

Working Paper Series
Congressional Budget Office
Washington, D.C.

Can Tax Rebates Stimulate Consumption Spending in a Life-Cycle Model?

Jonathan Huntley
(jon.huntley@cbo.gov)

Valentina Michelangeli
(valentina.michelangeli@cbo.gov)

Macroeconomics Analysis Division
Congressional Budget Office
Washington, D.C.

July 2011
Working Paper 2011-02

Working papers in this series are preliminary and are circulated to stimulate discussion and critical comment. These papers are not subject to CBO's formal review and editing processes. The analyses and conclusions expressed in them are those of the authors and should not be interpreted as those of the Congressional Budget Office. Citations in other publications should be cleared with the authors. Papers in this series are available at www.cbo.gov/publications (select Working and Technical Papers).

We thank Orazio Attanasio, Juan Contreras, Christopher D. Carroll, Robert Dennis, Ken Judd, Deborah Lucas, Damien Moore, Larry Ozanne, William Randolph, Felix Reichling, José-Víctor Ríos-Rull, Matthew D. Shapiro, Marika Santoro, Martino Tasso, and David Weiner for helpful comments and suggestions. An earlier version was circulated with the title "Tax Rebates and Fiscal Multipliers" and was presented at the 2010 NBER Summer Institute.

ABSTRACT

This paper presents a life-cycle model with earnings risk, liquidity constraints, and portfolio choice over tax-deferred and taxable assets to evaluate changes to household consumption in response to transitory, anticipated income shocks, such as the 2001 federal income tax rebate. Households optimally hold a large share of savings in tax-deferred assets, which are encumbered by withdrawal penalties, and choose to relinquish some taxable precautionary savings in exchange for higher after-tax returns. The model predicts an increase of several percentage points in the marginal propensity to consume out of the tax rebate compared with results from a standard frictionless life-cycle model. Moreover, the model's liquidity-constrained households—which include households with very few financial assets and those whose portfolios are expensive to reallocate because they contain sizeable tax-deferred investments—consume a higher fraction of the tax rebate than do other types of households. The results presented are in keeping with empirical findings published in the literature concerning the 2001 federal tax rebate.

1 Introduction

This paper presents a life-cycle model with asset choice and financial frictions that reproduces, more accurately than standard life-cycle models can, observed household responses to anticipated, transitory shocks to income. One-asset life-cycle models with borrowing constraints generate clear predictions about which households will respond to anticipated transitory income shocks and about how strongly households respond to those shocks. Some predictions of one-asset models, however, are inconsistent with the household responses identified in recent empirical studies. In our two-asset model, households choose between investing in lower-return taxable accounts or higher-return, tax-deferred accounts with withdrawal penalties. Some households may optimally choose to invest in tax-deferred accounts, leading to a higher average consumption response that is closer to empirical observation. In this two-asset framework, the households that respond to anticipated transitory income shocks are not limited to very young or very old households. We evaluated our model by simulating the 2001 federal tax rebate and comparing the predicted with the observed response in terms of household consumption.

Tax rebates, such as those under the Economic Growth and Tax Relief Reconciliation Act of 2001 (EGTRRA) and the Economic Stimulus Act of 2008 are used to stimulate household consumption and economic activity. Such rebates also have been used by researchers as tests of households' willingness to use anticipated transitory income shocks to fund consumption. Policymakers have included tax rebates in fiscal stimulus programs in keeping with the assumption that households consume a significant fraction of temporary anticipated income, and the result will be greater economic activity and the mitigation of the effects of economic downturns.

The assumption that households spend a significant share of transitory shocks to income has been studied extensively since Modigliani and Steindel (1977), and validated by several empirical studies using household-level data on the 2001 and 2008 tax rebates. Agarwal, Liu, and Souleles (2007); Johnson, Parker, and Souleles (2006); Parker, Souleles, Johnson, and McClelland (2011); and Shapiro and Slemrod (2003a, 2003b, 2009) show that temporary fiscal stimulus, in the form of a check sent from the Internal Revenue Service (IRS), can boost household consumption. Moreover, Johnson et al. and Agarwal et al. report larger responses for household consumption

of nondurable goods among liquidity-constrained households. They report that the permanent-income hypothesis (PIH) does not always hold, especially for households with low holdings of liquid assets or households that have significant constraints on borrowing. Furthermore, neither group reports a particular distinction in the household response conditioned on age.

