saving instrument. The discount factor required to replicate the median illiquid wealth-income ratio in the data is lower, and this lower degree of patience increases the MPC of all agents.

Information structure		Rebate coefficient					
		Baseline		Surprise for All		Anticipated by All	
Assets in model		One	Two	One	Two	One	Two
Borrowing rate	5%	<1%	19%	<1%	19%	<1%	19%
	10%	<1%	15%	3%	25%	<1%	11%
	15%	3%	25%	3%	30%	<1%	17%
Credit limit	0	<1%	28%	3%	30%	<1%	27%
	0.74	<1%	15%	3%	25%	<1%	11%
	1.48	<1%	14%	3%	24%	<1%	11%
Risk aversion	2	<1%	13%	3%	23%	<1%	10%
	4	<1%	15%	3%	25%	<1%	11%
	6	<1%	17%	4%	28%	3%	23%
IES	1.05	<1%	9%	3%	17%	<1%	7%
	1.5	<1%	15%	3%	25%	<1%	11%
	2	<1%	20%	3%	33%	<1%	16%
Return wedge	2.54	<1%	15%	3%	25%	<1%	11%
	3.54	<1%	15%	3%	25%	<1%	11%
	4.54	<1%	14%	3%	25%	<1%	11%
Housing service flow	0.02	<1%	14%	3%	22%	<1%	11%
	0.04	<1%	15%	3%	25%	<1%	11%
	0.06	<1%	18%	3%	25%	<1%	15%
Variance of shocks	0.002	<1%	14%	3%	23%	<1%	11%
	0.003	<1%	15%	3%	25%	<1%	11%
	0.004	<1%	16%	4%	27%	<1%	11%

Table D1: Robustness analysis. The borrowing rate is the nominal annual rate on unsecured credit. The credit limit is expressed as a fraction of quarterly income, as in the model. The return wedge is the differential after-tax return between illiquid and liquid assets. In all sensitivity analyses, the middle row is the value of the baseline calibration. For every parameterization, we recalibrate the discount factor β to match median illiquid wealth (as a fraction of average income).

E Numerical solution of the model

E.1 Detailed Description of Model

E.1.1 Preliminaries

An agent of age j can hold two assets in the model: an illiquid asset, a_j , that has an associated price q^a ; and a liquid asset, m_j , that has an associated price $q^m(m_{j+1})$, where dependence on m_{j+1} reflects the possibility of a wedge between the borrowing cost and the interest rate on liquid saving.

In this appendix we make the following modifications relative to the main text:

1. For ease of notation, we let $\psi_j \equiv (\alpha, z_j)$, and write earnings at age j as $y_j(\psi_j)$. We denote by $F(\psi_j|\psi_{j-1})$ the conditional probability distribution of earnings and assume ψ_j can only take a finite number of values.
2. In the main text we defined a tax function $\mathcal{T}(y_j, a_j, m_j)$. Since this tax function is separable between earnings and the two assets, in this appendix we express its earnings component as $T(y_j)$ to reflect the (non-linear) tax on earnings, and interpret the prices (q^a, q^m) as after-tax prices.
3. We use e_j to denote total expenditures before tax. That is $e_j \equiv c_j + h_j$ where c_j is non-durable expenditures and h_j is housing expenditures on the rental market. Because of the assumption of a frictionless rental market for housing, the model can be solved in two stages. In the first stage we solve for total expenditures, allowing for a flow of consumption services from the illiquid asset holdings in period j in the amount of ζa_{j+1} . In the second stage we solve the within-period problem of allocating total spending on non-durables and rental housing services, conditional on the optimal total expenditure and holdings of illiquid assets. In Section E.2 below, we show the solution to this second stage problem yields the indirect period utility function $e_{j+1} + \zeta a_{j+1}$, which we use in the first stage.

We define the following objects:

- x_j^N is total liquid funds available for consuming and saving, for an agent who is not adjusting:

$$x_j^N(m_j, a_j, y_j) \equiv m_j + y_j - T(y_j) + reb_j$$

reb_j is equal to 0 unless a rebate is received in period j .

