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, IFS and NBER
glv2@nyu.edu

Abstract

Fiscal stimulus payments (e.g., tax rebates) are a key instrument used by governments to stabilize aggregate economic activity and moderate the effects of recessions on household welfare. A wide body of empirical evidence based on randomized experiments concludes that nondurable household consumption rises by 20-40 percent in the quarter in which the rebates 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 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 holding sizeable quantities of illiquid assets. We find evidence for the presence of such households in Survey of Consumer Finances data. When the model is parameterized to match a number of targets - in particular the joint cross-sectional distribution of liquid and illiquid wealth - it is able to generate responses to fiscal stimulus payments of the observed order of magnitude.

Keywords: Consumption, Fiscal Stimulus Payments, Hand-to-Mouth, Liquidity.

JEL Classification: D31, D91, E21

---

1 Introduction

Fiscal stimulus payments (i.e., transfers to households such as tax rebates) are among the most common policy instruments used by governments for alleviating households' economic hardship. This type of fiscal intervention was authorized by the US Congress in the downturns of 1975, 2001, 2008, and 2009.[1] Households received one-off payments between \$300 and \$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, the latter recently renewed for another year.

On the empirical side there has been substantial progress in measuring the size of household-level consumption responses to the tax rebate episodes of 2001 and 2008. Johnson, Parker, and Souleles (2006, hereafter JPS), Parker, Souleles, Johnson, and McLelland 2011, hereafter PSJM), and Misra and Surico (2011) make clever use of the randomized timing of the receipt of rebate checks to estimate the effects on consumption expenditures from the Consumer Expenditure Survey (CEX). Shapiro and Slemrod (2003a, 2003b, 2009, hereafter SS) substantiate these findings with a detailed qualitative survey on how consumers use their rebate checks. This collective body of evidence –which we discuss below– suggests that between 20 and 40 percent of the rebate is spent by households on nondurables in the same quarter that it is received.

Yet while there is sound empirical evidence on the effect of stimulus payments on household consumption and saving behavior, there are virtually no studies that rationalize these episodes within a structural, dynamic, forward-looking macroeconomic model of household optimization. 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. The government has a number of alternative ways to use the same resources to moderate the impact of adverse aggregate shocks. Knowing the determinants of how consumers respond to stimulus payments helps in choosing among the policy options and in understanding whether the same fiscal instrument can be expected to have more or less bite in different macroeconomic conditions.[2]

In this paper we address this gap in the literature by developing a structural macroeconomic model that can be used to analyze the impact of tax rebates, and other similar measures of fiscal stimulus, on household consumption. From the perspective of 'off-the-shelf' consumption theory (the rational expectations, life-cycle, buffer-stock model), this is a challenging task: the theory predicts that the

---

[1]In the context of the most recent recession, Oh and Reis (2011) document that the large fiscal expansion of 2007-09 consisted primarily of growing social assistance, as opposed to government purchases, and that half of this increase was due to discretionary fiscal stimulus transfers.

[2]JPS (2006) conclude their empirical analysis of the 2001 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).

marginal propensity to consume (MPC) out of small, anticipated, transitory income fluctuations, such as the tax rebates, should be negligible. Plausible simulations of this class of models suggest a MPC not higher than 5% (e.g., Kaplan and Violante, 2010).

As emphasized by Modigliani and Steindel (1977) in their seminal study of the 1975 tax rebate, the data could deviate from standard theory so sizably only if many households are extremely impatient ('myopic') or borrowing constrained. Myopic households consume a large fraction of the rebate in spite of its transitory nature, and would have accumulated too little wealth to be able to respond to the announcement of the rebate before the receipt of the transfer. Constrained agents are similarly unable to respond to the news of the rebate, but will have a high propensity to consume when it is received. However, in plausibly calibrated life-cycle, heterogeneous-agents, incomplete markets models there are no myopic households, since the degree of patience is pinned down by median (or average) net worth in the data, which is large as a fraction of income. Moreover, the number of constrained agents is small since it is tied to the number of households with zero or negative net worth in the data (around 10%). This low fraction of households with negative net worth also severely limits the plausible amount of heterogeneity in discount factors one could use to generate myopic households.

We develop a quantitative framework that 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 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.[3] The model is parameterized to replicate a number of macroeconomic, life-cycle, and cross-sectional targets, of which the most important are the distributions of liquid and illiquid wealth across households.

To understand why our theory can be quantitatively consistent with the empirical evidence, reconsider the Modigliani-Steindel observations. First, in our two-asset model, in the periods between deposits or withdrawals from the illiquid account, consumption behavior is dictated by an Euler equation that features a very low effective discount factor (since the real return on the liquid asset is zero, or even negative). Hence, in the 'short-run' households are rather impatient. They prefer consuming over saving through either an inefficient storage technology (liquid wealth) or an efficient one that would

---

[3]Macroeconomic applications of the Baumol-Tobin model are mostly limited to monetary and financial issues. Within complete markets models, Alvarez, Atkeson and Kehoe (2002) and Kahn and Thomas (2009) study how transaction costs lead to endogenous asset market segmentation and real effects of monetary injections. Romer (1986) and Chatterjee and Corbae (1992) studied deterministic, overlapping-generations version of the Baumol-Tobin model that, as we explain below, bear some resemblance to our model during the retirement phase. Within Bewley-style two-asset economies with transaction costs, Aiyagari and Gertler (1991) focuses on the equity premium and the low frequency of trading equities; Imrohoroglu and Prescott (1991) introduce a fixed per-period participation cost of bond holdings to support equilibria where money has value; more recently, Bai (2005) studies welfare effects of inflation.

require paying a transaction cost (illiquid wealth). This disconnect between short- and long-run consumption dynamics is a novel feature of our environment that also helps in addressing a number of other facts about household behavior that we describe in Section 6.

Second, and more importantly, our model features a large number or "wealthy hand-to-mouth" households. These are households who, despite holding sizeable amounts of illiquid wealth, would like to consume more than their disposable income, i.e., they are constrained. They are predominantly working-age households who have optimally chosen to hold the majority of retirement saving in illiquid form because of the high return. Such households do not respond to the news of the rebate and have a high MPC when they receive it. It is through these households that the model can generate much higher average consumption responses to the tax rebate compared to the one-asset model. Indeed, when the data on net worth are interpreted through the lenses of the one-asset model, such households would not look like they are constrained.

When we replicate, by simulation, the randomized experiment associated with the tax rebate of 2001 within our structural model, we find contemporaneous consumption responses to the tax rebate of sizes comparable to those estimated by JPS (2006), i.e., around 25%. Instrumental in obtaining large MPC is incorporating in the model two key features of the macroeconomic environment of 2001: the Bush tax reform and the recession. By raising future permanent income relative to current income, both components contribute to exacerbate liquidity constraints at the time of the rebate.

Central to our economic analysis is the insight of Cochrane (1989): the welfare costs of failing to fully smooth consumption in the wake of small and temporary income fluctuations - like tax rebates - are themselves small.[4] Our mechanism is consistent with this insight: many households choose to consume their rebate precisely because saving it in illiquid wealth would entail a transaction cost that is larger than the welfare gain from smoothing. Households do not respond to the news about the rebate for the same reason.

There are two additional readings of Cochrane's important remark. The first one is that such 'small deviations' are unimportant and do not need to be studied by macroeconomists. We strongly disagree with this view. The JPS (2006) estimates feature prominently in the reports prepared by the Congressional Budget Office (CBO) and Council of Economic Advisors (CEA) that document and forecast the macroeconomic effects of fiscal stimulus.[5] Hence, understanding the behavior underlying the micro-evidence is crucial for a correct interpretation of CBO and CEA analyses. Moreover, al-

---

[4]Browning and Crossley (2001) make similar welfare calculations in the context of the evidence we are interested in. Chetty (2011) makes a similar argument in the context of labor supply behavior.

[5]Congressional Budget Office (2009); Council of Economic Advisors (2010).

though individual households are nearly indifferent about how they use their stimulus payments, in the aggregate this is not so: the implications can be substantial given the large amount of resources at stake.

The second implication of Cochrane's remark is that small tax rebates may not be a powerful validating data source for choosing between competing structural models of consumption behavior. We take this point very seriously: we describe a number of additional sources of data which are not subject to this critique, and show that our model is consistent with this evidence. Examples that we consider include the distribution of household portfolios, consumption responses to large, anticipated income fluctuations, evidence on the concentration of consumption responses across the population, and episodes of fiscal stimulus payment taking place in different macroeconomic conditions and under a different design.

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 describe our baseline 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 calibration and in Section 5 we provide a quantitative analysis of the 2001 rebates. In Section 6 we 'validate' the model using a number of additional sources of data. In Section 7 we extend our model to allow for an alternative characterization of earnings heterogeneity. Finally, in Section 8 we describe various possible extensions and uses of our framework, and conclude.

2 Measuring consumption responses to tax rebates

Direct cash transfers from governments to households are an increasingly commonly used policy too for stimulating consumption in the face of an economic downturn. The largest and most recent examples of this type of policy intervention were those associated with the 2008-09 recession.[6] For the majority of the paper, however, we focus our attention on the tax rebates of 2001. This is 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 about the nature of 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 most recent recession.

---

[6]The Economic Stimulus Act (2008) provided most households with payments of $300-$600 per adult and $300 per child, for a total outlay of $79 billion. The American Recovery and Reinvestment Act (2009) involved a refundable tax credit of up to $400 per adult.

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. However, the bill was ultimately renewed in December 2010 for a further two years.

The tax rebate   The reduction from 15 to 10 percent for the lowest bracket was deemed effective in January 2001 and would have implied a tax refund that would be received by households only in April 2002. In order to make this part of the reform highly visible during calendar year 2001, the Administration decided to pay an "advance refund" (informally called the tax rebate). The Treasury calculated that 92 million taxpayers received a rebate check, with 72 million receiving the maximum amount, ($600, or 5 percent of $12,000, for married couples). Overall, the payments amounted to $38B, i.e., just below 0.4% of 2001 GDP. The vast majority of the checks were mailed between the end of July and the end of September 2001, based on social security number (SSN) in a sequence which was announced by the Treasury. This sequence featured in the news in June, and the Treasury mailed every taxpayer a letter informing them about when they could expect to receive their check.

From the point of view of economic theory, the tax rebate of 2001 has three salient characteristics: 1) it is essentially a lump-sum, since almost every household received $300 per adult; 2) 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 in the news or from friends/family who had already received it; and 3) it has the feature of a randomized experiment because the checks were mailed according to the last two digits of SSN, which are uncorrelated with any individual characteristic. Whether it was perceived as permanent or transitory is harder to tell: first, it depends on what beliefs households held with respect to the sunset; second, the rebate was more generous that its long-run counterpart in the tax reform —according to Kiefer et al. (2002), the reduction in the effective marginal tax rate below $12,000 was roughly 3.3 percent, or $400. As a result, there was a sizeable temporary component in the rebate.

