

A Model of the Consumption Response to Fiscal Stimulus Payments*

Greg Kaplan

University of Pennsylvania and IFS
gkaplan@sas.upenn.edu

Giovanni L. Violante

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

THIS DRAFT: JANUARY 2012

Abstract

A wide body of empirical evidence, based on randomized experiments, finds that 20-40 percent of fiscal stimulus payments (e.g. tax rebates) are spent on nondurable household consumption in the quarter that they are received. We develop a structural economic model to interpret this evidence. Our model integrates the classical Baumol-Tobin model of money demand into the workhorse incomplete-markets life-cycle economy. In this framework, households can hold two assets: a low-return liquid asset (e.g., cash, checking account) and a high-return illiquid asset (e.g., housing, retirement account) that carries a transaction cost. The optimal life-cycle pattern of wealth accumulation implies that many households are “wealthy hand-to-mouth”: they hold little or no liquid wealth despite owning sizeable quantities of illiquid assets. They therefore 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 is able to generate consumption responses to fiscal stimulus payments that are in line with the data.

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 for his insights at the early stage of this research, and to numerous seminar participants for comments. This research is supported by grant no. 1127632 from the National Science Foundation.

1 Introduction

Fiscal stimulus payments (i.e., government transfers to households such as tax rebates) are frequently used by governments to alleviate the impact of recessions on households' welfare. In the last decade, this type of fiscal intervention was authorized by U.S. Congress in the downturns of 2001 and 2008.¹ Households received one-off payments that ranged from \$500 to \$1,000, depending on the specific episode. In the aggregate, these fiscal outlays were remarkably large: \$38B in 2001, \$79B in 2008, and \$60B in 2009, roughly equivalent to 0.5% 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. Using data from the Consumer Expenditure Survey (CEX), Johnson, Parker, and Souleles (2006, hereafter JPS), Parker, Souleles, Johnson, and McLelland (2011, hereafter PSJM), and Misra and Surico (2011) cleverly exploit the randomized timing of the receipt of payments to estimate the effects of the fiscal stimulus payments on consumption expenditures. Shapiro and Slemrod (2003a, 2003b, 2009, hereafter SS) substantiate these studies with a detailed qualitative survey on how consumers use their rebate.

This collective body of evidence convincingly concludes that between 20 and 40 percent of rebates are spent by households on nondurables in the quarter that they are received. This strong consumption response is measured relative to the (comparable, because of the randomization) group of households who are arguably aware of the rebate but have not yet received their check. Two crucial facts are encoded in this finding: 1) the marginal propensity to consume (MPC) out of the extra cash is large; 2) the MPC out of the news of the extra cash is small.

In spite of this large body of empirical research, there are virtually no quantitative studies of these episodes within structural, dynamic, forward-looking models. Such a gap in the literature is troubling because thoroughly understanding the effectiveness of tax rebates as a short-term stimulus for aggregate consumption is paramount for macroeconomists and policy makers.² Knowing the determinants of how consumers

¹In the context of the latest downturn, Oh and Reis (2011) document that, contrary to common belief, the large fiscal expansion of 2007-09 consisted primarily of growing social assistance, as opposed to government purchases. Half of this expansion comprised of discretionary fiscal stimulus transfers.

²The JPS (2006) estimates feature prominently in the reports prepared by the Congressional Budget Office (CBO, 2009) and the Council of Economic Advisors (CEA, 2010) documenting and forecasting the macroeconomic effects of fiscal stimulus.

respond to stimulus payments helps in choosing policy options and in understanding whether the same fiscal instrument can be expected to be more or less effective under different macroeconomic conditions.³

To develop a structural model that has some hope of matching the micro evidence is a challenging task: “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 MPC out of anticipated transitory income fluctuations, such as tax rebates, should be negligible in the aggregate. In that model, the only agents whose consumption would react significantly to the receipt of 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 below 10%) to generate a big enough response in the aggregate.

In this paper we overcome this challenge by proposing a quantitative framework that can speak 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 the workhorse incomplete-markets life-cycle economy. In our model, households can hold two assets: a liquid asset (e.g., cash, or bank account) and an illiquid asset (e.g., housing, or retirement wealth). Illiquid assets earn a higher 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. In particular, by matching the observed wealth-income ratio, we guarantee that our findings are not driven by implausibly small discount factors which make agents highly impatient.

Besides the usual small fraction of agents with zero net worth, our model features a significant number of what we call “wealthy hand-to-mouth” households. These are households who hold sizeable amounts of illiquid wealth, yet optimally choose to consume all of their disposable income during a pay-period. Examining asset portfolio and income data from the 2001 Survey and Consumer Finances (SCF) through the lens of our two-asset model reveals that between 1/4 and 1/3 of US households fit this profile.⁴ Although in our model these households act as if they are “constrained,” such

³JPS (2006) conclude 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.*” (page 1607). SS (2003a) conclude theirs with “*key parameters such as the propensity to consume are contingent on aggregate conditions in ways that are difficult to anticipate.*” (page 394)

⁴Our finding is consistent with recent survey evidence in Lusardi, Schneider and Tufano (2011) on financially fragile households. Their data show that almost half of US households would be probably

households would not appear constrained from the viewpoint of the one-asset model since they own substantial net worth.⁵

It is because of these wealthy hand-to-mouth households that the 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 the payment. When we replicate, by simulation, the randomized experiment associated with the tax rebate of 2001 within our structural model, we find consumption responses of comparable magnitude to those estimated in the micro data, i.e., around 25%. Two key features of the macroeconomic environment of 2001 act as important amplification mechanisms in the model: the Bush tax reform, and the 2001 recession. By raising future permanent income relative to current income, both components exacerbate liquidity constraints at the time of the fiscal stimulus payment.

The presence of wealthy hand-to-mouth households represents the crucial amplification mechanism relative to a plausibly calibrated one-asset model economy where, as we demonstrate, average consumption responses to the fiscal stimulus payments are only 3%.⁶ This pronounced amplification mechanism is present even when agents in the model are allowed access to credit.

Existing macroeconomic applications of the Baumol-Tobin model are essentially limited to financial issues and monetary policy.⁷ We argue that this is also a natural environment to analyze fiscal policy. Our paper shows that combining the Baumol-Tobin model with an incomplete-markets life-cycle economy gives rise, endogenously, to a sig-

or certainly unable to “come up with \$2,000 within a month.”

⁵Recently, Hall (2011) has also adopted the view that the degree of illiquidity in household wealth is useful to understand aggregate fluctuations and the effects of macroeconomic policy.

⁶Clearly, even one-asset models could, under certain extreme parameterizations where many agents hold close to zero net worth and are very often constrained, predict nontrivial consumption responses. This explains the sizeable MPC out of lump-sum tax cuts reported in some of Heathcote’s (2005) experiments on quantifying deviations from Ricardian neutrality in this class of models.

⁷Building on Miller and Orr (1966), Frenkel and Jovanovic (1980, 1981) study the optimal precautionary demand for money, and the optimal holdings of international reserves in a stochastic framework. Jovanovic (1982) analyzes the welfare cost of inflation; Rotemberg (1984) examines real effects of monetary policy. Romer (1986) and Chatterjee and Corbae (1992) develop a deterministic, OLG version of the Baumol-Tobin model that, as we explain below, bears some resemblance to our model during the retirement phase. More recently, within equilibrium complete markets models, Alvarez, Atkeson and Kehoe (2002) and Khan and Thomas (2010) study how transaction costs lead to endogenous asset market segmentation and real effects of monetary injections. Within incomplete-markets economies, Aiyagari and Gertler (1991) focus on the equity premium and the low frequency of trading equities; Imrohorglu and Prescott (1991) introduce a fixed per-period participation cost of bond holdings to support equilibria where money has value; Erosa and Ventura (2002) and Bai (2005) revisit, quantitatively, the question of welfare effects of inflation; Ragot (2011) studies the joint distribution of money and financial assets.

nificant presence of hand-to mouth households. Thus, properly designed government transfers and tax cuts can have substantial immediate impact on the macroeconomy, and short-run deviations from Ricardian neutrality can be large.

Our model is closely related to a recent and growing literature that studies portfolio choice in the presence of a frictional transactions technology. Alvarez and Lippi (2009) extend the classical Baumol-Tobin setup to a dynamic environment where investors can make free withdrawals at random, exogenously determined times. This creates a precautionary demand for holding liquid assets. Alvarez, Guiso and Lippi (2010) extend this model in the spirit of Abel, Eberly and Panageas (2009) to allow for observation costs as well as monetary transactions costs.⁸ Although our model is less analytically tractable than most of this literature, it contains a number of new features, all of which are crucial for generating wealthy hand-to-mouth households and empirically plausible rebate coefficients: endogenous non-durable consumption choices, borrowing constraints, non-financial income, and a lifecycle saving motive. This last feature leads to a problem of choosing when and how much to deposit into the illiquid account during the working life, rather than to withdraw, which we show in Section 3 to be at the heart of our model, and lacking from this literature.

Two recent papers examine whether the existence of two assets featuring different return and liquidity characteristics induces “excess sensitivity” in consumption. Li (2009) is an infinite-horizon model. The absence of a life-cycle prevents the model from featuring wealthy hand-to-mouth households. 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. This channel, working through impatience and the concavity of the consumption function –and not through transaction costs– is well known also in one-asset economies. The model of Huntley and Michelangeli (2011) focuses exclusively on the distinction between taxable and tax-deferred assets. Even though this is a life-cycle model, the amplification in the MPC, relative to the benchmark one-asset model, is barely significant (between 2 and 4 pct points), once the model is plausibly parameterized.

Since Campbell and Mankiw (1989), it has been argued that some aspects of the data are best viewed as generated not by a single forward-looking type of consumer, but rather by the coexistence of two types of consumers: one forward-looking and

⁸They use data on the frequency with which Italian investors observe and adjust their portfolios to estimate the size of both types of costs. They find that while observation costs are negligible, transaction costs average around 1% of the value of durable consumption, which is close to the calibrated costs we report in Section 5.

consuming its permanent income (the saver); the other following the “rule of thumb” of consuming its current income (the spender; see also Mankiw, 2000). Our model can be seen as a microfoundation for this view since it endogenously generates “wealthy hand-to-mouth” households alongside standard buffer-stock consumers.⁹

A natural question is: why would households with sizeable net worth optimally choose to consume all of their randomly fluctuating earnings every period, instead of smoothing shocks? The answer is that such households are better off taking the welfare loss rather than smoothing consumption because the latter option entails either (i) paying the transaction cost more often to withdraw cash when needed to smooth shocks, or (ii) holding large balances of cash and foregoing the high return on the illiquid asset. This explanation is reminiscent of calculations by Cochrane (1989), and more recently by Browning and Crossley (2001), who show that in some contexts the utility loss from setting consumption equal to income, instead of fully optimizing, is second order. Clearly, upon large downward income shocks, households in our model would withdraw from the illiquid account to smooth consumption. This rich adjustment pattern resembles that shown 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.