- x_j^A is total liquid funds available for consuming and saving, for an agent who is adjusting, before paying the adjustment cost:

$$\begin{aligned} x_j^A(m_j, a_j, y_j) &\equiv m_j + a_j + y_j - T(y_j) + reb_j \\ &= x_j^N(m_j, a_j, y_j) + a_j \end{aligned}$$

- $V_j^A(x_j, \psi_j)$ is the value function if the agent accesses the illiquid asset. $e_j^A(x_j, \psi_j)$ is the associated consumption policy function.
- $V_j^N(x_j, a_j, \psi_j)$ is the value function if the agent does not access the illiquid asset. $e_j^N(x_j, a_j, \psi_j)$ is the associated consumption policy function.
- We define the expected value function, where the expectation is taken over the current period shocks, and so is a function of the current period holdings of the two types of assets (since these are chosen the period before) and the previous period's realization of the persistent component of earnings. Note that cash-on-hand is only realized when earnings are realized and so is not a state variable for the expected value function. Dependence of (x_j^A, x_j^N) on (m_j, a_j, y_j) is implicit in this function and those defined below.

$$EV_j(m_j, a_j, \psi_{j-1}) = \sum_{\psi_j \in \Psi_j} \max \{V_j^A(x_j^A, \psi_j), V_j^N(x_j^N, a_j, \psi_j)\} F(\psi_j | \psi_{j-1})$$

- We define a new operator, $\tilde{\max}$. This operator chooses between two objects based on which of the corresponding value functions is higher. For example $\tilde{\max} \{e^A, e^N\}$ selects consumption expenditures e^A when $V^A > V^N$ at the corresponding point in the state space.
- We define the risk-adjusted expected value function, RV_j , as

$$RV_j(m_j, a_j, \psi_{j-1})^{1-\gamma} = \sum_{\psi_j \in \Psi_j} \tilde{\max} \left\{ V_j^A(x_j^A, \psi_j)^{1-\gamma}, V_j^N(x_j^N, a_j, \psi_j)^{1-\gamma} \right\} F(\psi_j | \psi_{j-1})$$

- We define the functions $FV_{a,j}$ and $FV_{m,j}$ as

$$\begin{aligned} FV_{a,j}(m_j, a_j, \psi_{j-1}) &= \sum_{\psi_j \in \Psi_j} \tilde{\max} \left\{ V_j^A(x_j^A, \psi_j)^{-\gamma} \frac{\partial V_j^A}{\partial a_j}, V_j^N(x_j^N, a_j, \psi_j)^{1-\gamma} \frac{\partial V_j^N}{\partial a_j} \right\} F(\psi_j | \psi_{j-1}) \\ FV_{m,j}(m_j, a_j, \psi_{j-1}) &= \sum_{\psi_j \in \Psi_j} \tilde{\max} \left\{ V_j^A(x_j^A, \psi_j)^{-\gamma} \frac{\partial V_j^A}{\partial m_j}, V_j^N(x_j^N, a_j, \psi_j)^{1-\gamma} \frac{\partial V_j^N}{\partial m_j} \right\} F(\psi_j | \psi_{j-1}) \end{aligned}$$

- We define $S_j = (m_j, a_j, \psi_{j-1})$.

E.1.2 Decision Problems

Problem if not adjusting

$$V_j^N(x_j, a_j, \psi_j) = \max_{e_j, m_{j+1}} \left\{ (1 - \beta) (e_j + \zeta a_{j+1})^{1-\sigma} + \beta RV_{j+1}(S_{j+1})^{1-\sigma} \right\}^{\frac{1}{1-\sigma}}$$

subject to:

$$\begin{aligned} q^m (m_{j+1}) m_{j+1} + (1 + \tau^c) e_j &\leq x_j \\ q^a a_{j+1} &= a_j \\ m_{j+1} &\geq \underline{m}_{j+1}(y_j) \end{aligned}$$