The empirical findings are difficult to reconcile with the traditional Aiyagari (1994) model or with a one-asset life-cycle model with a borrowing constraint without adding something like the Campbell and Mankiw (1990) “rule of thumb” consumers. In the Aiyagari model, infinitely lived agents face borrowing constraints, but most households optimally choose to save enough to smooth their consumption over time. The share of liquidity-constrained households under common parameterizations of the Aiyagari model is generally too low to generate the estimated consumption responses to the tax rebate. One-asset life-cycle models with heterogeneous households and borrowing constraints also cannot reproduce observed household consumption out of expected transitory shocks to income. Most households optimally set aside large precautionary or life-cycle savings and therefore are unwilling to spend a significant share of any expected and transitory increase in income, such as the 2001 tax rebate. The agents with the highest marginal propensity to consume (MPC) out of an additional unit of income tend to be those who are very impatient.¹ Some of the more impatient agents in a life-cycle model may have lower-than-typical survival probabilities, while others face very steep income profiles. Therefore, such models predict that most of the response to transitory shocks to income will occur among agents who are either very young or very old.

Our model offers a different mechanism to capture the observed MPC. We assume that the forward-looking agents in our model are subject to idiosyncratic income shocks and have access to separate tax-deferred and taxable assets, which provide different rates of return. Households can freely increase their holdings in either account; however, only the taxable accounts offer the household costless withdrawals. Some households might be liquidity constrained after a negative shock occurs because of the previously optimal choices to invest in illiquid tax-deferred assets. Those liquidity-constrained households are willing to spend a larger share out of anticipated

¹Carroll (2001) notes that impatient households, such as those with steep income profiles, are more likely to have a higher MPC than are those with shallower income profiles.

transitory shocks to income, thus driving the increase in the overall MPC. In addition to raising the MPC out of the transitory income shock, the friction in the model induces a change in the distribution of households that respond to anticipated income shocks. In this framework, more middle-aged households respond to the tax rebate than they would in a single-asset life-cycle model. By introducing the tax-deferred and taxable assets, we can reconcile some of the empirical observations with a life-cycle model in which households are forward-looking.

The importance of the distinction between investment in tax-deferred and taxable accounts has been studied by Amromin (2008); Bergstresser and Poterba (2004); Gomes, Michaelides, and Polkovnichenko (2009); and Shoven and Sialm (2004), among others. The distinction is important for household financial decisions: Although the tax-deferred accounts produce returns that are superior to those from identical assets held in taxable accounts, the investments in the tax-deferred assets cannot be liquidated without cost. Households make investments in tax-deferred retirement accounts through individual accounts like individual retirement accounts (IRAs) or employee-sponsored retirement accounts (such as Keogh, 401(k), and defined-benefit pension plans). And they can save in other tax-deferred instruments—such as the state-sponsored 529 education savings plans—that share many features with tax-deferred retirement accounts. The value of other tax-deferred assets, such as life insurance, can increase without triggering tax liabilities. Although there are a few exceptions, withdrawing assets from such accounts generally requires that the holders pay substantial penalties or income taxes. Marginal contributions to the accounts and marginal purchases of the assets costs little. Reversing the transactions, however, comes at a high marginal cost.

Despite the obstacles to withdrawing assets from the tax-deferred accounts, household-level data from the Survey of Consumer Finances (SCF) show that, in fact, most chose to hold a large share of total wealth in tax-deferred assets. The empirical data suggest that households optimally trade higher returns for lower liquidity, reducing their precautionary savings against short-run negative income shocks.

By incorporating similar transaction costs for withdrawing money from tax-deferred accounts into our life-cycle model with idiosyncratic income shocks, we endogenously generate a segment of households that has a higher MPC out of their anticipated tax rebates, particularly among

middle-aged households in the flatter portions of their lifetime income profile. We find, under a variety of scenarios, that incorporating a withdrawal penalty typically increases by 4 percentage points the marginal propensity to consume. In the benchmark scenario of a model incorporating the transaction cost, households consume about 21 percent of the tax rebate upon receipt. However, the ability of tax rebates to stimulate household consumption depends significantly on households' patience, rates of return on savings, and tax rates on capital income. Higher tax rates on capital income and higher interest rates cause households to boost their investments in illiquid tax-deferred assets as the returns from those assets increase substantially. Households with zero or very small holdings of liquid assets, regardless of total savings in illiquid assets, are more likely to behave like borrowing-constrained households when facing negative income shocks; specifically, they tend to spend a substantial fraction of the tax rebate. These predictions are more consistent with observed household-level data than are the predictions generated by a single-asset life-cycle model.