An important implication of Cochrane’s (1989) remark is that small tax rebates may not be a powerful validating source of data for choosing between competing structural models of consumption behavior. Chetty (2011) makes a similar argument in the context of labor supply choices. We take this point seriously and show that a number of additional implications of the model are consistent with the data: (i) the heterogeneity in consumption responses (Misra and Surico, 2011); (ii) the correlation between consumption response and holdings of liquid wealth (Broda and Parker, 2008); (iii) the size-asymmetry of the responses to large and small payments (Hsieh, 2003); (iv) their dependence on the aggregate state of the economy (JPS, 2009); (v) the increase in consumption inequality over the life-cycle (Heathcote et al., 2010).

⁹The model by Campbell and Hercowitz (2009) also features “wealthy constrained” agents endogenously, but through a very different mechanism. It assumes that, periodically, households discover they will have a special consumption need T periods ahead (e.g., education of their kids). This induces them to consume low amounts until they have saved enough for the special consumption need. During this phase, they may have large MPC out of unanticipated transitory income. However, in their model the response to the news would also be strong, and hence rebate coefficients estimated in the micro data as the difference between households who receive the check and those who receive the news (e.g., JPS, 2006) would be of negligible magnitude.

Our approach to explaining the consumption response to fiscal stimulus payments incorporates “frictions” in asset markets but abstracts from “behavioral” departures from standard consumption theory. A number of frameworks (e.g., myopia, hyperbolic discounting, self-control, mental accounting, rational inattention) have the ability to generate large MPCs, especially out of windfalls.¹⁰ However, without the addition of some form of transaction cost, they cannot immediately generate small consumption responses to *news* about future payments, a necessary condition to match the evidence we discuss.

The rest of the paper proceeds as follows. In Section 2, we describe the 2001 tax rebates and present the associated empirical evidence on consumption responses. In Section 3, we outline our model, and present a series of examples from a simplified version to convey intuition about how the model works. In Section 4, we describe our parameterization. Section 5 contains the quantitative analysis of the 2001 rebates and explores several additional implications of the model. In Section 6, we perform a number of robustness checks. Section 7 concludes.

2 Empirical estimates of the consumption response to the 2001 tax rebate

We focus our attention on the tax rebate episode of 2001 for two reasons. First, the 2001 tax rebates are the most extensively studied, and the ones for which we have the richest set of empirical evidence on household consumption responses. Second, we are able to obtain data from the Survey of Consumer Finances (SCF) about households’ balance sheets in 2001. As we discuss in Section 4, this is a crucial input into our analysis, and such data does not yet exist for the period surrounding the latest recession.

The 2001 tax reform 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 Congress. The reform decreased federal personal income tax rates at all income brackets, including a reduction in the tax rate on the first \$12,000 of earnings for a married couple filing jointly (\$6,000 for singles) from 15% to 10%. The majority of these changes were phased in gradually over the five years 2002-2006. According to the bill passed in Congress, the entire Act would “sunset” in 2011. Instead the bill was

¹⁰Existing applications to consumption behavior include, among others, Thaler (1990), Laibson (1998), Angeletos et al. (2001), Reis (2006); Luo (2008).

ultimately renewed in December 2010 for a further two years.

The tax rebate The reduction from 15% to 10% for the lowest bracket was effective retroactively to January 2001. In order to make this component of the reform highly visible during calendar year 2001, the Administration decided to pay an advance refund (informally called a tax rebate) to taxpayers for money they would have received from the Treasury only upon filing their tax returns in April 2002. The vast majority of the checks were mailed between the end of July and the end of September 2001, in a sequence based on the last two digits of 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 which week they would receive their check. The Treasury calculated that 92 million taxpayers received a rebate check, with 72 million receiving the maximum amount, (\$600, or 5% of \$12,000, for married couples). Overall, the payments amounted to \$38B, which is just below 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 received it; and (iii) the timing of receipt of the rebate has the feature of a randomized experiment because the last two digits of SSN are uncorrelated with any individual characteristics.¹¹

Empirical evidence JPS (2006) use questions added to the Consumer Expenditure Survey (CEX) that ask about the timing and the amount of the rebate check. Among the various specifications estimated by JPS (2006), we will focus on the most natural one:

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

where Δc_{it} is the change in nondurable expenditures of household i in quarter t , month_s is a time dummy, $X_{i,t-1}$ 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 comes from randomization

¹¹Whether the rebate was perceived as permanent or transitory is less clear. First, it depends on what beliefs households held with respect to the sunset. Second, the rebate was more generous than its long-run counterpart in the tax reform: according to Kiefer et al. (2002), there was a sizeable one-off component in the rebate amounting to approximately 1/3 of the total. In Section 6, we experiment with different sunseting scenarios.

Table 1: Estimates of the 2001 Rebate Coefficient ($\hat{\beta}_2$)

	Strictly Nondurable	Nondurable
JPS 2006, 2SLS ($N = 13,066$)	0.202 (0.112)	0.375 (0.136)
H 2008, 2SLS ($N = 12,710$)		0.242 (0.106)
MS 2011, IVQR ($N = 13,066$)		0.244 (0.057)

Notes: Strictly nondurables includes food (at home and away), utilities, household operations, public transportation and gas, personal care, alcohol and tobacco, and miscellaneous goods. Nondurables also includes apparel good and services, reading materials, and out-of-pocket health care expenditures. JPS 2006: Johnson, Parker and Souleles (2006); H 2008: Hamilton (2008); MS 2011: Misra and Surico (2011). 2SLS: Two-Stage Least Squares; IVQR: Instrumental Variable Quantile Regression.

in the *timing* of the receipt of rebate checks across households. However, since the size of the rebate is potentially endogenous, JPS (2006) estimate equation (1) by 2SLS using an indicator for whether the rebate was received as instrument. Table 1 reproduces estimates from JPS (2006). Their key finding is that β_2 is estimated between 0.20 and 0.40, depending on the exact definition of nondurables.

Since the original influential estimates of JPS (2006), others have refined this empirical analysis. Hamilton (2008) argues that the CEX is notoriously noisy, and one should trim the sample to exclude outliers. When a small number of extreme records are deleted from the sample (of roughly 13,000 observations), the rebate coefficient on nondurables drops to 0.24. Misra and Surico (2011) use quantile regression techniques to explicitly account for heterogeneity in the consumption response across households. Their point estimate is, again, around 0.24 and, more importantly, the rebate coefficient is much more precisely estimated.¹² We conclude that estimates of the rebate coefficient range between 0.20 and 0.40, with the most recent estimates putting more weight towards the lower bound.

It is crucial to understand the exact meaning of the rebate coefficient β_2 . Because households who do not receive the check at date t (the control group) may know they will receive it in the future, β_2 should be interpreted as the fraction of the rebate consumed in the quarter it was received, net of the consumption response of households in the control group which, in general, is not zero. As a result β_2 is not the MPC out of

¹²More precise estimates of the order of 20%-25% are obtained by PSJM (2011) in the context of the 2008 episode using the same econometric methodology as in JPS (2006). This additional evidence reinforces the finding that consumption responses to fiscal stimulus payments are significant, but since the economic environment and the program design in 2008 were different from those in 2008, one should be cautious in drawing conclusions. We return to this point in the Conclusions.

the rebate, a point on which the existing literature is somewhat fuzzy. This qualification has two consequences. First, to be able to generate a large value for β_2 , a model must feature a large MPC out of transitory shocks –but this is a necessary condition, not a sufficient one. The model must also feature a low MPC out of the news of the shock. In the absence of this second requirement, $\hat{\beta}_2$ would become small since it would be the difference between two equally large numbers. As argued in the Introduction, this observation is useful in distinguishing among competing theories. Second, since $\hat{\beta}_2$ is not a MPC, it cannot be used to draw inference on the impact of the policy on aggregate consumption. We return to this point in Section 5.3.

3 A life-cycle model with liquid and illiquid assets

We now describe our framework. The model integrates the Baumol-Tobin inventory-management model of money demand into an incomplete-markets life-cycle economy. In this section we outline the model in steady-state. We use a series of examples to highlight the economic mechanisms at work. Then, in Section 4, we introduce two additional features needed to model the rebate experiment: the tax reform and recession.

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 time-separable, expected lifetime utility given by

$$\mathbb{E}_0 \sum_{j=1}^J \beta^{j-1} \frac{c_{ij}^{1-\gamma} - 1}{1-\gamma} \quad (2)$$

where c_{ij} is consumption of nondurables for household i at age j .

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

$$\log y_{ij} = \chi_j + \alpha_i + \psi_i j + z_{ij}, \quad (3)$$

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

Financial assets Households can hold two assets: a liquid asset m_{ij} , and an illiquid asset a_{ij} . The illiquid asset pays a gross return $R^a = 1/q^a$, while positive balances of the liquid asset pay a gross return $R^m = 1/q^m$. We assume that returns are exogenous and note that, since they are real after-tax returns, they could be below one. When the household wants to make deposits into, or withdrawals from, the illiquid account, it must pay a transaction cost κ .¹³ When $R^a > R^m$, there is a meaningful trade-off between holding the two assets. Households start their working lives with an exogenously given quantity of each asset.

Illiquid assets are restricted to be non-negative, $a_{ij} \geq 0$, but we allow borrowing in the liquid asset to reflect the availability of credit up to an ad-hoc limit, $\underline{m}_{j+1}(y_{ij})$ which we allow to vary with age and individual income. The 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 $m_{i,j+1} < 0$.

Government Government expenditures G are not valued by households. Retirees receive social security benefits $p(\chi_{J^w}, \alpha_i, \psi_i, z_{iJ^w})$ 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. For retirees, the same tax function applies with $p(\chi_{J^w}, \alpha_i, \psi_i, z_{iJ^w})$ in place of y_{ij} . We denote the combined income tax liability function as $\mathcal{T}(y_{ij}, a_{ij}, m_{ij})$. 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 $\mathbf{s}_j = (m_j, a_j, \alpha, \psi, z_j)$ 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

¹³It 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. We return to this point in Section 6.

consuming.¹⁴

Consider a working household with age $j \leq J^w$. 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:

$$\begin{aligned}
V_j^0(\mathbf{s}_j) &= \max_{c_j, m_{j+1}} u(c_j) + \beta \mathbb{E}_j [V_{j+1}(\mathbf{s}_{j+1})] & (4) \\
\text{subject to} & : \\
q^m(m_{j+1}) m_{j+1} + (1 + \tau^c) c_j &= y_j - \mathcal{T}(y_j, a_j, m_j) + m_j \\
q^a a_{j+1} &= a_j \\
m_{j+1} &\geq \underline{m}_{j+1}(y_j) \\
y_j &= \exp(\chi_j + \alpha + \psi j + z_j).
\end{aligned}$$

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

If the household chooses to adjust its holding of illiquid assets because $V_j^0(\mathbf{s}_j) < V_j^1(\mathbf{s}_j)$, then it solves:

$$\begin{aligned}
V_j^1(\mathbf{s}_j) &= \max_{c_j, m_{j+1}, a_{j+1}} u(c_j) + \beta \mathbb{E}_j [V_{j+1}(\mathbf{s}_{j+1})] & (5) \\
\text{subject to} & : \\
q^m(m_{j+1}) m_{j+1} + q^a a_{j+1} + (1 + \tau^c) c_j &= y_j - \mathcal{T}(y_j, a_j, m_j) + m_j + a_j - \kappa \\
a_{j+1} &\geq 0 \\
m_{j+1} &\geq \underline{m}_{j+1}(y_j) \\
y_j &= \exp(\chi_j + \alpha + \psi j + z_j).
\end{aligned}$$

The problem for the retired household of age $j > J^w$ is similar, with benefits $p(\chi_{J^w}, \alpha_i, \psi_i, z_{iJ^w})$ in place of earnings y_j . Appendix A1 describes in detail the computational algorithm used to solve this problem.