Problem if adjusting

$$V_j^A(x_j, \psi_j) = \max_{e_j, m_{j+1}, a_{j+1}} \left\{ (1 - \beta) (e_j + \zeta a_{j+1})^{1-\sigma} + \beta RV_{j+1}(S_{j+1})^{1-\sigma} \right\}^{\frac{1}{1-\sigma}}$$

subject to:

$$\begin{aligned} q^m (m_{j+1}) m_{j+1} + q^a a_{j+1} + (1 + \tau^c) e_j &\leq x_j - \kappa \\ m_{j+1} &\geq \underline{m}_{j+1}(y_j) \\ a_{j+1} &\geq 0 \end{aligned}$$

E.1.3 First-Order Necessary Conditions

To solve the model, we derive the first-order conditions. Note that due to the non-convexity of the problem, these are not sufficient. Nonetheless, these conditions are necessary. Our computational approach is to look for all solutions to each set of FOCs, and then compare the associated value functions at each candidate solution.

No-adjust case When agents do not adjust, there is one FOC, a standard Euler Equation (EE):

$$\frac{1 - \beta}{1 + \tau^c} (e_j + \zeta a_{j+1})^{-\sigma} = \begin{cases} \frac{\beta}{q^m} RV_{j+1}(S_{j+1})^{\gamma-\sigma} FV_{m,j+1}(S_{j+1}) & \text{if } m_{j+1} > 0 \\ \frac{\beta}{q^m} RV_{j+1}(S_{j+1})^{\gamma-\sigma} FV_{m,j+1}(S_{j+1}) & \text{if } m_{j+1} < 0 \\ \in \left[\frac{1}{q^m}, \frac{1}{\bar{q}^m} \right] \cdot \beta RV_{j+1}(S_{j+1})^{\gamma-\sigma} FV_{m,j+1}(S_{j+1}) & \text{if } m_{j+1} = 0 \end{cases}$$

Adjust case For adjusting agents there are two FOCs. One is a standard Euler Equation (intuitively, the liquid asset can be adjusted costlessly the following period

so an EE holds), the other is a portfolio problem that equates the marginal value of investing in the two different assets:

$$\frac{1-\beta}{1+\tau^c} (e_j + \zeta a_{j+1})^{-\sigma} = \begin{cases} \frac{\beta}{\bar{q}_j^m} RV_{j+1} (S_{j+1})^{\gamma-\sigma} FV_{m,j+1} (S_{j+1}) & \text{if } m_{j+1} > 0 \\ \frac{\beta}{\bar{q}^m} RV_{j+1} (S_{j+1})^{\gamma-\sigma} FV_{m,j+1} (S_{j+1}) & \text{if } m_{j+1} < 0 \\ \in \left[\frac{1}{\bar{q}^m}, \frac{1}{\bar{q}^m} \right] \cdot \beta RV_{j+1} (S_{j+1})^{\gamma-\sigma} FV_{m,j+1} (S_{j+1}) & \text{if } m_{j+1} = 0 \end{cases}$$

$$\begin{aligned} \frac{1-\beta}{1+\tau^c} (e_j + \zeta a_{j+1})^{-\sigma} &= \frac{1-\beta}{q^a} \zeta (e_j + \zeta a_{j+1})^{-\sigma} \\ &\quad + \frac{\beta}{q^a} RV_{j+1} (S_{j+1})^{\gamma-\sigma} FV_{a,j+1} (S_{j+1}) \end{aligned}$$

with an inequality for the second FOC when the non-negativity constraint on illiquid assets ($a_{j+1} \geq 0$) binds.

Below we transform these two equations into an Euler equation and a portfolio constraint, so that they can be solved by (i) guessing the solution to the inter-temporal saving problem, and then (ii) solving the portfolio problem at each guessed value for savings.

E.1.4 Envelope Conditions

Here we derive the partial derivatives of value function that are required to evaluate $FV_{a,j}$ and $FV_{m,j}$. Our approach is to store these partial derivatives alongside the value function and policy functions, constructing them recursively. Of course, they may not be continuous, due to the discrete choice. However, (i) if there is enough uncertainty in the problem the jumps tend to be smoothed away; and (ii) there are a finite number points of discontinuity.