In this paper, we consider household portfolio choice over financial assets, and we choose not to model housing as a separate state and choice variable. Even though real estate shares some characteristics with other tax-deferred assets, the structure of any transaction costs differs substantially from most tax-deferred assets. Frictions in the housing market may generate similar results; agents making infrequent transactions in real estate may find it optimal to constrain themselves, producing a similar increase in the MPC out of anticipated transitory income shocks. Our goal for this paper is to show how widely available, commonly held, and easily transferable tax-deferred financial assets can affect optimal household savings and generate a higher MPC out of a transitory income shock.

In the next section, we review the literature in this area. In Section 3, we describe the empirical experiment that is replicated in our model. In Section 4, we describe the SCF data that motivate our modeling assumptions. In Section 5, we present the model and the calibration. In Section 6, we report the results. Section 7 reports our conclusions.

2 Literature Review

Many authors have analyzed predictable changes to household income and used data on household expenditures to test the permanent income hypothesis. Parker (1999) and Souleles (2002) consider changes in tax withholding to measure how changes in income affect consumption. Specifically, to demonstrate that changes in household-level consumption behavior are correlated with changes in withholding, Parker examines predictable changes in Social Security withholding. Souleles (1999, 2002) uses changes in withholding that resulted from income tax refunds and from the Reagan-era tax cuts. Instead of using changes in taxes, Stephens (2003) uses the timing of the issuance of Social Security checks—well-known by recipients—to evaluate whether the recipients’ spending patterns are based on the timing of benefits. Stephens reports that households tend to increase spending, and “the increase is sharpest on the day of the check arrival and is concentrated amongst households for whom Social Security is the primary source of income.”

Shapiro and Slemrod (2003a, 2003b, 2009) take a different approach. They rely on questions in the University of Michigan Survey Research Centers Monthly Survey, also known as the Survey of Consumers, to identify the effects of the 2001 and 2008 tax rebates. In almost all cases, households reported that they expected to spend a statistically and economically significant fraction of the tax rebate on consumption.

More recently, Agarwal et al., Johnson et al., and Parker et al. examined the effects of the 2001 and 2008 tax rebates by exploiting the random timing of the rebate. In each case, the rebate check was mailed on a date determined by the second-to-last digit of a tax filer’s Social Security number. All three studies use the random timing of the rebates to identify household spending, allowing them to separate the effects of the tax rebates from other economic conditions or unobserved variables. Johnson et al. and Parker et al. examine household spending patterns; Agarwal et al. examine credit card balance data. All three groups reject the PIH. Many studies also indicate that liquidity-constrained households tend to spend a greater share of the tax rebate on nondurable goods.

Our life-cycle model builds directly on those empirical studies. Frictionless life-cycle models

and infinite-horizon models of the kind described by Aiyagari cannot replicate the positive contemporaneous household response to a tax rebate. Moreover, life-cycle models with standard borrowing constraints and typical income profiles cannot reproduce empirical results. Carroll (2001) notes that households with higher income profiles are more impatient and more likely to spend out of a temporary anticipated shock. When frictions are not introduced in our model, we find that this group of consumers drives most of the observed response. However, in reality, this is not the only group that responds positively to a tax rebate when it is received. By introducing a household choice between taxable and tax-deferred assets, we can show that different groups of consumers spend some fraction of the rebate.

The choice to explicitly model tax-preferred and taxable assets builds on the literature on portfolio choice. Gomes et al. evaluate the impact of tax-deferred retirement accounts on life-cycle decisions over consumption, wealth accumulation, and savings. Longstaff (2009) evaluates the asset-pricing implications of having different types of illiquid and liquid assets that share some characteristics with our taxable and non-taxable assets. Dammon, Spatt, and Zhang (2004) study the problem of optimal asset location and allocation in the presence of taxable and tax-deferred accounts. Bergstresser and Poterba use the SCF to analyze households' allocation of funds between taxable and tax-deferred accounts.