Equilibrium Given initial conditions on asset holdings $\{m_0, a_0\}$, and asset prices $q^m(m)$ and q^a , households optimize. The intertemporal government budget constraint

¹⁴The timing of the earnings shock and adjustment decisions implies that our model *does not* feature a cash-in-advance (CIA) constraint. In models with a CIA constraint, during a period of given length the household is unable to use a certain fraction of its current income to consume. In our model, after the earnings shock, the household can always choose to pay the transaction cost, withdraw from the illiquid account, and use all its resources to finance consumption. Put differently, the time for which some of the agent's income is tied in the illiquid asset and unavailable for consumption expenditure is entirely under the agent's control. See Jovanovic (1982) for an exhaustive discussion of the difference between models with transaction costs and models with CIA constraints.

is balanced so that

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

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

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

Two Euler equations Consumption and portfolio decisions are characterized by a “short-run” Euler equation (EE-SR) that corresponds to (dis)saving in the liquid asset, and a “long-run” Euler equation that corresponds to (dis)saving in the illiquid asset (EE-LR). To understand, consider a deterministic version of the model without income risk, borrowing, and taxes.

In periods where the working household is unconstrained and does not adjust:

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

The slope of her consumption path is governed by β/q^m which, for plausible parameterizations, is below one. Hence consumption declines over time because of impatience and the low return on cash. A constrained household consumes all her earnings, i.e., $c_j = y_j$.

During the working life, the agent saves to finance retirement by making periodic deposits into the illiquid account. Given the fixed cost of adjusting, households accumulate liquid assets and choose dates at which to deposit some or all of their liquid holdings into the illiquid account (the “cake-baking” problem). Across two such adjustment dates N periods apart, consumption dynamics 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 in 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 accord-




Figure 1: Example of lifecycle in the two-asset model

ing to (EE-SR). The withdrawals are timed to coincide with the period where cash is exhausted. Across withdrawals, equation (EE-LR) holds.¹⁵

Figure 1 shows consumption and wealth dynamics in an example where an agent with logarithmic utility starts her working life with zero wealth, receives a constant endowment while working and a lower endowment when retired, and cannot borrow.¹⁶ Panel (a) shows that the agent in this example chooses to adjust his illiquid account at only three points in time: one deposit while working, and two withdrawals in retirement. In between, the value of the illiquid account grows at rate $1/q^a$. Panel (b) shows her associated income 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 $1/q^m$, 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, under this parameterization, the young agent is constrained for the initial phase of her working life, when her net worth is zero.

Endogenous “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,

¹⁵In this example, our problem during retirement with no discounting, log utility, and the transaction cost expressed in utility terms would coincide with Romer (1986) and withdrawal dates would be equidistant, subject to the “integer problem” intrinsic in the discrete-time formulation.

¹⁶To make this example as stark as possible, we impose a very large transaction cost.




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

except for a higher return on the illiquid asset R^a . This higher return leads to stronger overall wealth accumulation. Rather than increasing the *number* of deposits during the 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, instead of being constrained at the beginning of the life cycle, the household optimally chooses to hold zero liquid assets in the middle of the working life, after its deposit, while the illiquid asset holdings are positive and are growing in value. Intuitively, since its net worth is large, this household would like to consume more than its earnings flow, but the transaction cost dissuades it 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, cannot increase its expenditures without making a costly withdrawal from the illiquid account.

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? In this deterministic example, 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; and/or (ii) holding larger balances of liquid wealth hence foregoing the high return on the illiquid asset (and, therefore, the associated higher level of long-run consumption). This logic continues to apply in the presence of shocks, as we will see in our simulations: wealthy hand-to-mouth households are willing to absorb welfare losses

from random consumption fluctuations because smoothing consumption would require either (i) holding low precautionary liquid balances, but frequently paying the transaction cost, or holding large precautionary liquid balances with a significant loss in terms of return. This observation is reminiscent of Cochrane’s (1989) insight that, in a representative agent model with reasonable risk aversion, the utility loss from setting consumption equal to income is tiny.¹⁷

4 Calibration

We now present a calibration of the model without credit ($\underline{m} = 0$) and without heterogeneous earnings slopes ($\psi_i = 0$). In Section 6 we analyze these extensions of the model.

Demographics and preferences 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 assume a unitary intertemporal elasticity of substitution ($\gamma = 1$) and we calibrate the discount factor β to replicate median illiquid wealth in the SCF (see below).¹⁸

Earnings heterogeneity We estimate the parameters of the earnings process (common life-cycle profile, initial variance of earnings, and variance of earnings shocks) through a minimum distance algorithm that targets the empirical covariance structure of household earnings constructed from the Panel Study of Income Dynamics (PSID). Specifically, from the PSID we select a sample of households with heads 22-59 years old in 1969-1996, following the same criteria as in Heathcote, Perri, and Violante (2010). We construct the empirical mean age-earnings profile and covariance functions in levels and first differences, exploiting the longitudinal dimension of the data. We simulate

¹⁷See also Browning and Crossley (2001) for a similar calculation in the context of the life-cycle model of consumption.

¹⁸In 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 more closely the wealth distribution for the vast majority of households, with the exception of the upper tail that, notoriously, 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, therefore, understating 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 a plausible distribution of MPC across the population is far more important than aggregate price effects, for the question at hand.

the process in (3) at a quarterly frequency, also allowing for an i.i.d. shock, and aggregate quarterly earnings into annual earnings. From the implied annual earnings we construct the theoretical counterpart of the empirical moments and minimize the distance between the two set of moments. We interpret the quarterly i.i.d. component as measurement error in earnings and hence, in simulating the model, we abstract from it.

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 labor income for the bottom 95% is \$52,696, a number close to the one reported by JPS (2006, Table 1).¹⁹ Our definition of liquid assets comprises: money market, checking, savings and call accounts plus directly held mutual funds, stocks, bonds, and T-Bills net of revolving debt on credit card balances.²⁰ All other wealth, with the exception of equity held in private businesses, is included in our measure of illiquid assets. It comprises housing net of mortgages and home equity loans, vehicles net of installment loans, retirement accounts (e.g., IRA, 401K), life insurance policies, CDs, and saving bonds. Table 2 reports some descriptive statistics.

As expected, the bulk of household wealth is held in (what we call) illiquid assets, notably housing, vehicles, and retirement accounts. For example, the median of the liquid and illiquid asset distributions are, respectively, \$2,700, and \$70,000. Figure 3(a) shows the evolution of illiquid and liquid wealth over the life-cycle. It is clear that saving over the working life takes place in illiquid wealth, whereas liquid wealth is fairly constant.

Measuring hand-to-mouth households in the SCF These low holdings of liquid wealth in the data suggest that a number of households might be hand-to-mouth, i.e., we observe positive balances on their liquid accounts only because labor income is paid as cash and because of a mismatch in the timing of consumption and

¹⁹In our definition of household labor income, we included unemployment and disability insurance, TANF, and child benefits.

²⁰The SCF asks the following questions about credit card balances: (i) “How often do you pay on 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.

Table 2: Household Portfolio Composition

	Median (\$2001)	Mean (\$2001)	Fraction Positive	After-Tax Real Return (%)
Earnings plus benefits (age 22-59)	41,000	52,696	–	–
Net worth	77,100	164,463	0.95	5.5
Net liquid wealth	2,700	30,531	0.77	-1.1
Cash, checking, saving, MM accounts	1,880	12,026	0.87	-2.0
Directly held MF, stocks, bonds, T-Bills	0	19,920	0.28	4.1
Revolving credit card debt	0	1,415	0.33	–
Net illiquid wealth	70,000	133,932	0.93	6.2
Housing net of mortgages	31,000	72,585	0.68	7.1
Vehicles net of installment loans	11,000	14,562	0.86	5.8
Retirement accounts	950	34,431	0.53	4.5×1.35*
Life insurance	0	7,734	0.27	0.5
Certificates of deposit	0	3,805	0.14	1.3
Saving bonds	0	815	0.17	0.5

Source: Authors' calculations based on the 2001 Survey of Consumer Finances (SCF).

*The return on retirement assets is multiplied by a factor of 1.35 to account for the employer contribution. See the main text for details

earnings *within* a pay period; not because they hold cash *across* pay periods (i.e., they save through cash). To measure the fraction of these hand-to-mouth households, it is useful to note that *non* hand-to-mouth households would always hold average cash balances, during the pay-period, above half their earnings (where the “half” presumes resources being consumed at a constant rate).²¹ As a result, a lower bound estimate for hand-to-mouth households is the fraction of population with average balances of liquid wealth below half their earnings per pay-period.

Figure 3(b) provides an estimate of the number of such households, using three alternative assumptions on the frequency of payment: weekly, biweekly, and monthly.²²

²¹If 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. Alvarez and Lippi (2009) suggest a similar calculation as a test of the classic liquidity management model.

²²The amount on checking, saving, and money-market accounts reported in SCF is the “average over the last month.” However, for directly held mutual fund, stocks, and bonds (the other component of liquid wealth in our definition) the SCF reports the value at the date of the interview. Since holdings of this wealth component are minor and rare among US households (they are zero for over 80% of




Figure 3: Liquid and illiquid wealth over the life-cycle in the data (2001 SCF)

Between 30% and 42% of households appear to be hand-to-mouth in terms of liquid wealth according to this definition.²³ In the same figure, we report the same statistic for net worth, which is the relevant definition for comparison with one asset models. In terms of net worth, instead, only between 5.6% and 7.1% of households aged 22-79 are hand-to-mouth. The grey area in between the two solid lines (which refers to biweekly payments) is composed of households who have positive illiquid wealth but do not carry positive balances of liquid assets between pay periods. These are the empirical counterpart of the “wealthy hand-to-mouth” agents in our model. Note that this last group of households, which represents a sizeable fraction of the population (between 24% and 35%), is only visible through the lens of the two-asset model. From the point of view of the standard one asset model, these are households with positive net worth, and are hence unconstrained.

These findings are consistent with recent survey evidence in Lusardi, Schneider, and Tufano (2011) showing that almost one third 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.

households), our estimate should not be affected by this problem.

²³The plausibility of this estimate can be checked through another survey question: “Over the past year, would you say that your family’s spending exceeded your family’s income, that it was about the same as your income, or that you spent less than your income?” The fraction of households answering “exceeding” or “equalled” to this question, and answering “No” to the companion question “Did any of that spending include purchases of a home or automobile or spending for any investments?” is 37%, hence in line with our first estimate.