Empirical evidence   JPS (2006) use questions expressly added to the Consumer Expenditure Survey (CEX) that ask about the timing and the amount of the rebate check. Among the various

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

	Strictly Nondurable	Nondurable
2001 Tax Rebates		
JPS 2006, 2SLS ($N = 13,066$)	0.202 (0.112)	0.375 (0.136)
MS 2011, 2SLS ($N = 13,066$)		0.385 (0.120)
MS 2011, IVQR ($N = 13,066$)		0.244 (0.057)
2003 Child Tax Credit		
JPS 2009, 2SLS ($N = 15,069$)	0.020 (0.098)	0.232 (0.124)
2008 Fiscal Stimulus Payments		
PSJM 2011, 2SLS ($N = 11,239$)	0.207 (0.087)	0.252 (0.103)

Notes: Strictly non-durable consumption excludes expenditures on apparel, health care and reading materials. JPS 2006 and 2009 refer to Johnson, Parker and Souleles (2006, 2009); MS refers to Misra and Surico (2011), PJSM refers to Parker, Souleles, Johnson and McLelland (2011).

specifications estimated by JPS (2006), we will focus on the most natural one:

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

where Δc_{it} is the change in nondurable expenditures of household i in quarter t, month_s are set of time dummies, $X_{i,t-1}$ 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' – captures the fraction of the rebate spent by households who received the check. Identification comes from randomization in the *timing* of the receipt of rebate checks across households. Because households who do not receive the check at date t (the control group) may know they will receive it in the future or have already received it in the past, β_2 should be interpreted as the marginal propensity to consume (MPC) out of the rebate, net of the consumption responses for households in the control group which, in general, are not zero.

Table 1 reproduces estimates from JPS (2006), as well as estimates from other studies that are based on the same empirical strategy. The key finding of JPS (2006) is that households spend between 20 and 40 percent of the rebate on non-durable expenditures in the quarter they received it. Similar, and more precise estimates, are obtained by PSJM (2011) in the context of the 2008 episode. JPS (2009) found slightly smaller consumption responses for the 2003 child tax credit episode, which, unlike the other two episodes, took place during a period of aggregate economic growth.

Heterogeneity in consumption responses Misra and Surico (2011) use a quantile regression framework to obtain non-parametric estimates of the distribution of rebate coefficients. They find

6

evidence of substantial heterogeneity that is bi-modal in nature: around 45% of the sample saved the entire amount of the rebate, while around 15% spent the entire amount.[7] A complementary source of evidence comes from SS (2003a, 2003b) who added an ad-hoc module to the Michigan Survey of Consumers to assess the impact of the rebate. This survey asked households what 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.[8]

This collective body of evidence is difficult to reconcile with the rational expectations, life-cycle permanent income hypothesis (LCPIH), for two reasons. First, the LCPIH predicts that the MPC out of a small transitory income shock, such as the tax rebate, should roughly equal to the annuity value of the additional income, hence it should be very small except for elderly households with a short horizon. Second, because the news of the tax rebate was revealed to both the treatment and control groups at the same time, the two groups should respond similarly and hence the coefficient β_2 should be zero.

For households who are borrowing constrained, the above argument does not hold. Constrained agents do not respond to news about the rebate, and may consume a large fraction of the rebate when it is received. Thus, in principle, a standard life-cycle model with borrowing constraints (Deaton, 1991; Rios-Rull, 1995; Huggett 1996; Carroll, 1997) could account the rebate evidence. However, as we show in Section 5, plausibly calibrated life-cycle incomplete-markets economies do not feature nearly enough constrained households to generate rebate coefficients of the observed size. Such model economies are disciplined by the cross-sectional data on household net worth, for which very few households have zero or negative holdings –an upper bound on the number of constrained agents. In the next section, we modify this class of models in such a way as to expand the notion of constrained households to also include households with no liquid wealth, but possibly positive illiquid assets (and hence, positive net worth).

---

[7] JPS (2006) and PSJM (2011) test for differences in rebate coefficients across the population, by dividing the population into groups based on self-reported income and liquid wealth. They find that households in the bottom third of both distributions had substantially larger consumption responses. These effects are imprecisely estimates, 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. It is likely that respondents are a highly selected sub-sample.

[8] PSJM (2011) document 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.

3 A model with liquid and illiquid assets

We now describe our baseline framework. The model integrates the Baumol-Tobin inventory-theoretic model of money demand into an incomplete-markets life cycle economy (Aiyagari, 1994; Huggett, 1996). In this section we describe the model in steady-state, outside of the 2001 rebate environment. We use a series of examples to highlight the economic mechanisms at work. Then in Section 4, we introduce the additional features needed to model the rebate experiment.

3.1 Model description

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

Preferences   Households have time-separable, expected utility given by

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

where c_{ij} is consumption of non-durables for household i at age j.

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

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

where χ_j is a deterministic age profile that is common across all households; α_i is an individual-specific fixed effect; ψ_i is the slope of an individual-specific deterministic linear age-earnings profile, and z_{ij} is a stochastic idiosyncratic component that follows a discrete Markov process Γ_j^z.

Financial assets   There are two assets: a liquid asset (or cash) m and an illiquid asset a. 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 $R^a > R^m$ and note that since these are real returns, we may have $R^m < 1$. When the individual wants to make deposits into, or withdrawals from, the illiquid account, she must pay a transaction cost κ.[9] This creates 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 unsecured credit (e.g., credit card debt). The interest rate on borrowing

---

[9]It 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.

is denoted by $1/\bar{q}^m$ and we define the function $q^m\left(m_{j+1}\right)$ to encompass both the case $m_{j+1} \geq 0$ and $m_{j+1} < 0$. To rule out arbitrage across assets, we impose that $\bar{q}^m < q^a$.

Government   Government expenditures G are not valued by households. Retirees receive social security benefits $p\left(Y_{iJ^w}\right)$ which are a function of average gross lifetime earnings, $Y_{iJ^w} = \frac{1}{J^w}\sum_{j=1}^{J^w} y_{ij}$. The government levies proportional taxes on consumption expenditures (τ^c) and on asset income (τ^a, τ^m), a payroll tax τ^{ss} and a progressive tax on labor income $\tau^y\left(y\right)$. The combined income tax is:

$$\mathcal{T}\left(y, a, m\right) = \left[\tau^y\left(y\right) + \tau^{SS}\right] y + \tau^a(1 - q^a)a + \tau^m(1 - q^m)m \cdot \mathbf{1}_{m>0} \tag{4}$$

where the indicator function means that there is no deduction for interest paid on unsecured borrowing. For retirees, the same tax function applies with $p\left(Y\right)$ in place of y. Finally, we let the government issue one-period debt B at price q^a.

Household problem   We use a recursive formulation of the problem. Let $\mathbf{s}_j = \left(m_j, a_j, Y_j, \alpha, \psi, z_j\right)$ be the vector of individual states at age j. The value function of a household at age j is $V_j\left(\mathbf{s}_j\right) = \max\left\{V_j^0\left(\mathbf{s}_j\right), V_j^1\left(\mathbf{s}_j\right)\right\}$, where $V_j^0\left(\mathbf{s}_j\right)$ and $V_j^1\left(\mathbf{s}_j\right)$ 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 period endowment, but before consuming.[10]

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

$$V_j^0\left(\mathbf{s}_j\right) = \max_{c_j, m_{j+1}} \quad u\left(c_j\right) + \beta\mathbb{E}_j\left[V_{j+1}\left(\mathbf{s}_{j+1}\right)\right] \tag{5}$$

$$\text{subject to} \quad :$$

$$q_j^m\left(m_{j+1}\right)m_{j+1} + \left(1 + \tau^c\right)c_j = y_j - \mathcal{T}\left(y_j, a_j, m_j\right) + m_j$$

$$q^a a_{j+1} = a_j$$

$$a_{j+1} \geq 0$$

$$y_{ij} = \exp\left(\chi_j + \alpha_i + \psi_i j + z_{ij}\right), \quad \text{and} \quad \Gamma_j^z\left(z_{j+1}, z_j\right)$$

$$Y_{j+1} = \left(jY_j + y_j\right)/\left(j + 1\right).$$

---

[10]It is important to note that the timing of adjustment decisions *does not* imply a cash-in-advance constraint. Households realize all income shocks *before* making their adjustment and consumption decisions.

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:

$$V_j^1(\mathbf{s}_j) = \max_{c_j, m_{j+1}, a_{j+1}} u(c_j) + \beta \mathbb{E}_j[V_{j+1}(\mathbf{s}_{j+1})] \tag{6}$$

$$\text{subject to} \quad :$$

$$q_j^m(m_{j+1})m_{j+1} + q^a a_{j+1} + (1+\tau^c)c_j = y_j - \mathcal{T}(y_j, a_j, m_j) + m_j + a_j - \kappa$$

$$a_{j+1} \geq 0$$

$$y_{ij} = \exp(\chi_j + \alpha_i + \psi_i j + z_{ij}), \quad \text{and} \quad \Gamma_j^z(z_{j+1}, z_j)$$

$$Y_{j+1} = (jY_j + y_j)/(j+1)$$

The problem for the retired household, for age $j > J^w$, is similar, with benefits $p(Y_{J^w})$ in place of earnings y_j (and hence α, ψ and z_j omitted from the state vector), and with $Y_j = Y_{J^w}$ for all $j > J^w$.

Equilibrium   We take the returns on the two assets as exogenous since our focus is on understanding consumption dynamics given realistic values of returns, rather than on the determination of returns.[11] Given $(q^a, q^m(m))$, households optimize. Given G, $p(Y_{J^w})$ and $\mathcal{T}(\cdot)$, we let B adjust so that the intertemporal government budget constraint

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

is balanced, where μ_j is the measure of households of age j.

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

Two Euler equations   Consumption and portfolio decisions can be 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 clearly see the economics, consider a deterministic version of the model without income heterogeneity, borrowing, and taxes. In periods where the working household is unconstrained and does not adjust:

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

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 income, i.e., $c_j = y_j$.

---

[11]One simple way to close the model and determine returns in equilibrium would be to interpret the illiquid asset a as productive capital, as in Aiyagari (1994), and the liquid asset m as money. The return on capital would equal its marginal product and the return on money would then be pinned down by a policy rule for either the nominal interest

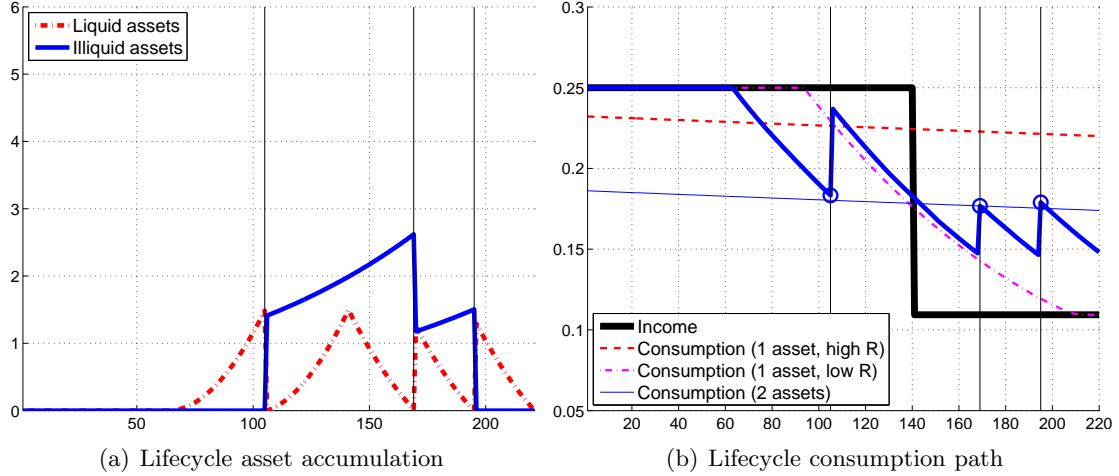(a) Lifecycle asset accumulation (b) Lifecycle consumption path

Figure 1: Example of lifecycle asset and consumption paths with low return differential

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 d periods apart, consumption dynamics are dictated by

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

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).[12] Consumption in excess of pension income is financed by making periodic withdrawals from the illiquid account. Between each withdrawal, the household runs down its cash-holdings and consumption falls according to (EE-SR). The withdrawals are timed to coincide with the period where cash is exhausted. Across withdrawals, equation (EE-LR) holds.