3 Empirically Evaluating the 2001 Tax Rebate

The Economic Growth and Tax Relief Reconciliation Act of 2001 was enacted to counteract the 2001 economic recession. One of the law's provisions altered the lowest nonzero marginal tax rate: The 15 percent tax rate was reduced to 10 percent for the first \$6,000 of taxable income for a single filer, for the first \$10,000 of income for a head of household, and for the first \$12,000 of income for joint filers.²

In lieu of applying the 10 percent marginal tax rate at filing time for the 2001 taxes, over the course of a 10-week period ending in September 2001, the IRS mailed checks to taxpayers

²Other provisions of EGTRRA included the gradual reduction of other income tax rates. The 28 percent, 31 percent, 36 percent, and 39.6 percent brackets were reduced between 3 and 4 percentage points. The limitation on itemized deductions was similarly phased out, being removed completely for the tax year 2010.

to stimulate consumption and economic recovery.³ As described by Johnson et al., each week the checks were mailed to a different cohort of households on the basis of the second-to-last digit of the filer's Social Security number. Hence, the timing of the tax rebate was random and orthogonal to other household features.

The checks were worth a maximum of \$300, \$500, or \$600 for filers falling into the categories of single, head of household, or married filing jointly, respectively. As described by Shapiro and Slemrod (2003a, 2003b), the IRS determined the amount of the rebate on the basis of the 2000 tax return, and any rebate in excess of that to which the filer was entitled on the 2001 return was forgiven. Any shortfall in the rebate could be claimed when the 2001 return was filed in 2002. According to Johnson et al., the total value of the rebate was a little under \$40 billion, roughly equal to 1.5 percent and 2.2 percent of total contemporaneous quarterly gross domestic product and personal consumption expenditures, respectively.

In their experiment, Johnson et al. exploit the exogenous and random timing of the tax rebate to estimate household consumption responses. Using questions in the Consumer Expenditure Survey, they identify the period in which households received the rebate and the amount of the rebate. They estimate linear relationships in which the change in various categories of consumption are regressed against the amount of the rebate and several other variables, and they report that the average fraction spent on nondurable goods by the household equals 0.373 (see Table 1).⁴ A second regression includes only those households that received the rebate; the fraction that group spent upon receipt is 0.247. That smaller value is closer to empirical results found for the 2008 tax rebate. In addition to the benchmark model, Johnson et al. also introduce other interactions, including asset holdings, income, and age. They report that the low-income and low-asset groups tended to spend more of the rebate. Households in the low-income group, with annual income below \$34,298, spent a statistically and economically significant 76 percent of the rebate on nondurable goods. The fraction of the rebate spent by middle-income

³See Joint Committee on Taxation, www.jct.gov/publications.html?func=startdown&id=2003, for a summary of EGTRRA's provisions.

⁴The nondurables category differs from the strictly nondurables by the inclusion of some semidurable goods, such as apparel, household health care expenditures, and reading materials. See Lusardi (1996) for more information on how Johnson et al. construct spending categories.

or high-income groups cannot be statistically differentiated from zero.⁵ Households with low taxable assets, defined as less than \$1,000, spent more of the rebate in the quarter of receipt than did either the middle- or the high-asset group. The low-asset group increased spending on nondurable goods by about 63 cents for every dollar of the rebate received.⁶

Hamilton (2008) reexamined the findings of Johnson et al. and shows that small changes in the analysis cause significant changes in the results. First, he modified the sample by dropping the top and the bottom 10 observations, which show excessively large changes in durable consumption. The new estimates resulting from that restricted sample are significantly smaller. The average fraction of the rebate spent on the consumption of nondurable goods upon receipt drops from 0.373 in Johnson et al. to 0.226 in Hamilton.⁷ Moreover, the estimates also are sensitive to the regression technique used. Hamilton carried out a median regression on the entire sample, and the fraction of the tax rebate spent on nondurable goods dropped to a statistically significant 0.126 in the quarter in which the rebate was received and a statistically significant 0.292 for the cumulative, two-quarter response.