Asset returns To measure returns on the various asset classes, we focus on the period 1960-2009. All of the following returns are expressed in annual nominal terms. We set the nominal return on checking accounts to zero and the return on savings accounts, T-Bills and savings bonds to the interest rate on 3-month T-Bills, which was 5.3% over this period (FRB database). 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%.²⁴ 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 bond components of each saving instrument. An important feature of some 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 1.35.²⁵

For housing, we follow Favilukis, Ludvigson, and Van Nieuwerburgh (2010). We measure housing wealth for the household sector from the Flow of Funds, and housing consumption from the National Income and Products Accounts. The return for year t is constructed as housing wealth in the fourth quarter of year t plus housing consumption over the year minus expenses in maintenance and repair divided by housing wealth in the fourth quarter of year $t - 1$.²⁶ We subtract population growth in order to correct for the growth in housing quantity. We also subtract the average property tax, 1% (Tax Foundation, 2011). As a result, we obtain an average annual nominal return of 13.2%.²⁷

Given the absence of data on the service flow from vehicles (autos, trucks and motorcycles), we adopt a user cost approach to calculate the return on vehicles. The sum of

²⁴After inflation (4.1% over this period), the real return is 7%. Favilukis, Ludvigson, and Van Nieuwerburgh (2010, Table 5), whose calculations we follow, report returns between 7.9% (1953-2008) and 6.6% (1972-2008).

²⁵Since we do not model explicitly the tax-deferred treatment of retirement accounts, we somewhat underestimate the effective return on this saving vehicle.

²⁶Our estimate of expenses in improvement and repair is based on a comprehensive study on housing compiled by the Joint Center for Housing Studies (2011). Figure 2 reports that these expenses amount to roughly 40% of total residential investment expenditures.

²⁷After inflation, the real return is 9.1%. Favilukis et al. (2010, Table 5) report returns between 9.8% (1953-2008) and 9.9% (1972-2008). They also report a housing return of 9.1% (1972-2008) computed based on the repeat-sale Freddie Mac Conventional Mortgage House Price index for purchases only (Freddie Mac) and on the rental price index for shelter (BLS).

the interest rate on T-Bills plus the rate of economic and physical depreciation (available from the BEA) yields an annual nominal return of 11.6%. For saving bonds and life insurance (assuming actuarially fair contracts for the latter) we use the return on T-Bills, and for CDs we obtain a nominal return of 6.3% (FRB database).

We apply these nominal returns to each household portfolio in the SCF and compute average returns in the population. The implied average nominal return on illiquid wealth is 12.1% and on liquid wealth is 3.8%. Finally, we set the annual inflation rate to 4.0 percent (the average over this period is 4.1). After inflation and taxes (see below), the after-tax real returns on liquid and illiquid wealth are 6.2% and -1.1%, respectively. The final after-tax real returns on liquid and illiquid wealth by category are reported in Table 2.

Initial asset positions We use observed wealth portfolios of SCF households aged 20 to 24 to calibrate the age $j = 0$ initial asset positions in the model. We divide these households into twenty-one sub-groups based on their earnings and calculate 1) the fraction with zero holdings, and 2) 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.²⁸

Government We set $\tau^{ss}(y)$ 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. To compute social security benefits, our proxies for individual average lifetime earnings $Y_{i,J^w} = \exp(\chi_{J^w} + \alpha_i + \psi_i J^w + z_{i,J^w})$ are run through a formula based on replacement rates and bend points as in the actual system in the year 2000. The effective consumption tax rate τ^c is set to 7.2% (McDaniel, 2007). The function $\tau^y(y)$ is 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. Kiefer 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 in 2000. We apply these tax rates on each household portfolio in our sample and derive an average effective tax of 22.9% on income from liquid assets (τ^m), and of 15.9% on income from illiquid assets (τ^a).²⁹ The price of government debt q^g is set equal to that on illiquid

²⁸For example, in the median initial earnings group, the fraction of households with zero liquid (illiquid) wealth is 24% (14%). For those with positive holdings, median liquid wealth is \$700, and median illiquid wealth is \$4,400.

²⁹We apply the interest/dividend tax rates on all assets except for housing and for the equity

assets. When we set government expenditure G to its value in the year 2000 (using aggregate labor income computed from “Wages and Salaries Disbursements” in the NIPA Table 2.1 as a metric), residual public debt B from the government budget constraint is close to its observed value for that year.

4.1 Modelling the 2001 tax rebate, tax reform, and recession

Tax rebate We assume that the economy is in steady state in 2001:Q1. The rebate checks are then randomly sent out to half the eligible population in 2001:Q2, and to the other half in 2001:Q3. We set the rebate size to \$500 based on JPS (2006) who report that the average rebate check was \$480 per household.

There are different views that one could plausibly take 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 the news arriving in 2001:Q1. Under this timing, the check is fully anticipated by both groups. At the other extreme, one could assume that households became aware of the rebate only after receiving their own check: under this assumption, both groups of households treat the rebate as a surprise. In our baseline, we take an intermediate position, i.e., all households learn about the rebate when the first set of Treasury checks are received, in 2001:Q2. Under this timing, the check was fully anticipated only by the second group who received the check in 2001:Q3. We explore the two alternative timing assumptions in Section 6. We assume throughout that the news/check reaches households before their consumption/saving and adjustment decisions for that quarter.

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 roughly 3% or more. We construct the sequence of effective tax schedules based on Kiefer et al. (2002).³⁰ Most of these changes were planned to be phased in gradually over the five years 2002:Q1-2006:Q1 and to “sunset” in 2011. It turned out that instead of sunsetting, the tax cuts were further extended. We consider two scenarios regarding the “sunset” clause: (i) households believe that the tax system will revert to its pre-reform state at the end of ten years; and (ii) households act as if the change in the tax system is permanent after the reform is fully phased in. A tax reform is defined as a sequence of

component of retirement accounts, on which we apply the capital gain tax rates.

³⁰Kiefer 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 (page 90).

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 2006:Q1 in absence of sunset, and 2011:Q1 when the tax reform sunsets as originally legislated.

Recession To model the downturn of 2001, we assume that in 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 1.5%, and will then fully recover over the following eight quarters.³¹

Rebate financing Our baseline assumption is that the government finances the rebate program by increasing debt for ten years, and then repays the rebate outlays and the accumulated interest on the new debt by introducing a permanent proportional tax on earnings. In our benchmark calibration, the required additional tax rate is around 0.05%. When we model the full tax reform, we reduce government expenditures accordingly. As we show in Section 6, our findings are virtually independent of the financing scheme.

5 Quantitative analysis of the 2001 tax rebate

We start by studying a baseline economy where the tax rebate occurs in isolation. Next, we incorporate the tax reform and recession. We analyze the model economy for values of the transaction cost κ ranging from zero to \$3,000. For each value of $\kappa > 0$, we re-calibrate β to match median holdings of illiquid wealth.³² The case $\kappa = 0$ corresponds to the standard one-asset model. At $\kappa = 0$, we set β to reproduce median net worth, and we set the interest rate to the average return on net worth in the SCF data (see Table 2).³³

Figure 4(a) shows that the fraction of households adjusting (i.e., accessing the illiquid account to withdraw or deposit) falls with the size of the transaction cost κ . As illustrated in the simulations of Section 3, retirees adjust more often than working-age

³¹The 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 the “Wages and Salaries Disbursements” series in NIPA Table 2.1.

³²The annualized values of β range between 0.950 and 0.956. Hence, our results are not driven by implausibly low discount factors which make households highly impatient.

³³This latter choice has no bearing on the findings since the discount factor is always calibrated to generate the same amount of net worth. When we set the interest rate to the calibrated value of R^m or R^a , we found virtually identical results.




Figure 4: Features of two-asset model, by transaction cost

households because they finance their consumption largely by withdrawing from the illiquid account. Holdings of liquid wealth increase with the transaction cost (Figure 4(b)), because when κ is larger households deposit into/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, around \$1,500.³⁴

Figure 4(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

³⁴Optimal holdings of liquid wealth are small in spite of the presence of plausibly calibrated earnings risk for three reasons. First, shocks are highly persistent. Second, the opportunity cost of holding cash is high, since $R^a - R^m$ is large. Third, households can always withdraw (at a cost) from the illiquid account in the event of a large negative shock. This view of earnings risk is consistent with the observed distribution of liquid wealth holdings: had we allowed for a large transitory earnings component, households would hold counterfactually high quantities of liquid assets.




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

wealth. The size of both groups is increasing in κ . As shown in Figure 4(d), the average length of spells in which hand-to-mouth households hold zero liquid assets and consume their income is also growing with the level of the transaction cost. Intuitively, a large value for κ makes it less likely that the household will withdraw from the illiquid account to smooth large negative earnings shocks.

Choice of transaction cost In what follows, we often focus on the range \$500-\$1,000 for the transaction cost: in this range, (i) the fraction of households that adjust each quarter is 15%-20% (4%-8% among workers and 35%-45% among retirees); (ii) median holdings of liquid wealth are close to their data counterpart in Table 2; and (iii) the fraction of hand-to mouth consumers is in line with the empirical estimates of Figure 3. Recall that the bulk of illiquid wealth in the data is composed by housing, vehicles, and retirement accounts. Transaction costs for housing are commonly estimated around 5% of the asset value (e.g., OECD, 2011). Alvarez and Lippi (2009) estimate 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). Relative to median illiquid wealth holdings (\$70,000), transaction costs between \$500 and \$1,000 appear therefore reasonable.

Baseline results Figure 5(a) displays the rebate coefficient in the model for different levels of the transaction cost. We compute the rebate coefficient exactly as in JPS (2006) –i.e., we run regression (1) on simulated panel data. The rebate coefficient grows steadily from 1.8% when $\kappa = 0$ (the one-asset model) to 21% when $\kappa = \$3,000$.

In the one-asset model with $\kappa = 0$, the vast majority of households in this model hold enough net worth to respond virtually identically to the news and to the check itself, as predicted by standard consumption theory. The action comes almost entirely from the constrained agents, most of which are very young (below age 35): those in the treatment group have a high MPC, while those in the control group do not respond at all. In our calibration, only 3 % of households have zero net worth and are hand-to-mouth, which explains the small rebate coefficient.³⁵

For transaction costs in the range \$500-\$1,000, the rebate coefficient is around 15%, or 8 times higher than its one-asset model counterpart. Figure 5(b) shows the marginal propensities to consume (MPC) out of the fiscal stimulus payment for two groups of households: those who are hand-to-mouth and those who are not. Note how, for the latter group, the average MPC is close to zero, while for the former group it is around 45%. Therefore, most of the households in the model behave exactly as predicted by the PIH and have zero MPC. The high rebate coefficients are driven by constrained households, and the two-asset model generates a larger fraction of hand-to-mouth consumers, some of which hold sizeable quantities of illiquid assets. Such households have significant MPC 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). As a back of the envelope calculation, the 15% rebate coefficient arises as a weighted average of zero MPC for 2/3 of the population and 45% MPC for the remaining 1/3 in the treatment group. In the control group, both hand-to-mouth and unconstrained agents have zero MPC.