Recall that

$$\begin{aligned} FV_{m,j}(S_j) &= E \left[\text{m}\ddot{\text{a}}\text{x} \left\{ V_j^A(x_j^A, \psi_j)^{-\gamma} \frac{\partial V_j^A}{\partial m_j}, V_j^N(x_j^N, a_j, \psi_j)^{-\gamma} \frac{\partial V_j^N}{\partial m_j} \right\} \right] \\ FV_{a,j}(S_j) &= E \left[\text{m}\ddot{\text{a}}\text{x} \left\{ V_j^A(x_j^A, \psi_j)^{-\gamma} \frac{\partial V_j^A}{\partial a_j}, V_j^N(x_j^N, a_j, \psi_j)^{-\gamma} \frac{\partial V_j^N}{\partial a_j} \right\} \right] \end{aligned}$$

where the partial derivatives with respect to assets and cash on hand are related by

$$\begin{aligned}\frac{\partial V_j^A}{\partial m_j} &= \frac{\partial V_j^A(x_j^A)}{\partial x_j} \equiv V_{x,j}^A \\ \frac{\partial V_j^N}{\partial m_j} &= \frac{\partial V_j^N(x_j^N)}{\partial x_j} \equiv V_{x,j}^N \\ \frac{\partial V_j^A}{\partial a_j} &= \frac{\partial V_j^A(x_j^A)}{\partial x_j} = V_{x,j}^A\end{aligned}$$

We denote the partial derivative with respect to illiquid assets when not adjusting by

$$\frac{\partial V_j^N}{\partial a_j} \equiv V_{a,j}^N$$

Next, we compute these partial derivatives of the choice-specific value functions. For the adjust case, it is given by

$$V_{x,j}^A(x_j, \psi_j) = \frac{1 - \beta}{1 + \tau^c} (e_j + \zeta a_{j+1})^{-\sigma} (V_j^A)^\sigma$$

For the no-adjust case, they are given by

$$\begin{aligned}V_{a,j}^N(x_j, a_j, \psi_j) &= \frac{\zeta}{q^a} \frac{1 - \beta}{1 + \tau^c} \left(e_j + \zeta \frac{a_j}{q^a} \right)^{-\sigma} (V_j^N)^\sigma + \frac{\beta}{q^a} RV_{j+1}(S_{j+1})^{\gamma-\sigma} FV_{a,j+1}(S_{j+1}) (V_j^N)^\sigma \\ V_{x,j}^N(x_j, a_j, \psi_j) &= \frac{1 - \beta}{1 + \tau^c} (e_j + \zeta a_{j+1})^{-\sigma} (V_j^N)^\sigma\end{aligned}$$

In these expressions, e_j , m_{j+1} and a_{j+1} on the RHS should be interpreted as the optimal decision rules at the point (x_j, a_j, ψ_j) .

E.1.5 Recursive Computation

To make progress in constructing these objects recursively, it is useful to define some intermediate functions:

$$\begin{aligned}d_j(S_j) &\equiv \frac{FV_{a,j}(S_j) RV_j(S_j)^{\gamma-\sigma}}{1 - \beta} \\ g_j(x_j, a_j, \psi_j) &\equiv \frac{V_{a,j}^N(x_j, a_j, \psi_j)}{(1 - \beta) (V_j^N)^\sigma} \\ \mu_j(S_j) &\equiv \frac{FV_{m,j}(S_j) RV_j(S_j)^{\gamma-\sigma}}{1 - \beta}\end{aligned}$$

By substituting into the envelope conditions, we obtain the following recursions:

$$\begin{aligned}\mu_j(S_j) &= RV_j(S_j)^{\gamma-\sigma} E \left[\text{m\ddot{a}x} \left\{ (V_j^A)^{\sigma-\gamma} (e_j^A + \zeta a_{j+1}^A)^{-\sigma}, (V_j^N)^{\sigma-\gamma} \left(e_j^N + \zeta \frac{a_j^N}{q^a} \right)^{-\sigma} \right\} \right] \\ g_j(x_j, a_j, \psi_j) &= \frac{\zeta}{q^a} \left(e_j^N + \frac{\zeta}{q^a} a_j \right)^{-\sigma} + \frac{\beta}{q^a} d_{j+1}(S_{j+1}) \\ d_j(S_j) &= RV_j(S_j)^{\gamma-\sigma} E \left[\text{m\ddot{a}x} \left\{ (V_j^A)^{\sigma-\gamma} (e_j^A + \zeta a_{j+1}^A)^{-\sigma}, (V_j^N)^{\sigma-\gamma} g_j(x_j, a_j, \psi_j) \right\} \right]\end{aligned}$$

These recursions reflect the expected marginal values of illiquid assets (d_j) and total assets (μ_j).