Figure 1 shows consumption and wealth dynamics in an example where an agent receives a constant endowment while working and a fraction of it when retired.[13] 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 on the illiquid account grows at rate $(1/q^a - 1)$. Panel (b) shows her associated income and consumption paths. In the same panel,

---

rate, money supply or inflation.

[12]In particular, in this example our problem during retirement would coincide with Romer (1986) with no discounting, log utility and the transaction cost expressed in utility terms.

[13]For clarity, we impose a very large transaction cost.

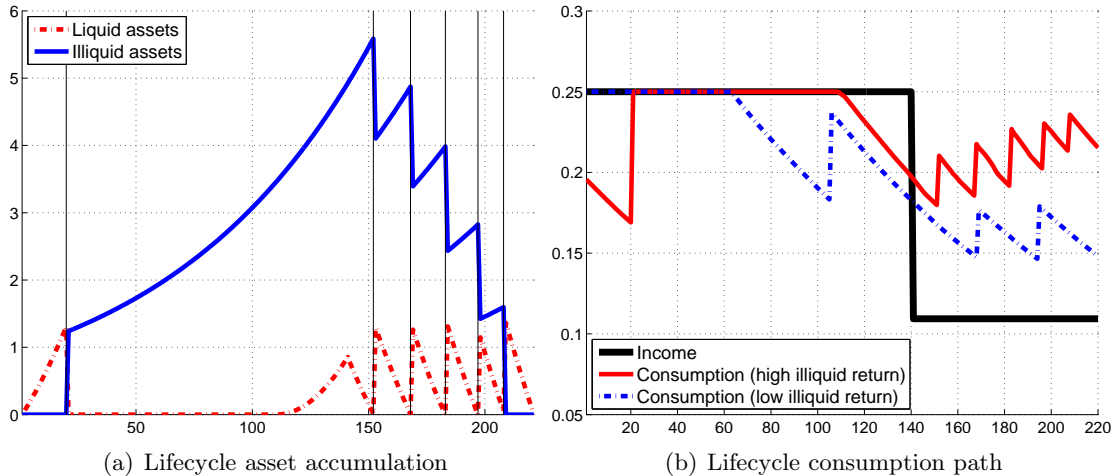(a) Lifecycle asset accumulation   (b) Lifecycle consumption path

Figure 2: Example of lifecycle asset and consumption paths with high return differential

we have also plotted the paths for consumption that arise 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' one $1/q^a$. The sawed pattern for consumption arising 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 consumption 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, except for a higher R^a. This higher return leads to higher overall wealth accumulation. Importantly, rather than increasing the *number* of deposits during the working life, the household changes the *timing* of its single deposit. The single 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, while the illiquid asset holdings are positive and are growing in value. Intuitively, since her net worth is large, this household would like to consume more than her income flow, however the transactions cost prevents her from doing this. This is a household that, upon receiving the rebate check, will consume a large part of it and, upon the *news* of the rebate, will not increase her expenditures.[14]

---

[14]The model by Campbell and Hercowitz (2009) also generates 'wealthy constrained' agents. However, their mechanism

Why would households choose to consume their income and deviate from the optimal consumption path imposed by the short-run Euler equation EE-SR? To understand, it is useful to recall Cochrane's (1989) insight that, in a representative agent model with reasonable risk aversion, the utility loss from setting consumption equal to income are very small. Our households are better off taking this loss than optimally smoothing consumption because the latter option entails accumulating more liquid wealth and foregoing the high return on the illiquid asset (and the associated higher future consumption).

4   Calibration

We now present a calibration of the model without idiosyncratic risk, but with heterogeneity in earnings growth. In Section 7 we present an extension of the model that allows for a realistic degree of idiosyncratic earnings risk. In our baseline model we exclude borrowing, however we show in Section 5 that our results are not sensitive to this assumption for realistic interest rates on unsecured credit.

Demographics and preferences   Decisions in the model take place at a quarterly frequency. Households begin their active economic life at age 25 ($j = 1$) and retire at age 60 ($J^w = 140$). 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).

Earnings heterogeneity   We estimate the parameters of the earnings process 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 PSID we select a sample of households with heads 25-59 years old in 1969-1996, following the same criteria as in Heathcote, Perri, and Violante (2010). We then construct the empirical mean age-earnings profile, and covariance functions in levels and first differences, exploiting the longitudinal dimension of the data. We then simulate the process in (3) at a quarterly frequency, with z_{ij} modelled as an i.i.d. shock, and aggregate quarterly earnings into annual earnings. We interpret the transitory component as measurement error in earnings and hence, in simulations, we abstract from it. From the implied annual earnings we construct the model counterpart of the empirical moments and minimize the distance between the two set of moments. The model's fit, exhibited in Figure 15 in Appendix A.1, demonstrates that we induce the right degree of heterogeneity in life-cycle earnings paths.

is different. They assume 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 that model the response of households to anticipated transitory income shocks will be zero and hence rebate coefficients will still be of negligible magnitude, as in standard one-asset models.

Table 2: Household Portfolio Composition (SCF)

	Mean	25th pct	50th pct	75th pct	Mean Return	Median Return
Earnings	58,036	21,586	41,116	71,953		
Net worth	219,387	16,600	86,650	255,000		
Liquid wealth	14,022	50	2,400	11,000	2.9%	3.2%
Liquid assets	15,412	710	3,200	11,700		
Revolving credit card debt	1,390	0	0	450		
Illiquid wealth	205,365	14,600	80,000	235,850	12.0%	12.8%
Housing net of mortgages	85,885	0	39,000	109,000		
Retirement accounts	43,304	0	1,000	32,000		
Directly held MF, stocks, bonds, T-Bills	33,139	0	0	3,800		
Vehicles net of installment loans	15,063	4,100	11,100	21,100		
Life insurance	8,668	0	0	1,000		
Certificates of deposit	5,449	0	0	0		
Saving bonds	1,121	0	0	0		

Households' portfolio data   Our data source is the 2001 wave of the Survey of Consumer Finances (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 2% of households by net worth and focus on households with heads aged 25 and over. We adopt the Federal Reserve Board's definition of liquid assets: money market, checking, savings and call accounts net of revolving debt on credit card balances.[15] All other wealth, excluding equity held in private businesses, is included in our measure of illiquid assets. It comprises CDs, directly held stocks, bonds and mutual funds, retirement accounts, housing net of mortgages and home equity loans, vehicles net of installment loans, life insurance policies, and other smaller categories. Table 2 reports some descriptive statistics.

The data provide overwhelming evidence that the majority of household wealth is held in (what we call) illiquid assets. For example, the 25th, 50th and 75th percentiles of the liquid asset distribution are $50, $2,400, and $11,000, compared with $14,600, $80,000 and $235,850 for the illiquid asset distribution.[16]

---

[15]The SCF asks the following questions about credit card balances: "How often do you pay on your credit card balance in full?" Possible answers are: Always or almost always, Sometimes, Almost never. And "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 Sometimes or Almost Never. Then, we 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.

[16]This finding is robust to other broader definitions of liquid wealth that we have considered. For example, treating directly held stocks, bonds, and T-Bills as liquid wealth, median liquid wealth is $3,600 and median illiquid wealth is $75,800.
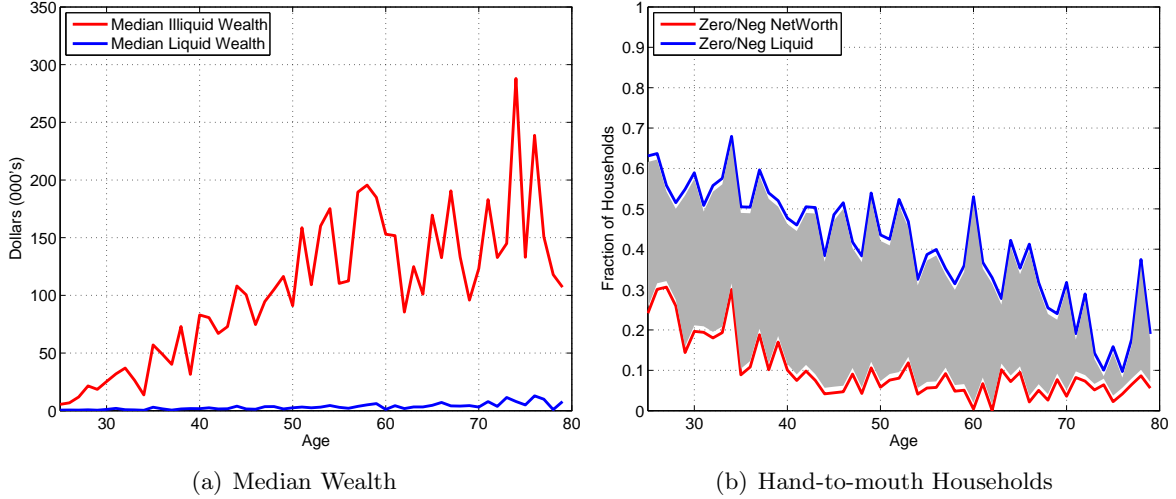
(a) Median Wealth   (b) Hand-to-mouth Households

Figure 3: Liquid and illiquid wealth over the lifecyle

Another important observation is that, for almost 70 percent of the population, illiquid wealth is entirely composed by housing, cars and retirement accounts, i.e., items whose transaction is quite costly. Our calibration strategy sets the level of impatience (i.e., the discount factor β), and the degree of illiquidity (i.e., the transaction cost κ) to match the median holdings of illiquid and liquid wealth, respectively.[17]

Figure 3(a) shows the evolution of illiquid and liquid wealth over the lifecycle. It is clear that the vast majority of the life-cycle saving over the working life takes place in illiquid wealth, whereas liquid wealth is fairly constant. An implication of these low holdings of liquid wealth is that a number of households are likely to be hand-to-mouth, i.e., households who hold assets only because of a mismatch in the timing of consumption and income *within* a pay period, and not to save *across* periods. Figure 3(b) provides an estimate of the number of such households, using two alternative definitions of wealth: net worth, which is the relevant definition for comparison with one asset models; and liquid assets, which is the relevant definition for comparison with our two asset model. The grey area in between the two lines is composed by households who have positive illiquid wealth but do not carry positive balances of liquid assets between pay periods. These are empirical counterpart of the 'wealthy hand-to-mouth' agents in our model. Note that this last group of households represents a sizeable fraction (27%) of the population.[18]

---

[17]An alternative approach would be to incorporate into the calibration of κ direct empirical evidence about frequency of transactions in mutual funds, retirement accounts, vehicles, and housing.