For low transaction costs, marginal propensities to consume out of moderate income changes (and hence rebate coefficients) can be negative. When transaction costs are low enough, upon receiving the rebate, some agents choose to anticipate the adjustment decision and save the rebate, together with their current holdings of cash, into the illiquid account. As a result, they save more than the rebate amount (which explains the negative MPC in Figure 5(b)) and consume less than the control group (which explains the negative rebate coefficient in Figure 5(a)).

Tax reform and recession Figure 6(a) shows that the consumption responses to the tax rebate are substantially higher when the full tax reform is modeled. On average, the rebate coefficient increases by 7-8 percentage points. The reason is that the

³⁵As reported in Figure 3(b), in the 2001 SCF data this fraction is around 6%. Hence, even though the model at $\kappa = 0$ slightly underestimates this fraction, there is essentially no scope for the one asset model to generate significant rebate coefficients, while remaining consistent with SCF data on the distribution of net worth.




Figure 6: Effect of tax reform and recession on rebate coefficient, by transaction cost


substantial reduction in future tax liabilities leads to a rise in the desired level of lifetime consumption. Since a substantial fraction of households (both poor and wealthy) are constrained in terms of liquid wealth, the rebate enables such households to start consuming out of this additional future income earlier than they would otherwise be able to.³⁶ As is clear from the figure, our finding is robust to whether the sunset clause is implemented or not.

A similar logic applies when we add the recession to the tax rebate experiment. Figure 6(b) shows that allowing for the fact that the economy was undergoing a recession at the time of the fiscal stimulus payments adds roughly 3 percentage points to the rebate coefficient. Overall, in the range \$500-\$1,000 for κ , the model generates a rebate coefficient between 22% and 29%, in line with the empirical estimates of Table 1.

5.1 Further implications on consumption responses


Cochrane's (1989) observation that there may be only small welfare differences between alternative consumption behavior when reacting to small transitory income shocks means that it is especially useful to investigate additional implications of our model, a point also raised by Card et al. (2007), and Chetty (2011). Our model has a number of implications about households' consumption responses that can be compared to the

³⁶When adding the tax reform, the economy with $\kappa = 0$ also yields a higher rebate coefficient (see Figure 6(a)) because the effect we describe does apply there as well. However, quantitatively, the additional kick of the tax reform is negligible in the one-asset model because the number of constrained households is so small.



(a) Distribution of rebate coefficients in the population

(b) Average earnings by percentile of rebate coefficient distribution



(c) Average rebate coefficients by tercile of earnings distribution

(d) Average rebate coefficients by tercile of liquid wealth distribution

Figure 7: Heterogeneity in rebate coefficients in the model ($\kappa = \$750$)

data. We explore: 1) the heterogeneity in consumption responses across households; 2) their correlation with households' liquid wealth; 3) their size-asymmetry; and 4) their state-dependence. When, in order to present more detailed features of our model, it is necessary to select a specific value for κ , we choose $\kappa = \$750$ corresponding to a rebate coefficient of 27% (with tax reform and recession).

Heterogeneity Misra and Surico (2011) apply quantile regression techniques to the same data set as JPS (2006) to study cross-sectional heterogeneity in rebate coefficients. They conclude that there is a large amount of heterogeneity in consumption responses. There are two main findings. First, the distribution of consumption responses is bimodal, with around 40% of households saving all of the rebate, and another sizeable group of households spending a high fraction of the rebate.³⁷ Second, high income

³⁷A complementary source of evidence comes from SS (2003a, 2003b) who added a module to the Michigan Survey of Consumers to assess the impact of the rebate. This survey asked households what

households are disproportionately concentrated in the two tails of the distribution of consumption responses. As we illustrate below, this second finding is not inconsistent with (and could potentially explain) the result reported by JPS (2006) that, when splitting the population into three income groups, differences in rebate coefficient across groups are not statistically significant. Similarly, SS (2003a, 2003b) find no evidence of higher spending rates among low income households. Figure 7 shows that our model can broadly reproduce these findings.

Figure 7(a) plots a histogram of rebate coefficients for the working age population, in which the bimodal nature is stark. The bimodality arises because of the coexistence of a substantial fraction of (wealthy) hand-to-mouth consumers together with unconstrained agents who fully save the rebate as predicted by standard theory. Figure 7(b) plots median earnings in each quantile of the cross-sectional distribution of the consumption responses. The reason why there are high earnings households at both ends of the distribution is that some of them are unconstrained and some are wealthy hand-to-mouth. The former are income-rich but their expected earnings growth is low, the latter are income-rich but their expected earnings profile is steep, which makes the liquidity constraint more likely to bind. Moreover, because the rebate is a lump sum, among constrained agents the income-richest have the highest MPC. Figure 7(c) shows that average rebate coefficients are similar in each tercile of the earnings distribution. This is consistent with findings in JPS (2006), who split their sample in this way and do not find different responses across earnings sub-groups.

Correlation with liquid wealth JPS (2006) also estimate rebate coefficients for sub-groups of households with different amounts of liquid assets. They find that households in the bottom half of the distribution have substantially larger consumption responses. These effects are imprecisely estimated, though, for two 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. For example, 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 sub-sample.

they would do with their rebate check. 22 percent of respondents said they would mostly spend it, while the rest said they would mostly save it or repay debt. From these studies, we can infer that the average estimate from JPS is likely to be the outcome of very high MPC among a relatively small group of households and very small (or zero) MPC among the majority of the population. PSJM (2011) validate this survey-based finding by documenting that CEX households who report that they mostly consumed the 2008 rebate had consumption responses almost twice as large as those households who report that they mostly saved it.

Even a lack of a strong correlation between rebate coefficients and CEX measures of liquid assets would not be inconsistent with the model. In the model, 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.³⁸ This explains why, empirically, the relationship between rebate coefficients and the *level* of liquid wealth can be statistically weak.

The model’s sharpest prediction is that households who carry low levels of liquid wealth *across pay periods* should have strong consumption responses. Figure 7(d) illustrates this by plotting average rebate coefficients from the model in each tercile of the distribution of liquid assets carried across periods. 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 earnings ratios*. This is because for a hand-to-mouth household, the quantity of liquid assets that are held for within pay-period expenditures is, on average, half its earnings. Broda and Parker (2008) exploit the AC Nielsen Homescan database, a sample fifty times larger than the CEX, to study the consumption response to the fiscal stimulus payment of 2008. They 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.³⁹ Souleles (1999) studies the consumption response to anticipated tax refunds (whose median size is around \$560 (Souleles, 1999, Table 1)). When the sample is split between low and high liquid wealth to earnings ratio households, the former group is found to have statistically significant larger responses to the refund (Souleles, 1999, Table 4).

Size asymmetry Hsieh (2003) shows that the same CEX consumers who “over-react” to small income tax refunds respond very weakly to payments received from the Alaskan Permanent Fund. These payments are, on average, much larger (around \$2,000 per household) than the rebate we examined. 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

³⁸The balances on checking, saving, and money market accounts in the CEX are not averages as in the SCF, but refer to holdings on the last day of the month.

³⁹They ask households “*In case of an expected decline in income or increase in expenses, do you have at least two months of income available in cash, bank accounts, or easily accessible funds?*” Hence their liquid wealth variable is relative to income.




Figure 8: Rebate coefficients by size of the stimulus payment

December do not alter their consumption growth significantly in those months. One interpretation of these results is that, although households spend large fractions of small anticipated shocks, they predominantly save large anticipated shocks.

Figure 8 shows how, in our baseline economy, the propensity to consume the rebate declines when the size of the rebate is increased above \$500, as a function of the transaction cost κ . With a \$750 transaction cost, the rebate coefficient drops by a factor of three (from 16% 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 will find it optimal to save a portion of their payment.⁴⁰ 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, since adjusting households are not constrained, they will save a significant fraction of the rebate, as in the one-asset model.

This latter effect may be strong enough to cause the consumption response to fall also in *absolute* terms as the rebate size increases, for given transaction cost. Figure 8 shows that, when the transaction cost is \$750, the average consumption expenditure is larger for a \$1,000 rebate (\$125) than for a \$2,000 rebate (\$100). This example underscores the relevance for policy of understanding the structural mechanism that lies behind the empirical evidence.⁴¹

⁴⁰This occurs if the size of the rebate is larger than the amount by which consumption would increase if the household were not currently constrained.

⁴¹Figure 5.1 shows how estimated rebate coefficients may become negative for large ranges of the transaction cost, as for the case of a \$5,000 fiscal stimulus payment. As explained earlier, this occurs when the stimulus payment is large relative to the transaction cost. In this case, a substantial fraction

State dependence As documented in Figure 6(b), our model implies that recessions exacerbate the consumption response to tax rebates. 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 significant). This leads them 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.⁴²

5.2 Life-cycle implications

Figure 9 compares the life-cycle means and variances of earnings, consumption and net worth across the one-asset and two-asset models. Both models reproduce reasonably well the key features of the data (e.g., see Heathcote et al., 2010). The models are hardly distinguishable along these dimensions. Note, however, that consumption inequality grows slightly more in the two-asset economy, as households tolerate larger consumption fluctuations to avoid paying the transaction cost and to hold assets in the high-return (illiquid) saving vehicle.

5.3 Aggregate implications of the rebate program

To isolate the impact of the rebate program on the aggregate economy, we study the dynamics of aggregate consumption relative to a counterfactual economy which has all the features of the baseline (including recession and tax reform) except for the fiscal stimulus payments. Figure 10(a) plots the time path of the rebate distribution and consumption for the two groups. Figure 10(b) shows the aggregate impact of the policy: 13% and 18% of total rebate outlays are spent on nondurables, respectively, in the

of working households choose to make a deposit into the illiquid account upon receipt of the payment, and hence reduce their consumption in that period relative to that of the control group.

⁴²The state dependence can be quite complex. For example, a sufficiently sharp recession may induce households to pay the transaction cost and withdraw from the illiquid account to smooth consumption. In such a scenario, rebate coefficients are likely to be very low. In other words, the size of the recession matters too. We return to this point in the Conclusions.





Figure 9: Life-cycle profiles (means and variances) in the one-asset and two-asset models

first two quarters. The impact, though, is very short lived: a year after the rebate disbursement differential consumption growth is essentially zero. Overall, within the first year, the cumulative fraction of the rebate outlays spent on nondurable consumption is around 40%.

Since, in our model, we know the counterfactual path of aggregate consumption in the absence of the rebate (but in the presence of the tax reform and recession), it is possible to express the impact of the rebate as a percentage increase in aggregate consumption. We find it to be 0.62% and 0.85% in the first and second quarter, respectively. As is clear from the figure, the effect on consumption beyond the first two quarters is negligible.