E.1.6 Euler Equations

We can now finally substitute these into the first-order conditions and obtain the Euler Equations that need to be solved.

For the no-adjust case, we have one Euler equation:

$$\left(e_j + \zeta \frac{a_j}{q^a} \right)^{-\sigma} = \begin{cases} \frac{\beta}{q^m} \mu_{j+1}(S_{j+1}) & \text{if } m_{j+1} > 0 \\ \frac{\beta}{q^m} \mu_{j+1}(S_{j+1}) & \text{if } m_{j+1} < 0 \\ \in \left[\frac{1}{q^m}, \frac{1}{q^m} \right] \times \beta \mu_{j+1}(S_{j+1}) & \text{if } m_{j+1} = 0 \end{cases}$$

For the adjusting agents, there are two Euler equations:

$$(e_j + \zeta a_{j+1})^{-\sigma} = \begin{cases} \frac{\beta}{q^m} \mu_{j+1}(S_{j+1}) & \text{if } m_{j+1} > 0 \\ \frac{\beta}{q^m} \mu_{j+1}(S_{j+1}) & \text{if } m_{j+1} < 0 \\ \in \left[\frac{1}{q^m}, \frac{1}{q^m} \right] \times \beta \mu_{j+1}(S_{j+1}) & \text{if } m_{j+1} = 0 \end{cases}$$

$$\begin{aligned}(e_j + \zeta a_{j+1})^{-\sigma} &= \frac{\beta}{q^a} d_{j+1}(S_{j+1}) + (1 + \tau^c) \frac{\zeta}{q^a} (e_j + \zeta a_{j+1})^{-\sigma} & \text{if } a_{j+1} > 0 \\ (e_j + \zeta a_{j+1})^{-\sigma} &> \frac{\beta}{q^a} d_{j+1}(S_{j+1}) + (1 + \tau^c) \frac{\zeta}{q^a} (e_j + \zeta a_{j+1})^{-\sigma} & \text{if } a_{j+1} = 0\end{aligned}$$

E.1.7 Recursive algorithm

The model is computed by recursively solving these Euler Equations backward from the last period of life $j = J$. At each point in the state space, we search for multiple solutions to the first-order conditions, compute the associated value functions and choose the solution with the highest value. We explicitly allow for the possibility of

solutions at each of the corners and compute the associated value function at these points.

E.2 Sub-problem for housing and non-durable consumption

In this section we outline the static sub-problem at age j that yields the optimal choice of housing services h_j bought/sold on the rental market, and non-durable consumption c_j . In this problem, total expenditures e_j and the allocation of illiquid assets a_{j+1} are predetermined. Recall that total housing services s_j which yields utility to the agent also include the flow from the illiquid asset. The household faces the problem:

$$u(e_j, a_{j+1}) = \max_{c_j, s_j, h_j} c_j^\phi s_j^{1-\phi}$$

subject to:

$$\begin{aligned} c_j + h_j &= e_j \\ s_j &= h_j + \zeta a_{j+1} \\ h_j &\geq -\zeta a_{j+1} \\ c_j &\geq 0 \end{aligned}$$

The interior solution to this problem is:

$$\begin{aligned} c_j &= \phi(e_j + \zeta a_{j+1}) \\ s_j &= (1 - \phi)(e_j + \zeta a_{j+1}) \\ h_j &= (1 - \phi)e_j - \phi\zeta a_{j+1} \end{aligned}$$

The resulting indirect utility function (modulo a multiplicative constant) used in the first-stage problem is:

$$u(e_j, a_{j+1}) = e_j + \zeta a_{j+1}.$$

E.3 Bounds, grids, and interpolation

We now describe the space for each of the state variables for the problem and our methods for interpolation.