[18]In constructing these fractions from SCF data, we take into account that self-reported balances of checking and saving accounts (as well as net worth) includes the periodic earnings flow. In Figure 3(b), we therefore report the fraction of agents with liquid balances (or net worth) less or equal than monthly earnings divided by two, under the assumptions

These findings are consistent with recent survey evidence in Lusardi, Schneider and Tufano (2011) showing that half of US households would "probably or certainly be unable to cope with a financial emergency in the next month that required them to come up with $2,000." Moreover, a high proportion of individuals at middle class levels of income report they are certainly or probably not able to cope with this type of financial emergency.

Asset returns   To measure returns, we focus on the period 1967-2006. We set the nominal return on checking accounts to zero and on savings accounts to the interest on 3-months T-bills, 5.93% over this period (FRB database). We apply this same return to the holdings of T-bills. For equities, we use CRSP value-weighted returns, assuming dividends are reinvested, and obtain an annualized return of 12.3%. For housing, we follow Ludvigson et al. (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 annual return for year t is constructed as housing wealth in the fourth quarter of year t plus housing consumption over the year divided by housing wealth in the fourth quarter of year $t-1$. We then subtract population growth in order to (partially) correct for the growth in housing quantity. We obtain a nominal return of 17.8%. Finally, for CDs we compute a nominal return of 6.7% (FRB database). We apply these nominal returns to each household portfolio in the SCF and compute average and median returns in the population. The final returns on liquid and illiquid wealth are reported in Table 2. The implied average (median) nominal return on illiquid wealth is 12.0% (12.8%) and on liquid wealth is 2.9% (3.2%). Finally, we set the annual inflation rate to 3 percent. Given inflation and taxes, the after-tax real return on the two assets are 6.9% and -0.8%, respectively.

Initial asset positions   We used observed wealth portfolios of households aged 23 to 27 to calibrate the age $j = 0$ initial asset holdings in the model. We divide the sample into five groups baed on earnings and assign the corresponding simulated households in the model the median liquid and illiquid assets in each group.[19]

Government   We set τ^{SS} to 12.4% in order to reproduce the Old-Age, Survivors, and Disability Insurance (OASDI) tax rates in 2000. To compute social security benefits, average lifetime earnings Y_i 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 $\tau^y(y)$

---

that (i) wages and salaries are paid at a monthly frequency, and (ii) consumption is uniformly spaced within a month. To adjust for this effect, we assume that households are paid monthly and that consumption is uniformly spaced within a quarter.

[19]The five groups correspond to the discretization of the fixed component of earnings in the model. From lowest to highest earnings these groups contain 6.5%, 25%, 37.25%, 35% and 6.25% of the population. The liquid and illiquid asset holdings are $0, $50, $300, $3,400, $14,000 and $0, $0, $7,900, $24,700 and $43,900, respectively.

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. We use the same source to compute tax rates on interests (for τ^m) and on dividends and capital gains (for τ^a). Based on this evidence, we set $\tau^m = 0.24$ and $\tau^a = 0.165$. When we set government expenditure G to its value in the year 2000 (using aggregate income as metric), residual public debt from the government budget constraint is close to its observed value.

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

Tax rebate   We set the rebate size to $500, since JPS (2006) report that the average check was $480 per household and most households received either $300 (single-filers) or $600 (if filing as married). The economy is in steady state in 2001:Q1. The rebate checks are randomly sent out to half the eligible population in 2001:Q2, and to the other half in 2001:Q3.

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.[20] The first group receives the check before making the consumption/saving and adjustment decision for that quarter.

We then compute the transitional dynamics of our economy. We begin by assuming that the government finances the rebate by increasing debt for ten years, and then repays the rebate outlays and the accumulated interest by introducing a permanent proportional tax on income. In our benchmark calibration the required additional tax rate is around 0.05%.

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. Most of these changes were planned to be phased in gradually over the five years 2002-2006 and to 'sunset' in 2011. It turned out that instead of sunsetting, the tax cuts were extended in 2011 for a further year.

We consider two scenarios regarding the 'sunset' clause: (i) households believe that the tax system

---

[20]We discuss how to relax this assumption in Section 7.

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 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 is 2002:Q1, and date t^{**} is 2006:Q1 without the sunset, and 2011:Q1 when the tax reform sunsets as legislated. We construct the sequence of effective tax schedules based on Kiefer et al. (2002, Table 5).

Recession To model the downturn of 2001, we assume that in 2001:Q2 households become aware that they are in a recession. At this time they learn that labor income will fall linearly for the next three quarters, generating a cumulative drop of 1.5%, and will then recover to its initial level over the following eight quarters.[21]

5 Quantitative analysis of the 2001 tax rebate

We start by studying the special case of the one-asset economy, where $\kappa = 0$. Then, we study our two-asset economy. Our baseline results refer to the tax rebate in isolation. Next we incorporate the EGTRA tax reform and 2001 recession.

5.1 One-asset model

In the one-asset version of the model, corresponding to the special case of $\kappa = 0$, we choose $\beta = 0.985$ to reproduce median net worth in the SCF data and we set the interest rate to the calibrated value of R^a.[22] The left panel of Figure 4 displays the dynamics of consumption (and the rebate receipts) for the two groups along the transition. The contemporaneous response of consumption to the rebate is tiny. When we compute the rebate coefficient exactly as in JPS (2006) –i.e., we run regression (1) on simulated panel data– the estimated rebate coefficient is 3.7%.

Only young agents (who are more likely to be constrained), and elderly agents (who have a short horizon) have a substantial MPC out of the transitory rebate. However, recall that the rebate coefficient measures the difference between the treatment group who receives the stimulus payment and everyone else. This latter control group comprises of households who already received the payment and households who know they will receive it. If households in the control group are unconstrained –and old households would be– consumption smoothing dictates that their response should approximately

---

[21]The 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 NIPA Wages and Salaries series.

[22]This latter choice has no bearing on the findings. When we set the interest rate to the calibrated value of R^m, we found very similar results.

(a) Time path of rebate and consumption   (b) Households with zero or negative wealth
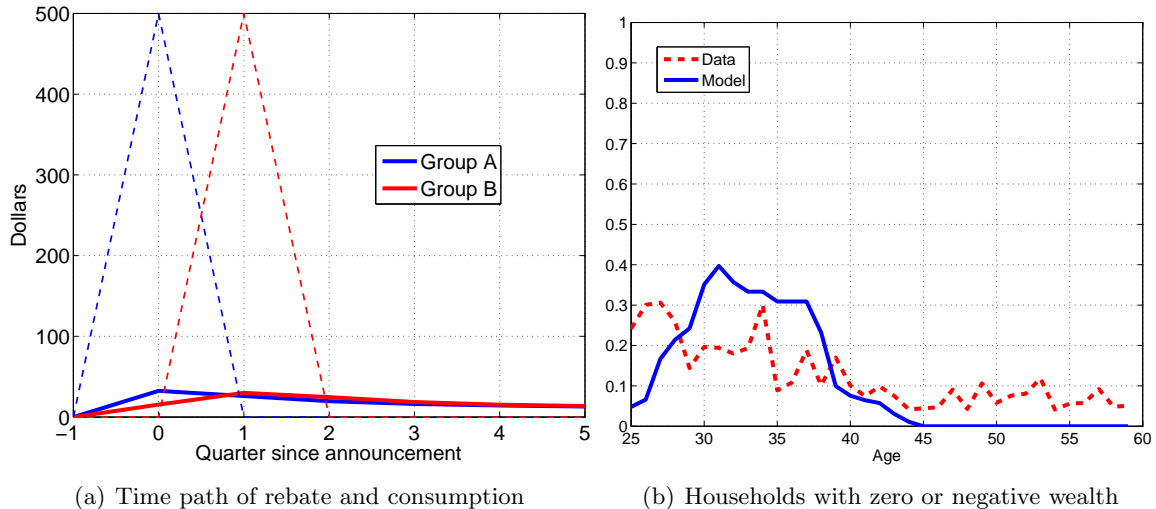
Figure 4: Rebate experiment in the one asset model

equal that of the first group and the estimated rebate coefficient would be zero. Therefore, the action comes only from the constrained agents. In our calibration (right panel of Figure 4), 8% of the model agents have zero or negative net worth. As reported in the Figure, in the 2001 SCF data this fraction is also around 8%.[23] Hence, there is essentially no scope for the one asset model to generate high rebate coefficients, while remaining consistent with SCF data on the distribution of net worth.[24]

5.2   Two-asset model

We begin by analyzing the model economy for different values of the transaction cost κ ranging from zero to $1,000. We then zoom into the economy with $\kappa = \$500$ For each value of κ, we reset β to match median holdings of illiquid wealth. In all the plots which follow, the one-asset economy corresponds to the point $\kappa = 0$.

Baseline results   The fraction of households adjusting (i.e., accessing the illiquid account to withdraw or deposit) falls with κ, as shown in Figure 5(a). For the intermediate case of $\kappa = \$500$, one fifth of the population adjusts in every quarter. As is clear from the simulations of Section 3, retirees adjust more often than working-age households because they finance their consumption largely by withdrawing from the illiquid account.[25] Holdings of liquid wealth increase with the transaction cost

---

[23]As explained above, to compare data and model, we take explicitly into account that the balances of liquid wealth reported by SCF respondents include the average earning flow accruing periodically to households.

[24]As a back of the envelope calculation, half of the constrained agents are in group A and half in group B. If their consumption response is zero to the news and one to the rebate, then the rebate coefficient is roughly half the size of the fraction constrained, i.e., roughly 4% which is very close to the estimated 3.7%.

[25]In the economy with $\kappa = \$500$, workers adjust roughly every 10 quarters and retirees every 2 quarters.