These calculations highlight two benefits of having a structural model that is consistent with the micro evidence. First, because the tax reform and the recession do affect the aggregate economy even in the absence of the rebate, it is incorrect to measure



(a) Time path of rebate (dashed line) and consumption expenditures (solid line) for the two groups

(b) Percentage of total rebate outlays spent on nondurables by quarter


Figure 10: Aggregate consumption dynamics following the rebate

the additional aggregate consumption growth due to the rebate relative to a baseline of consumption in the period before the rebate announcement, as done for example in JPS (2006) or Misra and Surico (2011). Rather, the full path of counterfactual consumption (without rebate, but with recession and tax reform) should be taken into account. Second, to calculate the total portion of the rebate that is spent, it is incorrect to multiply the rebate coefficient by the total size of the rebate (as a fraction of aggregate consumption), as done also by the CBO (2009). The reason is that, as explained, the rebate coefficient β_2 is not the MPC out of the rebate, but the difference between the MPC for those who receive the check and that of the rest of the population. Interestingly though, Figure 10(a) shows that the rebate coefficient estimated for the *first* quarter is very close to the true MPC –since our model reveals that the first-quarter response of the control group is close to zero– and hence it represents a good proxy for simple policy analysis.⁴³


6 Robustness

In this Section, we discuss the robustness of our findings with respect to the timing of arrival of the news, the financing scheme of the rebate, the specific process for

⁴³The first-quarter response of the control group (who has the news but not the check) is close to zero because the unconstrained consumers have low MPC out of anticipated transitory income shocks, and the hand-to-mouth consumers cannot respond at impact.



(a) Rebate coefficients under alternative assumptions on timing of the arrival of news



(b) Rebate coefficients under alternative models for earnings heterogeneity

Figure 11: Robustness analysis with respect to arrival of news and earnings heterogeneity

individual earnings, the form of the transaction cost, and the availability of consumer credit.

Arrival of news In Figure 11(a), we report the consumption response to the tax rebate under alternative assumptions about when the news of the rebate enters households' information sets. When the rebate is a surprise for everyone, the rebate coefficient increases by around 5 percentage points on average. When it is anticipated by every household (the news arrives one quarter ahead for the first group and two quarters ahead for the second group), the estimated rebate coefficient drops by a similar amount. However, it still remains of a sizeable magnitude, between 17% and 25% in the \$500-\$1,000 range for κ . The reason is that liquidity constrained households are those driving the results and, for such households, learning that they will receive some income in the near future has little or no impact on their current consumption.⁴⁴

Financing of rebate We experimented by shortening from ten years to one year the duration of the transitory phase during which, before raising taxes, the government uses debt to finance the stimulus payments and found nearly identical results. There are two reasons. First, the financing scheme affects the unconstrained households in the treatment and the control group equally. Second, the behavior of hand-to-mouth

⁴⁴When the rebate is anticipated by all households, even though the measured rebate coefficient is smaller, the aggregate cumulative effect on consumption is very similar to the baseline specification. This result reinforces the view that one should be cautious in using the value of the estimated rebate coefficient to draw inference on aggregate implications for consumption, as we argued in Section 5.3.




Figure 12: Robustness analysis with respect to alternative models for the transaction cost

households is unaffected by the timing of future rises in taxes.

Earnings heterogeneity In our baseline calibration we took the view that the increase in earnings dispersion over the life-cycle is the result of the accumulation of highly persistent idiosyncratic shocks. An alternative view, also quite common in the literature (e.g., Guvenen, 2009), is that it is the result of heterogeneity in deterministic idiosyncratic earnings profiles. We re-estimated the parameters of our log household earnings process omitting the unit root component z_{it} and reinstating the household-specific slope ψ_i by matching the same set of cross-sectional moments described in Section 4. The discount factor β is always set to match median illiquid wealth, as in the baseline. Figure 11(b) compares the rebate coefficients in the baseline model (with tax reform and recession). The implications of these two models are very similar, but the heterogenous earnings profile model seems capable of generating slightly larger rebate coefficients for a wide range of κ .

Form of transaction cost In the benchmark model, the fixed cost κ is the same across households. Here we explore two alternative forms of transaction cost: a cost proportional to individual earnings to represent the opportunity cost of time (i.e., entering the budget constraint as $-\kappa y_{ij}$), and a cost paid directly in terms of utility (i.e., entering period utility as $\log(c_{ij}) - \kappa$). Given log utility, the latter specification also corresponds to a cost proportional to individual consumption (as in Romer, 1986). The key difference with the baseline is that the cost is not lump sum, but is larger for high income (or consumption) households. Figure 12(a) reports the rebate coeffi-

cients for these two models. Transaction costs of around 6% of quarterly earnings or consumption yield values for the rebate coefficient around 25%, as in the benchmark. For the median household whose annual earnings are \$41,000 (see Table 2), this cost corresponds to roughly \$600, and hence on the low end of our preferred range in the baseline specification.

Consumer credit There are two reasons why allowing households access to unsecured consumer credit may reduce rebate coefficients, and we consider each in turn.


The first conjecture is that if households can take negative positions in the liquid asset, then their Euler equation would hold with equality and the MPC out of transitory income shocks would be negligible, unless they borrowed all the way to their debt limit.⁴⁵ Since the number of households at their available credit limit is arguably small, so is the aggregate MPC out of the stimulus payment. However, in our model this group of households is not the only one with strong MPC the rebate, and as a result this argument is flawed. No arbitrage opportunities means that borrowing in the liquid asset must take place at an interest rate R^b that satisfies $R^b > R^a > R^m$. This, in turn, implies that there still exists a corner in the household problem at $m_{t+1} = 0$, and in general there will be a mass of households located at zero liquid wealth in the resulting distribution. Stimulus payments in the vicinity of \$500 are not large enough to push a substantial fraction of households off this corner into costly borrowing, and so MPCs out of the rebate for this group remain substantial.

Figure 13(a) plots the MPC out of the rebate for various borrowing rates.⁴⁶ The figure shows that with an expensive borrowing premium of 25% p.a. the results are very similar to the benchmark. Even with a modest premium of 10% p.a., the MPC remains around 0.20 and there is still significant amplification (by a factor of three) relative to the one asset model.

The second conjecture is that, even if MPCs out of the rebate are largely unaffected, the estimated rebate coefficients would be reduced: households in the control group (who are paid the rebate later) may find it worthwhile to borrow in anticipation of the fiscal stimulus payment, and consume as much as the treatment group. Clearly, for high enough rates, households are better off waiting for the check next quarter

⁴⁵Using data on credit card payments, Agarwal, Liu and Souleles (2007, Table 3f) found that households at their borrowing limits did use the tax rebates to increase spending, and the effect is statistically significant.

⁴⁶In these experiments we have set the individual borrowing limit equal to their current quarterly earnings, but the results are robust to changes in this parameter. Borrowing rates are expressed as a wedge in the after-tax-real return between the liquid return and the borrowing rate.



(a) Marginal propensity to consume out of the rebate under alternative borrowing rates. The reported wedge is relative to the return on the liquid asset

(b) Rebate coefficients under alternative borrowing rates. The reported wedge is relative to the return on the liquid asset

Figure 13: Robustness analysis with respect to credit

than pre-empting consumption at the cost of large interest payments. Figure 13(b) illustrates the strength of this effect in the model for borrowing premia in the range of 10% to 25%. We conclude that unsecured borrowing can shrink the model’s rebate coefficients only when (i) the rebate is anticipated by some, (ii) the interest rate on borrowing is modest, and (iii) every household can access credit at zero transaction cost.

Finally, it is useful to emphasize two features of the model extended to credit. First, in all versions of the model with non-zero borrowing limits, the amplification of rebate coefficients relative to the one asset model remains strong, even when the rebate coefficients are smaller than in the no borrowing economy. Second, the dampening effect on rebate coefficients is not monotonic in the cost of credit. For low enough wedges (e.g. corresponding to “cheap” borrowing through home equity loans or introductory credit card rates) many households in the model find it optimal to borrow up to the limit and then save into the high-return illiquid account.⁴⁷ As a result, even more households in this version of the model are at a kink and behave as hand-to-mouth than in the baseline model with no borrowing. For example, when the borrowing wedge $R^b - R^m$ is 5% p.a., the model generates rebate coefficients up to 37% in the \$500 – \$1,000 range.

⁴⁷This behavior, akin to an arbitrage opportunity, has a data counterpart. According to Greenspan and Kennedy (2008), 2/3 of the net proceeds of home equity loans are used for home improvements and purchase of durables.

7 Concluding remarks

In this paper we have demonstrated how, 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 anticipated temporary income changes is large – an empirical finding that is substantiated by robust quasi-experimental evidence. Going forward, we plan to expand our analysis in several directions.

In using public funds to stimulate consumption, policy makers face a broad array of options, of which across-the-board tax rebates are only one specific example. A key benefit of having a fully structural model is that one can investigate the counterfactual welfare and aggregate effects of alternative policies. Among the policies that we plan to study are more targeted stimulus payments and temporary reductions in consumption taxes –the latter recently implemented in the UK (see Blundell, 2009).

The model can be used to address the 2008 episode of fiscal stimulus payments. Under the Economic Stimulus Act of 2008, households received on average nearly \$1,000. PJSM (2011) and Broda and Parker (2011) estimate rebate coefficients for nondurable expenditures between half and 2/3 of the size of the 2001 estimates. The 2008 episode differs from the 2001 episode in three ways. First, its magnitude is roughly twice as large. Second, it is not part of any broader tax reform. Third, eligibility phases out quickly starting at \$75,000 of gross individual income. Qualitatively, the lack of a tax reform and the larger rebate size suggest that the model would generate a somewhat lower response in 2008, while the phasing out would induce a higher response as more of the hand-to-mouth households are among the low-income ones (although, as discussed in Section 5.1, we expect this effect to be weak). Moreover, the 2008 recession was much deeper than its 2001 counterpart: as explained in Section 5.1, it is *a priori* unclear whether a deeper recession is associated with a higher or lower consumption response. Overall, only a full quantitative analysis that contains all of these ingredients can shed light on what accounted for the smaller impact on consumption of the 2008 stimulus program.

In calibrating our model, we have included durables as part of illiquid wealth. Modeling durables explicitly as a third asset (with a consumption flow proportional to the stock in the period utility function) beyond the liquid and the illiquid assets would push computational complexity beyond feasibility. Moreover, it would not affect the response of nondurable consumption to the rebate, as long as utility is separable in the two

consumption goods. However, such an extension could be useful to study the aggregate implications of the policy for consumption and output in more detail. Unsurprisingly, as documented by JPS (2006) and PSJM (2011), when durable goods are included among expenditures, the response of household consumption to the rebate almost doubles.⁴⁸

Given the high-frequency OLG structure, solving a stochastic version of our model in general equilibrium (i.e., with asset returns determined endogenously and aggregate shocks) 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 (possibly with households transiting randomly between work and retirement phases). To close the model, one would interpret the illiquid asset a as productive capital with a return equal to its marginal product, whereas the return on the liquid asset m could be pinned down by a monetary policy rule (as in Ragot, 2011). In this set up, one could also endow agents with rational beliefs over the probability that the government will make stimulus transfers during a recession, as in Heathcote (2005).