E.3.1 (m_j, a_j) space

The risk-adjusted expected value function RV_j and the expected marginal values of the two assets (μ_j, d_j) are defined over the space (m_j, a_j) . We discretize this space as

follows. Let the lower bound for liquid assets, m_j , be given by \underline{m}_j . Let M_j and A_j be an exogenous, age-dependent upper bound on liquid and illiquid assets, that will be chosen so that they never bind in the solution. Then the feasible set for (m_j, a_j) is

$$\begin{aligned} m_j &\in [\underline{m}_j, M_j] \\ a_j &\in [0, A_j] \end{aligned}$$

i.e., a rectangular space. We choose grid points in the a dimension to be polynomial spaced with more points closer to $a = 0$. We choose grids in the positive m dimension to be polynomial spaced between $m = 0$ and $m = M$, with an explicit point at $m = 0$. For the negative m dimension, the grid points are polynomial spaced between \underline{m} and $\underline{m}/2$, and between $\underline{m}/2$ and 0, with more points closer to 0 and \underline{m} .

E.3.2 (x_j, a_j) space

The value functions (V_j^A, V_j^N) and the decision rules are defined separately for the adjust and no-adjust cases.

When the agent is adjusting, these are defined over the space of cash on hand conditional on adjusting, x_j^A . This space is discretized as follows. The lowest possible value of x_j^A is

$$\underline{x}_j^A = \underline{m}_j + \min \{y_j - T(y_j)\}$$

and the highest possible value is

$$X_j^A = M_j + \max \{y_j - T(y_j)\}$$

We choose grids in the positive dimension to be polynomial spaced between 0 and X_j^A , with an explicit point at $x_j^A = 0$. For the negative x_j^A dimension, the grid points are polynomial spaced between \underline{x}_j^A and $\underline{x}_j^A/2$, and between $\underline{x}_j^A/2$ and 0, with more points closer to 0 and \underline{x}_j^A .

When the agent is not adjusting, these functions are defined over the space (x_j^N, a_j) . We use the same space as defined above for a_j . The x_j^N space is discretized as follows. The lowest and highest possible values of x_j^N are

$$\begin{aligned} \underline{x}_j^N &= \underline{m}_j + \min \{y_j - T(y_j)\} \\ X_j^N &= M_j + \max \{y_j - T(y_j)\} \end{aligned}$$

subject to these not violating the borrowing limit. The grid points are chosen in analogous manner to the adjust case.

E.3.3 Grid sizes

In the models without borrowing, we use 30 points each in the grids for a_j , m_j and x_j^N , and 50 points in the grid for x_j^A . In the models with borrowing, we retain the same grid points as for the models without borrowing, but add 16 points in the negative regions for each of m_j , x_j^N and x_j^A . We use 21 points in the grid for the realization of the permanent shock. Polynomial spaced grids with points concentrated at the lower bound are constructed by taking an equally spaced partition, z , of $[0, 1]$, then constructing a grid for x as $x_L + (x_H - x_L) z^{\frac{1}{k}}$. We use $k = 0.4$.

E.3.4 Interpolation

We use linear and bilinear interpolation. When using bilinear interpolation over the (m_{j+1}, a_{j+1}) space, we interpolate along the m_{j+1} dimension and a diagonal that holds total assets, $m_{j+1} + a_{j+1}$ constant. This provides much more accurate interpolations than standard bilinear interpolation since m_{j+1} is the relevant dimension if the agent does not adjust at $j + 1$, while $m_{j+1} + a_{j+1}$ is the relevant dimension if the agent does adjust at $j + 1$.

E.4 Computation of rebate coefficients

To compute the rebate coefficients implied by the model, we simulate two consumption paths for each of 200,000 individuals. Thus, the size of the simulated economy in the policy experiments is 400,000, two identical groups of size 200,000 each. In the first path, the timing of the arrival of the information and payment of the rebate check is as described for group A in the text. In the second path, the timing of the arrival of the information and payment of the rebate check is as described for group B in the text. These paths depend on the assumed information structure.