19

(a) Fraction of households adjusting

(b) Distribution of liquid wealth

(c) Fraction of hand-to-mouth households

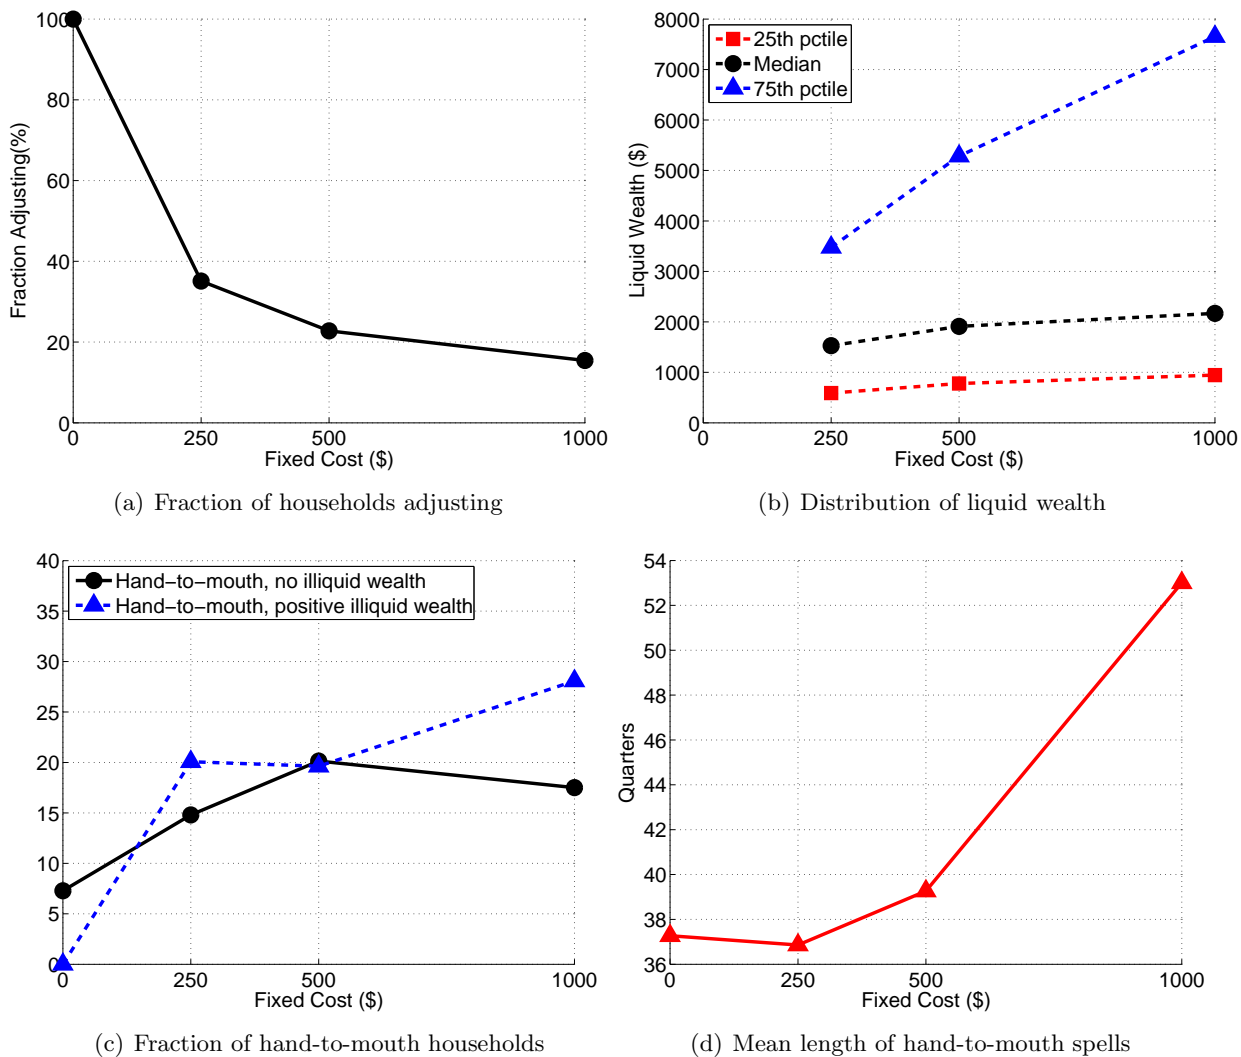(d) Mean length of hand-to-mouth spells

Figure 5: Features of two asset model, by transaction cost

(5(b)), because households deposit into/withdraw from the illiquid account less often and carry larger balances of liquid assets.

A corollary of the wealth distribution in Figure 5(b) is that a substantial fraction of agents do not carry any balances of liquid assets across periods. In the model, there are two types of such agents. Some agents have have zero liquid assets at the end of the period because they are adjusting.[26] Others have zero liquid assets at the end of the period because they are hand-to-mouth and consume their earnings every period. Figure 5(c) plots the fraction of households in the latter group - the hand-to-mouth consumers - and divides them into those who also have zero illiquid wealth and those with

---

[26]When households adjust they optimally choose to hold zero liquid wealth in the period before (after) adjusting if retirees (workers).

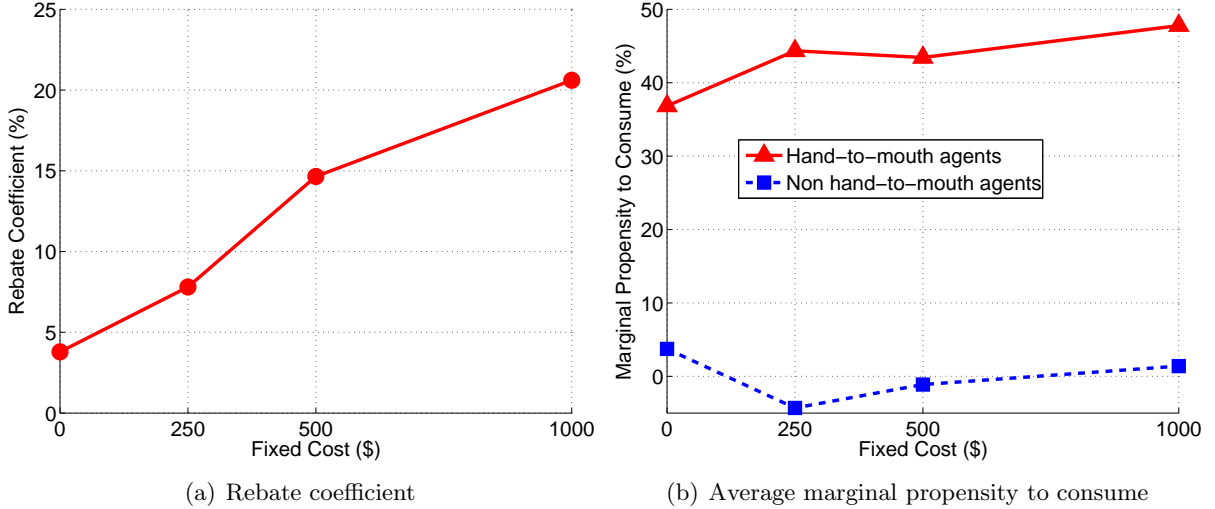(a) Rebate coefficient   (b) Average marginal propensity to consume

Figure 6: Rebate coefficient and marginal propensity to consume in two asset model, by transaction cost

positive illiquid wealth. The size of this latter group is increasing in κ. Finally, Figure 5(d) reports the average length of spells in which hand-to-mouth households hold zero liquid assets and consume their income.

Figure 6(a) displays the rebate coefficient in the model for different levels of the transaction cost. The rebate coefficient grows steadily from 3.7% when $\kappa = 0$ (the one-asset model) to 20.6% when $\kappa = \$1000$. For moderate transaction costs, e.g., $\kappa = \$500$, the rebate coefficient is increased by a factor of 4 relative to the one-asset model. Figure 6(b) shows the marginal propensities to consume (MPC) out of the rebate check for two groups of households, unconstrained and constrained. Note how for the former group the average MPC is close to zero, and sometimes negative, while for the latter group it is around 50%.[27] Therefore, the high rebate coefficients are mainly driven by constrained households, as in the standard one-asset model. However, the two-asset model generates a much larger fraction of liquidity constrained households, some of which hold sizeable quantities of illiquid assets.

The JPS (2006, 2009) estimates by sub-groups of households with different amounts of liquid assets are suggestive that low wealth households spend more of the rebate, but given the small sample size, they are not statistically significant. Broda and Parker (2008) perform a similar estimation on the AC Nielsen Homescan database, a sample 50 times larger, and find very strong evidence that low-liquid assets households spend at least twice as much as the average.

Borrowing   In the baseline experiment, we ruled out access to unsecured credit. Figure 7(a) plots

---

[27]In the two-asset model, marginal propensities to consume out of moderate amounts (in this case $500), can be negative because agents change the optimal path of future adjustments, thus changing the effective time horizon.

(a) Rebate coefficient by borrowing rate   (b) Rebate coefficient for different assumptions about information sets
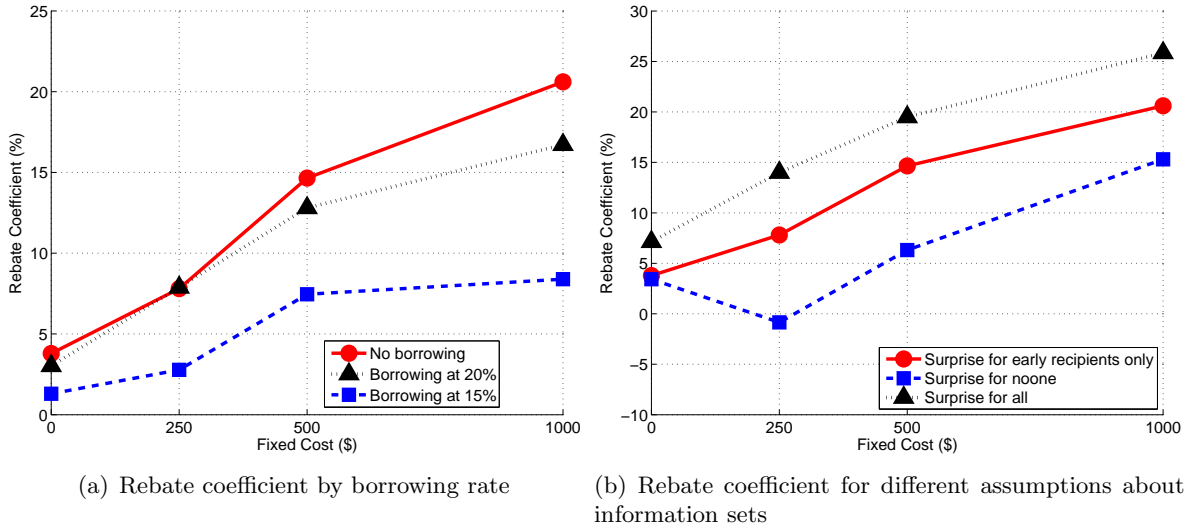
Figure 7: Effect of borrowing rate and information set, by transaction cost

the rebate coefficient as a function of the transaction cost when we allow costly borrowing at annual nominal interest rate of 20% and 15%. For borrowing rates around 20% or higher, the rebate coefficient is barely affected. The reason why costly borrowing does not affect the rebate coefficients is that the no arbitrage assumption $q^b \leq q^a$ necessarily implies that $q^m > q^a$. The strict inequality means that there is always a range of outcomes where a substantial fraction of households continues to act as hand-to-mouth consumers even with the possibility of borrowing. Rebate coefficients are reduced by the ability to borrow if those households receiving the rebate later find it worthwhile to borrow for one quarter in anticipation of the payment. The figure shows that this happens for low enough interest rates on unsecured borrowing.

Information sets   In Figure 7(b), we report the consumption response to the tax rebate under alternative assumptions about when the rebate enters households' information sets. When we assume that the rebate is a surprise for everyone, the rebate coefficient increases by around 5 percentage points on average. When we assume it is anticipated by every household (by one quarter for the first group and two quarters for the second group), the estimated rebate coefficient drops significantly. With this latter assumed information structure, although the measured rebate coefficients are smaller, the aggregate cumulative effect on consumption is very similar to the baseline specification.