⁴⁸See Guerrieri and Lorenzoni (2011) for a recent example of an incomplete-markets model with durable and nondurable consumption. As they emphasize, a crucial parameter affecting households' response to an income shock is the degree of liquidity of durable goods.

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.
- AGARWAL, S., C. LIU, AND N. SOULELES (2007): “The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium,” *mimeo*.
- 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,” Discussion paper, National Bureau of Economic Research.
- 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.
- BAI, J. (2005): “Stationary Monetary Equilibrium in a Baumol-Tobin Exchange Economy: Theory and Computation,” *Mimeo*.
- BLUNDELL, R. (2009): “Assessing the Temporary VAT Cut Policy in the UK,” *Fiscal Studies*, 30(1), 31–38.
- BRODA, C., AND J. PARKER (2008): “Estimates of the Impact of the 2008 Economic Stimulus Payments on Consumption Demand,” *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 Z. HERCOWITZ (2009): “Liquidity Constraints of the Middle Class,” *Chicago Fed Working Paper 2009-20*.
- CAMPBELL, J., AND N. MANKIW (1989): “Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence,” *NBER Macroeconomics Annual*, 4, 185–216.
- CARD, D., R. CHETTY, AND A. WEBER (2007): “Cash-On-Hand and Competing Models of Intertemporal Behavior: New Evidence from the Labor Market*,” *Quarterly Journal of Economics*, 122(4), 1511–1560.
- 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*, 1992, 61–156.
- (1997): “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis,” *Quarterly Journal of Economics*, 112(1), 1–55.
- CHATTERJEE, S., AND D. CORBAE (1992): “Endogenous Market Participation and the General Equilibrium Value of Money,” *Journal of Political Economy*, 100(3), 615–646.
- CHETTY, R. (2011): “Bounds on Elasticities with Optimization Frictions: A Synthesis of Micro and Macro Evidence on Labor Supply,” *NBER Working Paper*, 15616.
- CHETTY, R., AND A. SZEIDL (2007): “Consumption Commitments and Risk Preferences*,” *The 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,” *The 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*.
- DEATON, A. (1991): “Saving and Liquidity Constraints,” *Econometrica*, 59(5), 1221–1248.
- EROSA, A., AND G. VENTURA (2002): “On Inflation as a Regressive Consumption Tax,” *Journal of Monetary Economics*, 49(4), 761–795.
- 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*.
- FRENKEL, J., AND B. JOVANOVIĆ (1980): “On Transactions and Precautionary Demand for Money,” *The Quarterly Journal of Economics*, 95(1), 25–43.
- (1981): “Optimal International Reserves: a Stochastic Framework,” *The Economic Journal*, 91(362), 507–514.
- GREENSPAN, A., AND J. KENNEDY (2008): “Sources and Uses of Equity Extracted from Homes,” *Oxford Review of Economic Policy*, 24(1), 120.
- GUERRIERI, V., AND G. LORENZONI (2011): “Credit Crises, Precautionary Savings and the Liquidity Trap,” *Chicago-Booth Working Paper*.
- GUVENEN, F. (2009): “An Empirical Investigation of Labor Income Processes,” *Review of Economic Dynamics*, 12(1), 58–79.
- HALL, R. (2011): “The Long Slump,” Discussion paper, National Bureau of Economic Research.
- HAMILTON, D. (2008): “A Reexamination of Johnson, Parker, and Souleles 2001 Tax Rebate Estimate,” *mimeo*.
- 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.
- 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*, CBO.
- IMROHOROGU, A., AND E. PRESCOTT (1991): “Seigniorage as a Tax: A Quantitative Evaluation,” *Journal of Money Credit and Banking*, 23(3), 462–475.
- JOHNSON, D., J. PARKER, AND N. SOULELES (2006): “Household Expenditure and the Income Tax Rebates of 2001,” *The 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*.
- JOINT CENTER FOR HOUSING STUDIES (2011): “A New Decade of Growth for Remodeling: Improving America’s Housing,” *mimeo*.
- JOVANOVIC, B. (1982): “Inflation and Welfare in the Steady State,” *The Journal of Political Economy*, 90(3), 561.
- KAPLAN, G., AND G. VIOLANTE (2010): “How Much Consumption Insurance Beyond Self-Insurance?,” *American Economic Journals: Macroeconomics*, 2(4), 53–87.
- KHAN, A., AND J. THOMAS (2010): “Inflation and Interest Rates with Endogenous Market Segmentation,” *Mimeo*.
- 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.
- LAIBSON, D. (1998): “Life-Cycle Consumption and Hyperbolic Discount Functions,” *European Economic Review*, 42(3-5), 861–871.
- LI, G. (2009): “Transaction Costs and Consumption,” *Journal of Economic Dynamics and Control*, 33(6), 1263–1277.
- LUO, Y. (2008): “Consumption Dynamics under Information Processing Constraints,” *Review of Economic Dynamics*, 11(2), 366–385.

- LUSARDI, A., D. SCHNEIDER, AND P. TUFANO (2011): “Financially Fragile Households: Evidence and Implications,” *Brookings Papers on Economic Activity*, Forthcoming.
- MANKIW, N. (2000): “The Savers-Spenders Theory of Fiscal Policy,” *American Economic Review*, 90(2), 120–125.
- MCDANIEL, C. (2007): “Average Tax Rates on Consumption, Investment, Labor and Capital in the OECD: 1950–2003,” *Mimeo*, Arizona State University.
- MILLER, M., AND D. ORR (1966): “A Model of the Demand for Money by Firms,” *The Quarterly Journal of Economics*, 80(3), 413.
- MISRA, K., AND P. SURICO (2011): “Heterogeneous Responses and Aggregate Impact of the 2001 Income Tax Rebates,” *CEPR DP*, 8306.
- OH, H., AND R. REIS (2011): “Targeted Transfers and the Fiscal Response to the Great Recession,” *NBER Working Paper*, 16775.
- 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.
- PARKER, J., N. SOULELES, D. JOHNSON, AND R. MCCLELLAND (2011): “Consumer Spending and the Economic Stimulus Payments of 2008,” *NBER Working Paper*, 16684.
- 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,” *The Quarterly Journal of Economics*, 101(4), 663–685.
- ROTEMBERG, J. (1984): “A Monetary Equilibrium Model with Transactions Costs,” *The Journal of Political Economy*, 92(1), 40–58.
- 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,” *The American Economic Review*, 89(4), 947–958.
- TAX FOUNDATION (2011): “Property Taxes on Owner Occupied Housing 2004-2009 by State,” <http://www.taxfoundation.org/research/show/1913.html>.
- TELYUKOVA, I. (2011): “Household Need for Liquidity and the Credit Card Debt Puzzle,” *mimeo*.
- THALER, R. (1990): “Anomalies: Saving, Fungibility, and Mental Accounts,” *The Journal of Economic Perspectives*, 4(1), 193–205.

A1 Appendix for Online Publication

This appendix contains a detailed description of the algorithm used for computation of the model.

A1.1 Preliminaries

It is useful to define the following objects:

- The vector ζ_j summarizing the three components of earnings, i.e., $\zeta_j = (\chi_j, \alpha, z_j) \in \mathcal{L}_j$ with conditional c.d.f. $F_j(\zeta_j | \zeta_{j-1})$. Therefore individual earnings can be expressed as $y_j(\zeta_j)$.
- In the main text we defined a linear consumption tax τ^c and a tax function $\mathcal{T}(y_j, a_j, m_j)$. In what follows, we exploit that this tax function is separable between the progressive tax on earnings $T(y_j)$ and the linear taxes on the two assets (τ^m, τ^a) .
- Recall that the price of the illiquid asset q^a in the main text is interpreted as an after-tax price. Holdings of the illiquid asset cannot be negative, or $a_{j+1} > 0$. The after-tax price of the liquid asset $q^m(m_j)$ in the main text is allowed to depend on m_j to reflect the possibility of a spread between the price of loans and that of deposits in the liquid asset, i.e.

$$q^m(m_{j+1}) = \begin{cases} q^{m+} & \text{for } m_{j+1} \geq 0 \\ q^{m-} & \text{for } m_{j+1} < 0. \end{cases}$$

with $q^{m-} < q^{m+}$. As a result, there is a corner in all decision rules at $m_{j+1} = 0$. The other corner is at the ad-hoc borrowing limit $\underline{m}_{j+1}(\zeta_j)$, where we allow dependence from earnings $y_j(\zeta_j)$.

- x_j^0 is total cash-on-hand available for consuming and saving, for an agent who is not adjusting:

$$x_j^0(m_j, a_j, \zeta_j) \equiv m_j + y_j(\zeta_j) - T(y_j(\zeta_j)) + reb_j,$$

where reb_j is equal to 0 unless windfall income from a tax rebate is received in period j .

- x_j^1 is total cash-on-hand available for consuming and saving, for an agent who is adjusting, before paying the transaction cost κ :

$$\begin{aligned} x_j^1(m_j, a_j, \zeta_j) &\equiv m_j + a_j + y_j(\zeta_j) - T(y_j(\zeta_j)) + reb_j \\ &= x_j^0(m_j, a_j, \zeta_j) + a_j. \end{aligned}$$

- $V_j^1(x_j, \zeta_j)$ is the value function if the agent accesses the illiquid asset, and $c_j^1(x_j, \zeta_j)$ is the associated consumption policy function.

- $V_j^0(x_j, a_j, \zeta_j)$ is the value function if the agent does not access the illiquid asset, and $c_j^0(x_j, a_j, \zeta_j)$ is the associated consumption policy function.
- $\Omega_j(m_j, a_j, \zeta_{j-1})$ is the expected value function at age j , where the expectation is taken over the current period j earnings shocks. It is a function of the current period holdings of the two assets (since these are predetermined in period $j - 1$) and of the previous period's earnings.

$$\Omega_j(m_j, a_j, \zeta_{j-1}) = \int_{\mathcal{L}_j} \max \{V_j^1(x_j^1(m_j, a_j, \zeta_j), \zeta_j), V_j^0(x_j^0(m_j, a_j, \zeta_j), a_j, \zeta_j)\} dF(\zeta_j | \zeta_{j-1})$$

Clearly, during retirement, the function Ω_j does not depend on ζ_{j-1} .

- A new operator, $\tilde{\max}$. This operator chooses between two objects based on which of the corresponding value functions is higher. For example $\tilde{\max}\{f^1, f^0\}$ chooses function f^1 if $V^1 > V^0$, and function f^0 if $V^0 \geq V^1$ at the corresponding point in the state space.

A1.2 Decision Problems

No-adjust case

$$V_j^0(x_j, a_j, \zeta_j) = \max_{c_j, s_j} \{u(c_j) + \beta \Omega_{j+1}(m_{j+1}, a_{j+1}, \zeta_j)\}$$

subject to

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

Adjust case

$$V_j^1(x_j, \zeta_j) = \max_{c_j, s_j, a_{j+1}} \{u(c_j) + \beta \Omega_{j+1}(m_{j+1}, a_{j+1}, \zeta_j)\}$$

subject to

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

A1.3 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 potential solution.