We extended the previous analysis in several ways. Following Hamilton, we restricted the sample by dropping the outlier observations. When only those households that receive the rebate are included, the average fraction of the rebate spent drops from the 0.247 reported in Johnson et al. to 0.093. Removing the outliers has little impact on any other results, however. The results in samples restricted by income or assets are effectively identical to those reported by Johnson et al. Households with low income and low liquid assets still exhibit a strong, statistically significant consumption response upon receipt of the rebate.

3.1 Household Response to Tax Rebate by Age

Johnson et al. tested whether the marginal propensity to consume varies with age. A life-cycle model would predict that young households, with low liquid wealth and high income growth,

⁵The middle-income group contains households above the \$34,298 cutoff for the low-income group and below \$69,000. The high-income group contains households with incomes over \$69,000.

⁶Agarwal et al. suggest that consumption increases, upon tax rebate receipt, mostly for liquidity-constrained households. Instead, liquidity-unconstrained households prefer to increase their savings. Agarwal et al. conclude that liquidity constraints are the main driving force of the consumer response to “lumpy” changes in income.

⁷Hamilton reports that this estimate is relatively insensitive to the number of outliers removed.

are relatively impatient and are expected to have a high MPC. Older households with lower survival probabilities also are impatient and may also have a higher MPC relative to middle-aged households. Johnson et al. tested those hypotheses by estimating regressions in which they interacted an age variable with the rebate to discern differences in the MPC out of a transitory shock among three subsamples of households, segregated by age. They report no distinction in response to the tax rebate among the ages, however, and thus fail to reject the null hypothesis that age is not a determining factor in household response to the income shock.

However, after comparing the fraction of the rebate consumed by households of different ages, Johnson et al. report that young (younger than 40) and old (older than 55) households spend a little more of the tax rebate relative to middle-aged households but that the differences are not statistically significant. Parker et al. evaluated the consumption response in a similar experiment with the 2008 tax rebate. Their results were similarly inconclusive, and they were unable to reject the null hypothesis that the households exhibit similar behavior.

Examining the same set of issues, Shapiro and Slemrod (2009) evaluated responses from the Survey of Consumers. That survey includes direct questions about whether a household would mostly spend, mostly save, or use the rebate to pay off debt. The authors identify a correlation between the age of the household and a report that they would mostly spend the rebate. Specifically, older households are more likely than the other age groups to plan to spend the rebate. About 28.4 percent of older-than-65 households reported that they would mostly spend the rebate; only about 17 percent of the younger-than-65 households said that they would mostly spend the rebate.

Using the credit card data, Agarwal et al. evaluated the consumption response to the 2001 tax rebate. They report that the long-run cumulative response is a monotonic decreasing function of age. Younger households consume a larger fraction of the rebate relative to middle-aged and older households. The marginal and cumulative coefficients for spending are jointly significant if the three age groups (young, middle, old) are considered at the same time.

Despite those results, however, it is difficult to find agreement among the studies about whether the MPC varies with age. The one feature common to all, however, is that middle-aged households do not report uniformly spending less of their transitory and expected income than

the young and old households.

4 Data on the Household Portfolio

To evaluate our model and assess its implications, we explored household financial data in the SCF. We focused on the household-level decision to hold assets in tax-deferred, as opposed to taxable, accounts.

Taxable assets include those on which a yearly tax is assessed on the income earned. Following the work of Bergstresser and Poterba, we classify taxable assets as those including money held in savings, money market, and call accounts; certificates of deposit; directly held mutual funds; stocks; bonds; and savings bonds. Cash and checking accounts are not included because neither earns significant income or incurs more than minimal taxes. Alternatively, assets can be held in tax-deferred accounts, which are characterized by growth and income that does not trigger yearly taxes on capital gains, income, or distributions. Withdrawing assets from tax-deferred accounts, however, typically triggers a large income tax assessment, substantial financial penalties, and potentially large transaction costs.⁸ Some penalties and transaction costs are written into the U.S. tax code. Others, such as those involving life insurance, are written into the contracts signed at the time of investment. Tax-deferred assets include defined-contribution retirement accounts, defined-benefit pension plans, whole-life insurance plans, and other nonfinancial assets (see Table 2).