Tax reform and recession   Figure 8(a) shows that the consumption responses to the tax rebate are substantially higher when the full tax reform is modeled, regardless of whether the sunset clause is implemented. With a transaction cost of $500, the rebate coefficient increases from 15% to 25%. The reason is that the substantial reduction in future tax liabilities leads to a rise in the desired level of

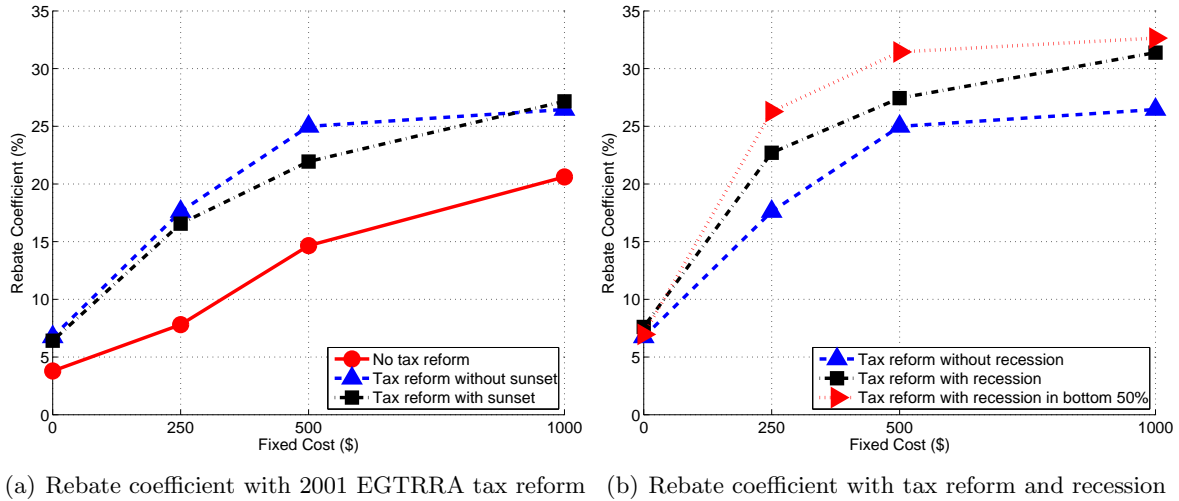(a) Rebate coefficient with 2001 EGTRRA tax reform   (b) Rebate coefficient with tax reform and recession

Figure 8: Effect of borrowing rate and information set, by transaction cost

lifetime consumption. Since a substantial fraction of households (poor and wealthy) are constrained, the rebate enables such households to start consuming out of this additional future income earlier than they would otherwise be able to.[28]

A similar logic applies when we add the recession to the tax rebate experiment. Figure 8(b) shows that allowing for the fact that the economy was undergoing a recession at the time of the rebates adds roughly 3% to the rebate coefficient. When the earnings drop induced by the downturn is concentrated in the bottom half of the earnings distribution, the effect is twice as large.[29] 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. Compared to the spending response to the 2001 rebates, their point estimates of the contemporaneous response of consumption are about 60 percent 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 suggests a mechanism why this force may be at work.

To explore cross-sectional properties of our model, we need to select a specific parameterization. We

---

[28]When adding the tax reform, the economy with $\kappa = 0$ also yields a higher rebate coefficient because the effect we describe does apply to the agents who are constrained in that model. However, the two-asset model magnifies the interaction between constraints and the tax reform (i.e., the distance between the lines is larger for $\kappa > 0$), because more households are constrained.

[29]The recession experiment with $\kappa = 0$ is the closest to Heathcote (2005, Table 4) who finds that, in an infinite horizon standard incomplete market model, the MPC out of a temporary reduction in taxes is 13.5%.

(a) Time path of rebate and consumption

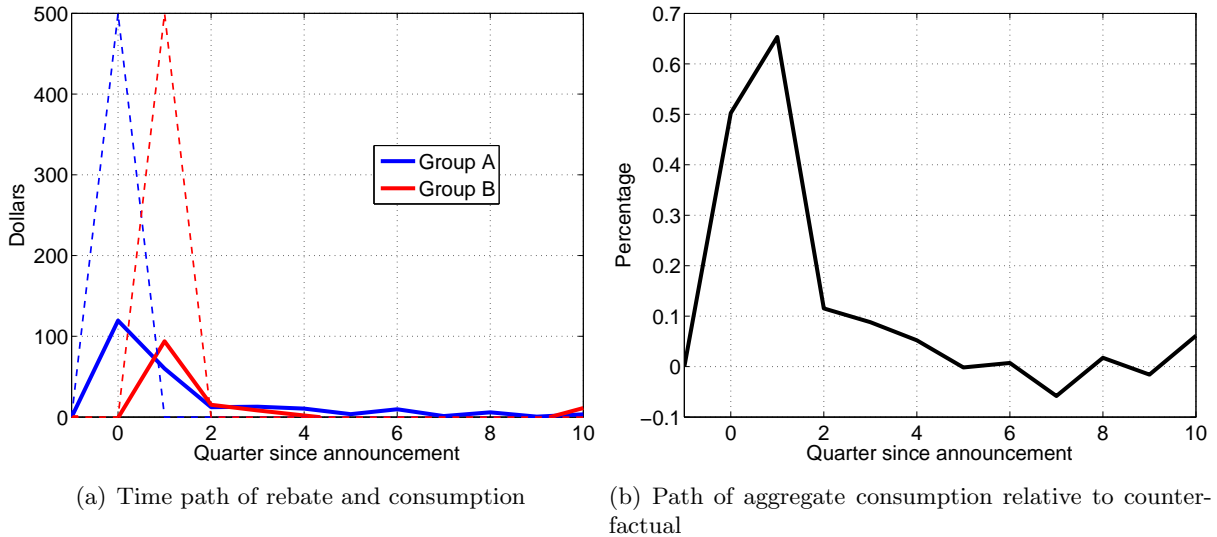(b) Path of aggregate consumption relative to counter-factual

Figure 9: Rebate experiment in the two asset model: consumption dynamics

choose the economy without borrowing, $\kappa = \$500$, where the rebate coefficient is a surprise for half of the population, the recession is uniformly distributed, and the tax reform does not sunset. The estimated rebate coefficient in this economy is 27%, in line with the empirical estimates.

Response of aggregate consumption   Figure 9(a) plots the dynamics of 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 9(b) demonstrates that the impact of the rebate is to increase aggregate quarterly consumption expenditures by 0.5% and 0.65%, respectively, in the first two quarters. The impact is very short lived: a year later differential consumption growth is essentially zero.

We note that, to calculate the impact on aggregate consumption, it is incorrect to multiply the rebate coefficient by the total size of the rebate (as a fraction of aggregate consumption). The reason is that the rebate coefficient β_2 is not the MPC the rebate, but the difference between the MPC of those who receive the check and that of the rest of the population.

Heterogeneity in rebate coefficients   In Section 2, we discussed the evidence on the heterogeneity in rebate coefficients across the population. Misra and Surico (2011) use the same data set as JPS (2006) to conduct a detailed analysis of cross-sectional heterogeneity in rebate coefficients. They conclude that there is a large amount of heterogeneity in consumption responses. In particular: (i) the distribution of consumption responses is bimodal, with around 40% of households saving all of the rebate, and another large group of households spending a high fraction of the rebate; and (ii)
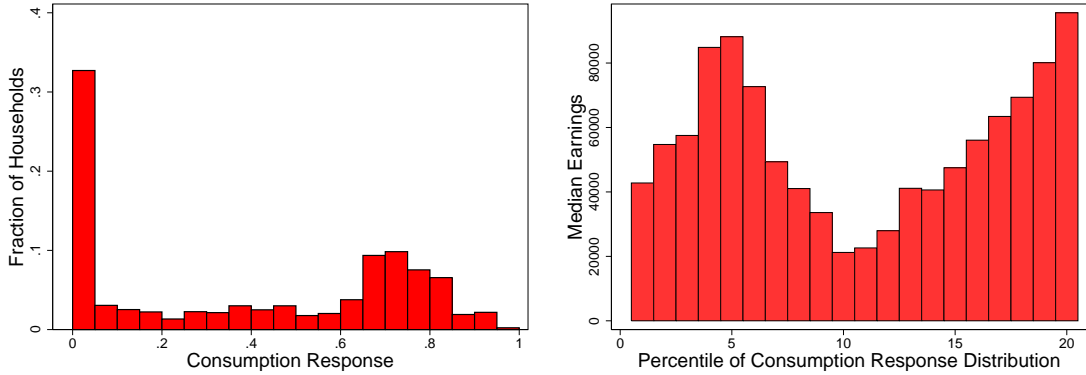
(a) Distribution of rebate coefficients for working age population

(b) Median earnings by quantile of rebate coefficient distribution

Figure 10: Heterogeneity in rebate coefficients

high income households are disproportionately concentrated in the two tails of the distribution of consumption response. Figure 10 shows that the model can broadly reproduce all of these findings. Figure 10(a) plots a histogram of rebate coefficients for the working age population, in which the bimodal nature is quite stark. The bimodality arises because of the coexistence of a substantial fraction of (wealthy) hand-to-mouth consumers with unconstrained agents who save the rebate as predicted by the LCPIH.

Figure 10(b) plots median earnings in each quantile of the cross-sectional distribution of the consumption responses. The U-shape pattern is consistent both with Misra and Surico (2011), and also with the results in JPS (2006) when they estimate average rebate coefficients separately in each third of the income distribution. The reason why there are high earnings households at both ends of the distribution is that some of them are unconstrained and some constrained. The former are income-rich because their earnings profile is high but flat, the latter are income-rich because their earnings profile is steep, which makes the liquidity constraint more likely to bind.