No-adjust case When agents do not adjust, there is one FONC:

$$\frac{1}{1 + \tau^c} u'(c_j) = \begin{cases} \frac{\beta}{q^{m_+}} \Omega_{m,j+1}(m_{j+1}, a_{j+1}, \zeta_j) & \text{if } m_{j+1} > 0 \\ \in \left[\frac{1}{q^{m_+}}, \frac{1}{q^{m_-}} \right] \times \beta \Omega_{m,j+1}(m_{j+1}, a_{j+1}, \zeta_j) & \text{if } m_{j+1} = 0 \\ \frac{\beta}{q^{m_-}} \Omega_{m,j+1}(m_{j+1}, a_{j+1}, \zeta_j) & \text{if } m_{j+1} \in (\underline{m}_{j+1}(\zeta_j), 0) \\ > \frac{\beta}{q^{m_-}} \Omega_{m,j+1}(m_{j+1}, a_{j+1}, \zeta_j) & \text{if } m_{j+1} = \underline{m}_{j+1}(\zeta_j) \end{cases} \quad (\text{A1})$$

Adjust case When agents adjust, there are two FONCs, one for each asset:

$$\frac{1}{1 + \tau^c} u'(c_j) = \begin{cases} \frac{\beta}{q^{m_+}} \Omega_{m,j+1}(m_{j+1}, a_{j+1}, \zeta_j) & \text{if } m_{j+1} > 0 \\ \in \left[\frac{1}{q^{m_+}}, \frac{1}{q^{m_-}} \right] \times \beta \Omega_{m,j+1}(m_{j+1}, a_{j+1}, \zeta_j) & \text{if } m_{j+1} = 0 \\ \frac{\beta}{q^{m_-}} \Omega_{m,j+1}(m_{j+1}, a_{j+1}, \zeta_j) & \text{if } m_{j+1} \in (\underline{m}_{j+1}(\zeta_j), 0) \\ > \frac{\beta}{q^{m_-}} \Omega_{m,j+1}(m_{j+1}, a_{j+1}, \zeta_j) & \text{if } m_{j+1} = \underline{m}_{j+1}(\zeta_j) \end{cases} \quad (\text{A2})$$

and

$$\frac{1}{1 + \tau^c} u'(c_j) \geq \frac{\beta}{q^a} \Omega_{a,j+1}(m_{j+1}, a_{j+1}, \zeta_j) \quad (\text{A3})$$

where the strict inequality for the second FONC holds when the non-negativity constraint on illiquid assets binds.

A1.4 Envelope Conditions

We now compute the required partial derivatives of the expected value function Ω_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 of points of discontinuity.

For the derivative with respect to liquid assets we have

$$\Omega_{m,j}(m_j, a_j, \zeta_{j-1}) = \mathbb{E} \left[\max \left\{ V_{m,j}^1(x_j^1(m_j, a_j, \zeta_j), \zeta_j), V_{m,j}^0(x_j^0(m_j, a_j, \zeta_j), a_j, \zeta_j) \right\} \right]$$

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

$$\begin{aligned} V_{m,j}^1(x_j^1(m_j, a_j, \zeta_j), \zeta_j) &= V_{x,j}^1(x_j^1(m_j, a_j, \zeta_j), \zeta_j) \\ V_{m,j}^0(x_j^0(m_j, a_j, \zeta_j), \zeta_j) &= V_{x,j}^0(x_j^0(m_j, a_j, \zeta_j), \zeta_j). \end{aligned}$$

For the partial derivative with respect to illiquid assets we have

$$\Omega_{a,j}(m_j, a_j, \zeta_{j-1}) = \mathbb{E} [\text{m}\ddot{\text{a}}\text{x} \{V_{x,j}^1(x_j^1(m_j, a_j, \zeta_j), \zeta_j), V_{a,j}^0(x_j^0(m_j, a_j, \zeta_j), a_j, \zeta_j)\}].$$

Next, we compute the partial derivatives of the choice-specific value functions (V_j^0, V_j^1) . For the adjust case, it is given by

$$V_{x,j}^1(x_j, \zeta_j) = \frac{1}{1 + \tau^c} u'(c^1(x_j, \zeta_j)).$$

For the no-adjust case, they are given by

$$\begin{aligned} V_{a,j}^0(x_j, a_j, \zeta_j) &= \frac{\beta}{q^a} \Omega_{a,j+1}(m_{j+1}, a_{j+1}, \zeta_j) \\ V_{x,j}^0(x_j, a_j, \zeta_j) &= \frac{1}{1 + \tau^c} u'(c^0(x_j, a_j, \zeta_j)). \end{aligned}$$

In the RHS of the first line, m_{j+1} should be interpreted as the optimal decision rule for savings at the point (x_j, a_j, ζ_j) .

To make progress in constructing these objects recursively, we define the following objects:

$$\begin{aligned} h_j(m_j, a_j, \zeta_{j-1}) &\equiv \Omega_{a,j}(m_j, a_j, \zeta_{j-1})(1 + \tau^c) \\ g_j(x_j, a_j, \zeta_j) &\equiv V_{a,j}^0(x_j, a_j, \zeta_j)(1 + \tau^c) \\ \mu_j(m_j, a_j, \zeta_{j-1}) &\equiv \Omega_{m,j}(m_j, a_j, \zeta_{j-1})(1 + \tau^c). \end{aligned}$$

By substituting into the envelope conditions, we obtain

$$\mu_j(m_j, a_j, \zeta_{j-1}) = \mathbb{E} [\text{m}\ddot{\text{a}}\text{x} \{u'(c^1(x_j^1(m_j, a_j, \zeta_j), \zeta_j)), u'(c^0(x_j^0(m_j, a_j, \zeta_j), a_j, \zeta_j))\}]$$

together with the recursion

$$\begin{aligned} g_j(x_j, a_j, \zeta_j) &= \frac{\beta}{q^a} h_{j+1}(m_{j+1}, a_{j+1}, \zeta_j) \\ h_j(m_j, a_j, \zeta_{j-1}) &= \mathbb{E} [\text{m}\ddot{\text{a}}\text{x} \{u'(c^1(x_j, \zeta_j)), g_j(x_j^0(m_j, a_j, \zeta_j), a_j, \zeta_j)\}]. \end{aligned}$$

Note that by defining objects in this way we can construct $h_j \equiv \Omega_a$ recursively, which avoids the need to compute numerical derivatives of the expected value function. This allows us to

restate the optimality conditions in terms of functions that can be computed directly from the optimal consumption choices in later periods, as described in the following subsection.

A1.4.1 Euler Equations

Armed with these recursions for the expected marginal values of illiquid assets (h_j) and total assets (μ_j), we can finally derive the Euler Equations that need to be solved.

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

$$u'(c_j) = \begin{cases} \frac{\beta}{q^{m_+}} \mu_{j+1}(m_{j+1}, a_{j+1}, \zeta_j) & \text{if } m_{j+1} > 0 \\ \in \left[\frac{1}{q^{m_+}}, \frac{1}{q^{m_-}} \right] \times \beta \mu_{j+1}(m_{j+1}, a_{j+1}, \zeta_j) & \text{if } m_{j+1} = 0 \\ \frac{\beta}{q^{m_-}} \mu_{j+1}(m_{j+1}, a_{j+1}, \zeta_j) & \text{if } m_{j+1} \in (\underline{m}_{j+1}(y_j), 0) \\ > \frac{\beta}{q^{m_-}} \mu_{j+1}(m_{j+1}, a_{j+1}, \zeta_j) & \text{if } m_{j+1} = \underline{m}_{j+1}(y_j) \end{cases}$$

Adjust case For the adjust case, we have the two Euler equation:

$$u'(c_j) = \begin{cases} \frac{\beta}{q^{m_+}} \mu_{j+1}(m_{j+1}, a_{j+1}, \zeta_j) & \text{if } m_{j+1} > 0 \\ \in \left[\frac{1}{q^{m_+}}, \frac{1}{q^{m_-}} \right] \times \beta \mu_{j+1}(m_{j+1}, a_{j+1}, \zeta_j) & \text{if } m_{j+1} = 0 \\ \frac{\beta}{q^{m_-}} \mu_{j+1}(m_{j+1}, a_{j+1}, \zeta_j) & \text{if } m_{j+1} \in (\underline{m}_{j+1}(y_j), 0) \\ > \frac{\beta}{q^{m_-}} \mu_{j+1}(m_{j+1}, a_{j+1}, \zeta_j) & \text{if } m_{j+1} = \underline{m}_{j+1}(y_j) \end{cases}$$

and

$$\begin{aligned} u'(c_j) &= \frac{\beta}{q^a} h_{j+1}(z_{j+1}, a_{j+1}, \zeta_j) & \text{if } a_{j+1} > 0 \\ u'(c_j) &> \frac{\beta}{q^a} h_{j+1}(z_{j+1}, a_{j+1}, \zeta_j) & \text{if } a_{j+1} = 0. \end{aligned}$$

A1.4.2 Recursive algorithm

The model is computed by recursively using these Euler Equations to solve for the policy functions back from the last period in the life-cycle $j = J$, when the agent faces a simple static problem. At each point in the state space, we search for multiple solutions to the Euler equations, compute the associated value functions and select the solution with the highest value. In the case of no-adjustment, we use a version of the method of endogenous

grid points. To find multiple solutions, at each point in the state space we divide the choice set into equally sized partitions and search for local solutions in each partition. We verify the second-order conditions at each local optimum. We check that increasing the number of partitions does not lead to an increase in the number of local optima. We explicitly allow for the possibility of solutions at each corner, and compute the associated value function at these points.

A1.5 Bounds, Grids and Interpolation over Asset Space

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

A1.5.1 (m_j, a_j) asset space

Let the lower bound for liquid assets, m_j , be given by \underline{m}_j . We verify that \underline{m}_j never violates the natural borrowing limit. 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 along 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} .

A1.5.2 (x_j, a_j) asset space

For adjusting agents, the lowest possible value of x_j^1 is

$$\underline{x}_j^1 = \underline{m}_j + \min_{y_j} \{y_j - T(y_j)\}$$

and the highest possible value is

$$X_j^1 = M_j + A_j + \max_{y_j} \{y_j - T(y_j) + reb_j\}.$$

For non-adjusting agents, we have to construct a 2-dimensional state space for assets. Values

of a_j lie in $[0, A_j]$. Over the x_j dimension, we have

$$\begin{aligned}\underline{x}_j^0 &= \underline{m}_j + \min_{y_j} \{y_j - T(y_j)\} \\ X_j^0 &= M_j + \max_{y_j} \{y_j - T(y_j) + reb_j\}.\end{aligned}$$

A1.5.3 Grid sizes

In the models without borrowing, we use 30 points each in the grids for a_j , m_j and x_j^0 , and 50 points in the grid for x_j^1 . 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^0 and x_j^1 . 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 $k = 0.4$.

A1.5.4 Interpolation

We use linear over x_{j+1}^1 and bilinear interpolation over (x_{j+1}^0, a_{j+1}) . For the (m_{j+1}, a_{j+1}) space, we use a variant of bilinear interpolation that interpolates 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 $j + 1$.