We compute the average rebate coefficient by regressing consumption growth of all individuals (combining both paths) on the amount of rebate received in that period (either \$500 or zero), a full set of quarter dummies, and a quadratic polynomial in age. We use only the quarters in which some individuals receive a check. This approach is equivalent to regression (1) in the main text. To mitigate the effects of outliers, we estimate this regression on a truncated sample of individuals whose individual-specific rebate coefficients are within 2 standard deviations either side of the mean, a procedure that results in dropping approximately the top and bottom 1% of individual-specific rebate coefficients. We compute the individual-specific rebate coefficients as the

individual’s average consumption growth in the periods when they receive the check, minus their average consumption growth in the periods when they do not receive the check, using only the periods where they receive the check in one of the paths. So, for example in the baseline informational configuration, we use the average between consumption growth when the individual is in group A at Q2 and consumption growth when it is in group B at Q3, minus the average between consumption growth when it is in group A at Q3 and consumption growth when it is in group B at Q2.

To compute the aggregate consumption response to the policy, we simulate a third counterfactual consumption path for each of the 200,000 individuals in which they never receive a stimulus payment. We compute the aggregate consumption response as the average of the aggregate consumption for groups A and B minus the aggregate consumption along the counterfactual path.

E.5 Other computational details

Our model is very computationally intensive. However, by working with the first order conditions directly, rather than using value function iteration, and by parallelizing the computation of decision rules and simulations, we are able to compute the model in a reasonable time on New York University’s High Performance Computing Bowery cluster. Using 16 processors, it takes roughly 1-2 hours to solve one parameterization of the model. This involves iterating over the steady state of the model (to calibrate the discount factor, which is computationally equivalent to solving for the interest rate in a general equilibrium economy), iterating over the transition path induced by the policy change (to find the payroll tax that balances the government budget constraint), simulating the economy, and computing rebate coefficients.

The amount of memory (RAM) that is required to store the large number of decision rules -for each quarter along the transition at every quarter of the lifecycle over a very large state space- and the large number of simulations is significant. Our baseline model requires around 50GB of RAM to run.

References

- BARRO, R. J., AND C. J. REDLICK (2011): “Macroeconomic Effects from Government Purchases and Taxes,” *Quarterly Journal of Economics*, 126(1), 51–102.
- FAVILUKIS, J., S. LUDVIGSON, AND S. VAN NIEUWERBURGH (2010): “The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium,” *Mimeo*.
- FOSTER, K., E. MEIJER, S. SCHUH, AND M. ZABEK (2011): “The 2009 Survey of Consumer Payment Choice,” *Federal Reserve Bank of Boston Public Policy Discussion Paper*, 11-1.
- KIEFER, D., R. CARROLL, J. HOLTZBLATT, A. LERMAN, J. MCCUBBIN, D. RICHARDSON, AND J. TEMPALSKI (2002): “The Economic Growth and Tax Relief Reconciliation Act of 2001: Overview and Assessment of Effects on Taxpayers,” *National Tax Journal*, 55(1), 89–118.
- LAGOS, R. (2010): “Asset Prices and Liquidity in an Exchange Economy,” *Journal of Monetary Economics*, 57(8), 913 – 930.
- PAPKE, L., AND J. POTERBA (1995): “Survey Evidence on Employer Match Rates and Employee Saving Behavior in 401 (k) Plans,” *Economics Letters*, 49(3), 313–317.
- PASTOR, L., AND R. STAMBAUGH (2003): “Liquidity Risk and Expected Stock Returns,” *Journal of Political Economy*, 111(3), 642–685.
- PIAZZESI, M., AND M. SCHNEIDER (2007): “Asset Prices and Asset Quantities,” *Journal of the European Economic Association*, 5(2-3), 380–389.
- POTERBA, J., AND T. SINAI (2008): “Income Tax Provisions Affecting Owner-Occupied Housing: Revenue Costs and Incentive Effects,” *NBER Working Paper*, 14253.
- TELYUKOVA, I. (2013): “Household Need for Liquidity and the Credit Card Debt Puzzle,” *Review of Economic Studies*, 80(3), 1148–1177.