6   Further validation of the model

As discussed in the introduction, Cochrane (1989) and others have observed that there may be only small welfare differences between alternative consumption behavior in reacting to small transitory income shocks. While we argued that this, in no way, belittles the importance of *understanding* such consumption responses, it does imply that we must bring more evidence in order to validate our model, a point also raised by Card et al. (2007). There are a number of ways that we can do this. We begin
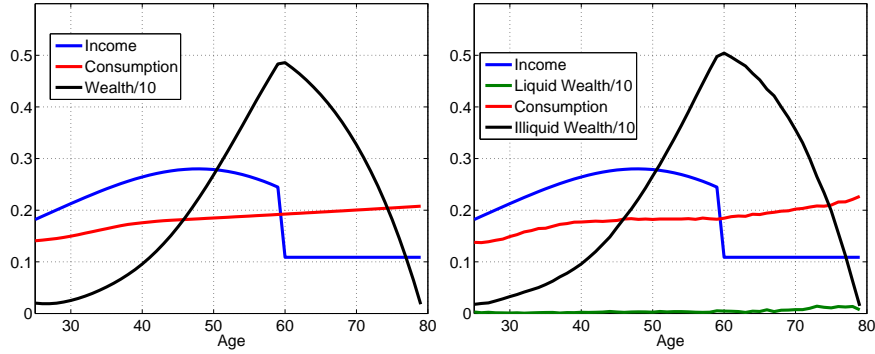
(a) Mean earnings, consumption and wealth in one asset model

(b) Mean earnings, consumption and wealth in two asset model
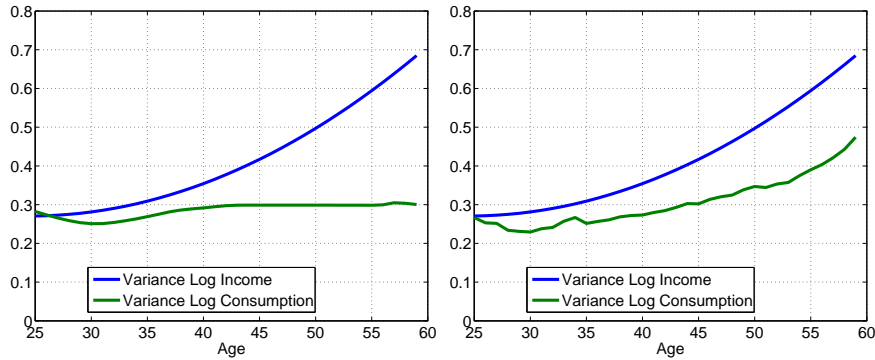
(c) Variance log earnings and log consumption in one asset model

(d) Variance log earnings and log consumption in two asset model

Figure 11: Lifecycle profiles (means and variances) in the one-asset and two-asset models

by looking at a set of cross-sectional and life-cycle implications

Life-cycle implications Figure 11 compares the life-cycle means and variances of earnings, consumption and net worth across the one-asset and two-asset models. The models are hardly distinguishable along these dimensions. Both models reproduce the well-known features of the data (e.g., see Heathcote, Perri and Violante, 2010) quite well, and the two-asset model is able to generate a sizeable rise in consumption inequality even in absence of shocks.

In Figure 12, we contrast the implications of the model for the lifecycle distribution of liquid and illiquid assets to the SCF data. Figure 12(a) plots the P25, P50 and P75 of the distribution of illiquid wealth in the model and in the data, and Figure 12(b)) plots the same percentiles for the liquid wealth distribution. Figures 12(c) and 12(d) reproduce the fraction of hand-to-mouth households and those with positive illiquid assets, respectively, over the life cycle in the data and in the two-asset model. Even though all these cross-sectional lifecycle profiles were not targeted by our calibration procedure, the model's counterpart are quite close to the data.
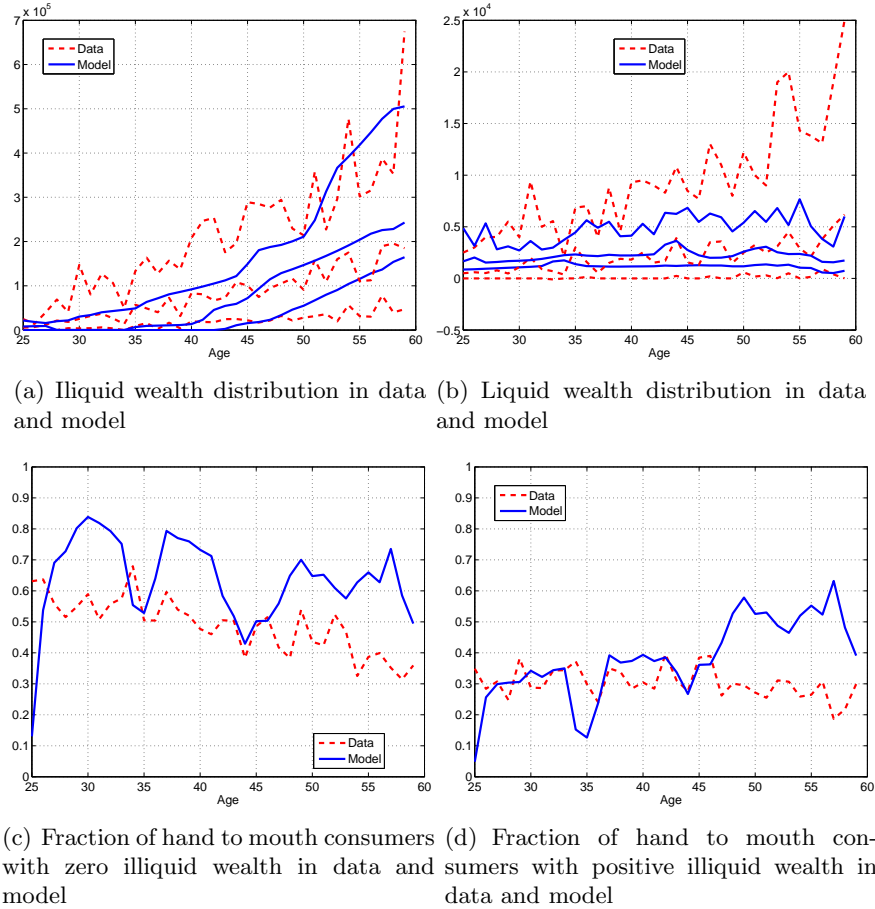
(a) Iliquid wealth distribution in data and model

(b) Liquid wealth distribution in data and model

(c) Fraction of hand to mouth consumers with zero illiquid wealth in data and model

(d) Fraction of hand to mouth consumers with positive illiquid wealth in data and model

Figure 12: Wealth distribution over the life cycle in data and model

Size of rebate   Hsieh (2003) shows that the same CEX households that 'overreact' to the 2001 tax rebates respond in line with standard theory to payments received from the Alaskan Permanent Fund, which are, on average, much larger than the rebate. One interpretation of this fact is that, although households spend large fractions of small anticipated shocks, they predominantly save large anticipated shocks. Browning and Collado (2001) document similar evidence from Spanish survey data: workers who receive anticipated double-payment bonuses in the months of June and December do not alter their consumption growth significantly in those months.

Figure 13 shows how the propensity to consume the rebate declines when the size of the rebate is increased above $500. For a $500 transaction cost, the rebate coefficient drops by over a factor of five (from 15% to 2%) 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.[30] Moreover, for rebates that are sufficiently high relative to the transaction

---

[30]This occurs if the size of the rebate is larger than the amount by which consumption would increase if the household
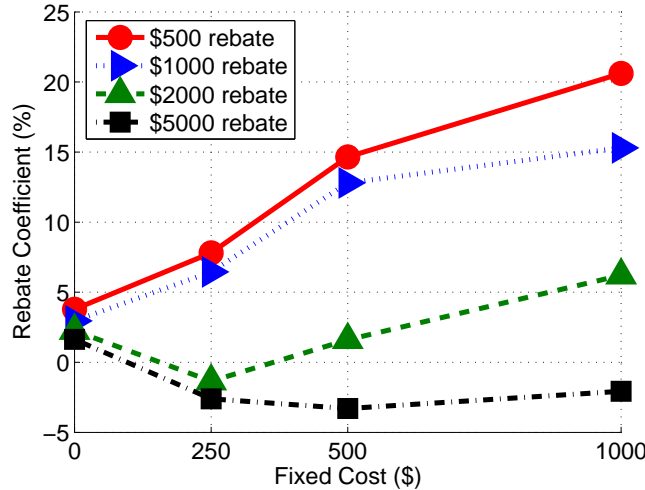
(a) Rebate coefficient by size of the rebate

Figure 13: Rebate coefficients by size of the stimulus payment

cost, many working households will choose to pay the 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 in *absolute terms* as the rebate size increases, for given transaction cost.In the example in Figure 13, when the transaction cost is $500, the average consumption expenditure is larger for a $500 rebate than for a $2,000 rebate. This example underscores the relevance for policy of understanding the structural mechanism that lies behind the empirical evidence.[31]

7   Idiosyncratic Risk

In baseline calibration we took view that the increase in earnings dispersion over the life cycle is the result heterogeneity in deterministic idiosyncratic earnings profiles. An alternative view, more common in the literature, is that it is the result of the accumulation of highly persistent idiosyncratic shocks. In this section we show that our results are not sensitive to this choice.

We model the log household earnings process as the sum of two orthogonal components: a unit root

---

were not currently constrained.

[31]Figure 13 shows how estimated rebate coefficient may even become negative. This occurs when the stimulus payment is so large enough relative to the transaction cost that a substantial fraction 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 the control group.

(a) Rebate coefficients by transaction cost   (b) Median liquid wealth by transaction cost
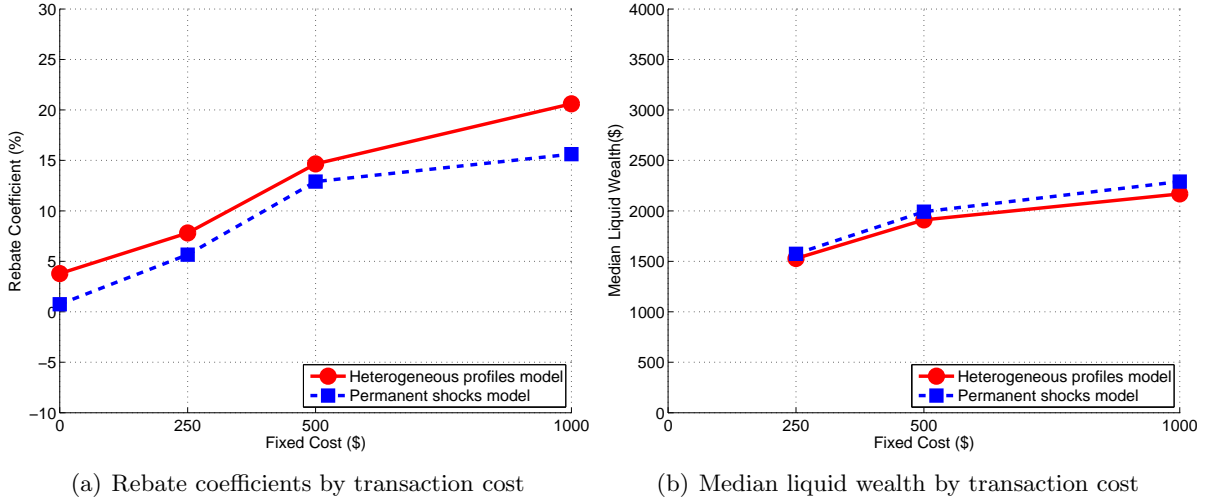
Figure 14: Rebate coefficients and liquid wealth with idiosyncratic risk

and an i.i.d shock that, as in the heterogeneous slopes model, we interpret as measurement error. Also this statistical model matches features of the distribution of earnings well.

Figure 14 compares the rebate coefficients and median liquid wealth in the baseline model (without tax reform and recession). The implications of these two models are very similar. The reason is that with highly persistent shocks there is little incentive to hold liquid wealth for precautionary reasons, a well known result in the literature. This view of risk is consistent with the observed distribution of liquid wealth holdings: had we allowed for a large transitory component, households would hold counterfactually high quantities of liquid assets, and accordingly would have low marginal propensities to consume out of the rebate.

8 Concluding remarks

In this section, we list several directions which we plan to explore in future work.

8.1 Counterfactual policy analysis

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 the model can be used to investigate the counterfactual effects of other feasible stimulus policies. This exercise is essential for guiding future policy decisions. Among the various alternative policy options, we plan to analyze the following three.

Targeted tax rebates   An implication of incomplete markets and heterogeneous agents is that there is a distribution of individual propensities to consume tax rebates, which are highest for constrained agents. For this reason, Elmendorf et al. (2008), among others, have argued that stimulus policies aimed at lower income households would deliver a larger aggregate consumption response for the same fiscal cost. Our structural model is a natural laboratory to study various alternative designs of tax rebates (e.g., different phase-in/out thresholds, means testing, rebate size dependent on the tax base, etc...) that have the same net effect on the government budget constraint.

Unemployment benefits   Since our model is calibrated at a high frequency, it can be applied to questions about changes in the UI benefit system. Policies that temporarily increase the duration of unemployment benefits, as done repeatedly in the course of the last recession, may be especially effective in raising aggregate consumption. We plan to simulate the effects of this type of policy and contrast it to non-targeted fiscal stimulus payments.

Temporary reductions in consumption tax   In November 2008, the Chancellor of the Exchequer in the United Kingdom announced a temporary (13 month) reduction in the standard rate of Value Added Tax (VAT) from 17.5% to 15%.[32] This type of fiscal intervention is an important alternative to the direct payments that have been favored in the US. We will simulate a temporary reduction in the consumption tax in the context of the 2001 and 2008 rebates, by reducing the tax rate by an amount that has the same effect on the government budget constraint as the actual tax rebates, and compare the short-run effects on consumption.

8.2   Behavioral approach

Our approach to explaining the consumption response to fiscal stimulus payments incorporates 'frictions' in asset markets but abstracts from 'behavioral' biases. Even though expanding our analysis in this direction is beyond the scope of our project, we devote our final thoughts to this class of models to argue that there is no obvious candidate among them.

A number of these frameworks (e.g., myopia, hyperbolic discounting, self-control) provide behavioral foundations for the existence of impatient consumers, and can hence generate large MPCs.[33] However, without the addition of some form of transaction cost, they cannot generate small consumption responses to *news* about future payments. To match the evidence on anticipated income shocks, models based on information processing costs (or 'mental accounting') have been proposed. We favor associating the adjustment costs with financial transactions, rather than the processing of information, for

---

[32]See Blundell (2009) for a description of the policy and its likely impact on consumers' demand.
[33]Examples include Laibson (1998), Mankiw (2000) and Angeletos et al. (2001).

two reasons. First, financial transactions costs are more easily connected to observable data. Second, models of inertia in savings that lead to *large* consumption responses can equally be formulated as inertia in consumption leading to *small* consumption responses.[34]

8.3 Extensions

Aggregate fluctuations and beliefs over the rebate   Tax rebates policies are typically implemented during recessions. We plan to extend our model to include aggregate fluctuations. Given the high frequency OLG structure, solving a stochastic version of the model in general equilibrium (with asset returns determined endogenously) would not be numerically feasible, as explained by Krueger and Kubler (2004). A version of the economy with stochastic fluctuations which we are able to solve is one where the age-earnings profile and the return on liquid and illiquid asset vary stochastically, but exogenously. In this version of the economy, we can also endow agents with correct beliefs over the probability that the government will make fiscal stimulus payments to households during a recession.

Durables   PSJM (2011) document that when durable goods are included among expenditures, the response of household consumption to the 2008 rebate almost doubles. We plan to extend our model to address this finding. First, we note that models of durable consumption without non-convexities (either in the form of indivisibilities or transaction costs for durable purchases) suffer from the same shortcomings as the one asset model described above. As long as households purchase durables every period, they will generate small responses for durable expenditures. Hence infrequent adjustment of durables is the key element that is needed. A natural extension of our model is to embed the intuition from the model of illiquid assets to a setting where the infrequent adjustment instead applies to durables. Second, our definition of illiquid assets includes some items that could be classified as durable consumption (e.g. cars, housing). Our model can allow for an exogenous fraction of the illiquid asset to generate a utility flow for the household, which captures the idea that durable consumption is itself a vehicle for saving.

Automatic deposits into the illiquid account   Two of the largest categories in our definition of illiquid wealth are housing and private retirement accounts. A key feature of housing is that, after the purchase, households can increase its net value over time by regularly making mortgage payments. While the purchase entails a sizeable transaction cost, the fee on mortgage payments is rather small. Similarly, in the case of retirement wealth, households can set automatic deposits from their paycheck into their retirement account at no extra cost. We are extending our model to allow working households

---

[34]For example Reis (2006) is a model of inattentiveness with this property.

to choose a fixed percentage of their earnings to be diverted directly into the illiquid account in every period at no cost. Similarly, retired household can choose a percentage of their funds held in the illiquid account balance to be withdrawn at no cost each period. The transaction cost is only paid by households upon (re-)setting this percentage. As in the baseline model, households are still permitted to make lump-sum deposits and withdrawals into the illiquid account by paying the transactions cost.

A   Appendix

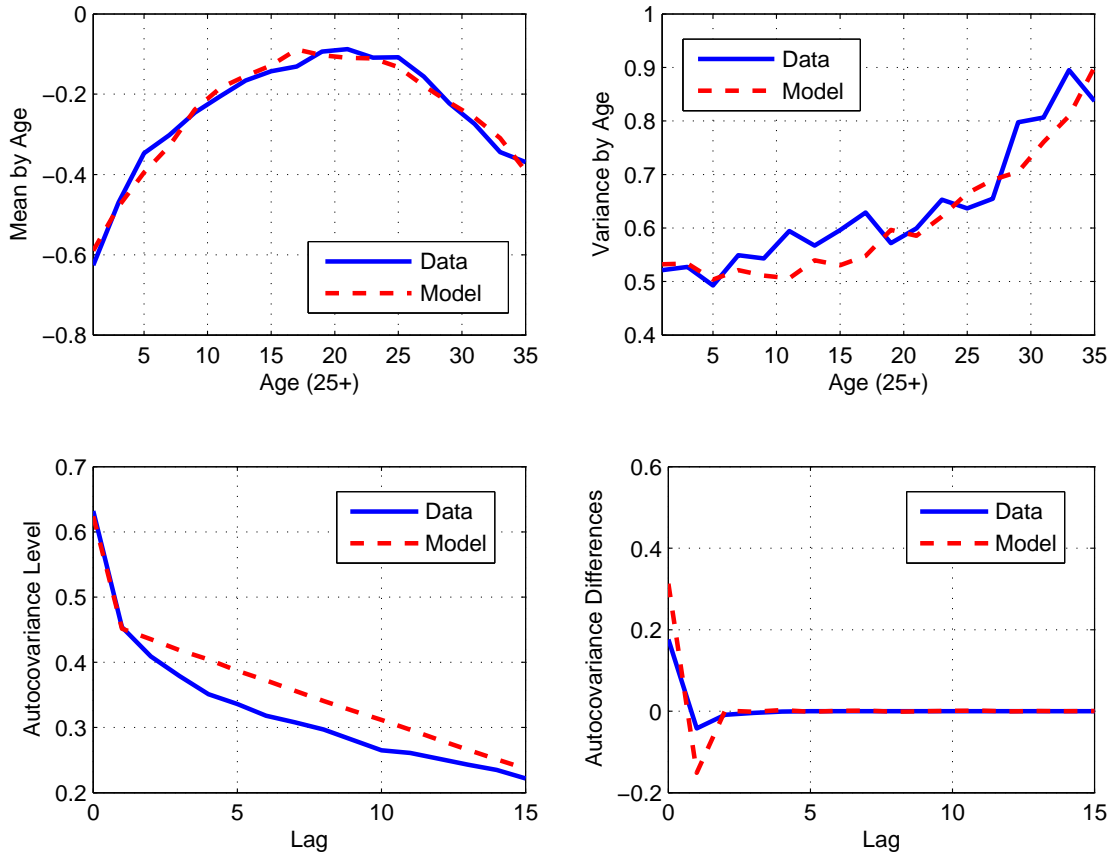## A.1   Quarterly earnings models

Figure 15: Fit of estimated earnings model with heterogeneous levels and slopes

References

AGARWAL, S., C. LIU, AND N. SOULELES (2007): "The Reaction of Consumer Spending and Debt to Tax Rebates-Evidence from Consumer Credit Data," *Journal of Political Economy*, 115(6), 986–1019.

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.

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.

ATTANASIO, O., C. MEGHIR, AND A. SANTIAGO (2009): "Education choices in Mexico: using a structural model and a randomized experiment to evaluate Progresa," *Mimeo, University College London.*

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. D. (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.

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.

ELMENDORF, D., J. FURMAN, W. GALE, AND B. HARRIS (2008): "Distributional Effects of the 2001 and 2003 Tax Cuts: How Do Financing and Behavioral Responses Matter?," *National Tax Journal*, 61(3), 365–380.

FAVILUKIS, J., S. C. LUDVIGSON, AND S. V. NIEUWERBURGH (2011): "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium," *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. L. 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 Reviws 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.

IMROHOROGLU, 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.*

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., T. MUKOYAMA, A. SAHIN, AND A. SMITH JR (2009): "Revisiting the welfare effects of eliminating business cycles," *Review of Economic Dynamics*, 12(3), 393–404.

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.

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.

MARTIN, J. (1996): "Measures of replacement rates for the purpose of international comparisons: a note," *OECD Economic Studies*, 26(1), 98–118.

MCDANIEL, C. (2007): "Average Tax Rates on Consumption, Investment, Labor and Capital in the OECD: 1950–2003," *Mimeo, Arizona State University*.

MISRA, K., AND P. SURICO (2011): "Heterogeneous Responses and Aggregate Impact of the 2001 Income Tax Rebates," *CEPR DP*, 8306.

MODIGLIANI, F., AND C. STEINDEL (1977): "Is a tax rebate an effective tool for stabilization policy?," *Brookings Papers on Economic Activity*, 1977(1), 175–209.

OH, H., AND R. REIS (2011): "Targeted Transfers and the Fiscal Response to the Great Recession," *NBER Working Paper*, 16775.

PARKER, J., N. SOULELES, D. JOHNSON, AND R. MCCLELLAND (2011): "Consumer Spending and the Economic Stimulus Payments of 2008," *NBER Working Paper*, 16684.

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.

SAHM, C., M. SHAPIRO, AND J. SLEMROD (2010): "Check in the mail or more in the paycheck: does the effectiveness of fiscal stimulus depend on how it is delivered?," *Working Paper, University of Michigan*, 1001, 48109–1220.

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.

SUZUKI, M. (2009): "Consumption Smoothing without Secondary Markets for Small Durable Goods," *IMES Discussion Paper Series*.

TELYUKOVA, I. A. (2011): "Household Need for Liquidity and the Credit Card Debt Puzzle," *mimeo*.

TODD, P., AND K. WOLPIN (2006): "Using experimental data to validate a dynamic behavioral model of child schooling: Assessing the impact of a school subsidy program in Mexico," *American Economic Review*, 96(5), 1384–417.

TOWNSEND, R., AND S. URZUA (2009): "Measuring the impact of financial intermediation: Linking contract theory to econometric policy evaluation," *Macroeconomic Dynamics*, 13(S2), 268–316.