We classified the sample of households in the 2001 SCF into four groups. The first, consisting of 16.2 percent of the population, holds neither tax-deferred nor taxable assets. The second, 8.9 percent of the U.S. population, includes households that invest in tax-deferred assets only. Asset holdings by the second group are described in Tables 3 and 4. The third group, 18.6 percent of

⁸Some employers allow employees to borrow money from a 401(k) account. But according to Li and Smith (2008), in 2004, only 16 percent of eligible households—those with access to 401(k) plans that permit withdrawals—had taken out such loans. Obstacles to borrowing money from the 401(k) accounts, such as origination fees, timeliness, minimum and maximum loan sizes, and other expenses could account for the relatively low usage rate. Typically, there is some significant, fixed loan origination fee that, for small loans, could represent a substantial fraction of the principal. Furthermore, the household must pay interest on the loan, which ends up being taxed twice (see Li and Smith). The term of such loans typically is short—between one and five years—and, if the borrower defaults, the household must still pay any taxes and penalties that would have been assessed for an early withdrawal.

the population, invests only in taxable assets and has no tax-deferred accounts (see Tables 5 and 6). The majority of the U.S. population, 56.3 percent, includes households that hold taxable and tax-deferred assets (see Tables 7, 9, 8, and 10).

In each group, there are many households with small holdings of assets—the kinds of households identified by Johnson et al. as more likely to spend a tax rebate. Among those with tax-deferred assets only, young households that belong to the first quartile in earnings hold very few assets (Table 3). Middle-aged households hold larger amount of assets; however, those assets are not easily accessible, because penalties are charged upon withdrawal. A large share of households with taxable assets alone has almost no assets saved (see Table 5). Through age 45, half of that group has less than \$1,200 in taxable assets, although the availability of taxable assets increases substantially as the households approach retirement. The distribution of assets in that group shows that a large fraction of the population of households that save only in taxable assets has very little savings to draw from in the event of a negative income shock.

The households that hold taxable and tax-deferred assets tend to have more assets than the other groups that hold only one type of asset (see Tables 7 and 9). Nonetheless, the value of taxable assets held by a sizable fraction of that group is still quite low early in life. More than one quarter of the households that are younger than 30 years old have less than \$1,000 in taxable assets. Of the households observed in the empirical studies of the 2001 federal tax rebate, those are the types that were more likely to spend a tax rebate.

The SCF also provides more detailed disaggregated data on household asset holdings. (See Table 11 for the U.S. average share of each asset by age, weighted by the sample weight. The taxable assets are the first six categories listed; the tax-deferred assets are the others). Households hold a smaller quantity of taxable than tax-deferred assets as a share of total assets. (Table 12 shows average asset holdings as a function of total household financial assets.) For households with total assets of less than \$10,000, the share of total wealth held in taxable assets varies between 1 percent and 99 percent. The share decreases for households with total assets between \$10,000 and \$300,000. For households with the most financial assets, \$300,000 and higher, the share increases slightly as the households increase holdings of stocks and bonds.

5 The Model

We built a life-cycle model with portfolio choice over taxable and tax-deferred assets. The economy is populated by a continuum of households, and the state variables are household age j , stock of tax-deferred assets k^d , stock of taxable assets k^t , and labor productivity status θ . We also include an additional exogenous and determined state variable s that represents a Social Security number and does not change over time. Including this additional state allows us to simulate the experiment in Johnson et al. in which multiple cohorts of households receive tax rebates at exogenous, predetermined times. Then, we constructed moments that are close to those estimated by Johnson et al. using data from the Consumer Expenditure Survey, and we evaluated the effect of tax-deferred assets on household elasticity with respect to predetermined transfers.

After receiving news of their labor productivity status, households purchase consumption and invest in the two assets. The household's problem is as follows:

$$V(j, k^d, k^t, \theta, s) = \max_{c, k^{d'}, k^{t'}} u(c) + \beta E_t S(j+1) [V(j+1, k^{d'}, k^{t'}, \theta', s)] \quad (1)$$

The household's choices are consumption c , next-period tax-deferred assets $k^{d'}$, and next-period taxable assets $k^{t'}$. The function $S(j)$ equals the likelihood that the individual survives to age j conditioned on survival at age $j-1$. The utility function $u(c)$ takes the following form:

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

The household is subject to the borrowing constraints

$$\kappa k^{d'} \geq 0; k^{t'} \geq \underline{b}$$

Furthermore, the household is subject to a budget constraint

$$\begin{aligned} c(1 + \tau^c) + k^{d'} + k^{t'} + \phi_t(j, k^{d'} - k^d) \leq \\ (1 + r(1 - \tau^{kt}))k^t + (1 + r)k^d + w\theta(1 - \tau^l) + tr(j, e) \end{aligned} \quad (2)$$

The after-tax consumption, next-period tax-deferred assets, next-period taxable assets, and penalties must be no greater than the rental earned on capital, net of taxes, stock of capital, wage income, and government transfer. We assume that government transfers, such as tax rebates, are a function of age and earnings $tr(j, e)$. Labor earnings are taxed at the effective tax rate τ^l , which includes labor income and the payroll tax rate. Taxable assets are taxed at the effective tax rate τ^{kt} . r is the real return on household investments.

The lifetime labor earning process evolves as follows:

$$\ln \theta_t = f(t) + z_t + \epsilon_t$$

where

$$z_t = \rho z_{t-1} + \eta_t$$

$f(t)$ is an age-polynomial, which represents the typical life-cycle earning profile and is entirely known to the individual at the time of entry into the labor force. In addition to this deterministic component, we introduce two random components: a persistent component z_t that evolves as an autoregressive process with persistence ρ , and a transitory component ϵ_t . η_t and ϵ_t are i.i.d. with variance σ_η^2 and σ_ϵ^2 , respectively. We follow Cocco, Gomes, and Maenhout (2005) for the parameterization of the deterministic component $f(t)$ and Storesletten, Telmer, and Yaron (2004) for the parameterization of the stochastic idiosyncratic labor earning shocks.

5.1 Tax-Deferred and Taxable Assets

Taxable and tax-deferred assets differ principally in the tax treatment and in households' ability to liquidate their positions (Table 2). We model those differences by specifying different tax rates for the two types of assets and by including a penalty function in the household budget

constraint.

In the model, any income earned on funds invested in taxable assets is taxed in the year it is received. Even though interest, dividends, and capital gains might be subject to different tax rates, for simplicity, we apply a uniform linear tax rate τ^{kt} to the income derived from this class of assets.

Funds invested in tax-deferred assets k^d grow at an after-tax rate of return r . Taxes on the income earned on those assets are not paid as long as the earnings remain in the accounts. Instead, the household pays taxes on earned income at the time of withdrawal. However, some companies pay taxes on their profits after deducting interest payments on new investments and depreciation, therefore reducing overall household earnings. That feature of the model reflects the U.S. tax treatment of those assets. For example, an IRA can generate a stream of dividends and capital gains that, as long as they remain in the account, are untaxed at the individual level. Although the tax treatment differs slightly across the set of tax-deferred assets listed in Table 2, we impose a linear tax τ^{kd} on the income derived from such assets.

Tax-deferred assets are substantially more illiquid than taxable assets. We address the distinction by applying an asymmetric penalty function $\phi_t(j, k^{d'} - k^d)$, which depends on household age j and the household's change in the tax-deferred asset holdings. The age affects the penalty function in that many of the penalties for drawing down tax-deferred assets decline substantially upon retirement. Before age 59-1/2, individuals can contribute to their IRAs and 529 education savings accounts or increase their amount of life insurance with practically no transaction costs and in increments that are approximately continuous. Even though there are limits in the maximum amounts that households can contribute to some tax-deferred assets, those limits vary among assets. We choose not to explicitly model the maximum amount that can be invested in those assets, as this constraint is never binding in the model at the high end of the range. Even though households may reach the maximum that can be invested in retirement accounts, they still have the option of investing in other tax-deferred accounts, such as life insurance. On the other side, there is generally no penalty for increasing one's liquid position before retirement.

Working-age people who want to divest their portfolios of tax-deferred assets confront several costly obstacles. Retirement funds and education savings accounts are taxed and penalized for

early withdrawal. Thus, we consider a penalty function that imposes a significant marginal cost when reducing the stock of these assets. The benchmark penalty equals 30 percent of the decline in assets and represents the transaction costs, the 10 percent penalty for early withdrawal, and taxes accompanying a decline in assets. After retirement, households do not pay the 10 percent early-withdrawal penalty, and the overall cost of withdrawal equals 20 percent of the decline in assets. In the sensitivity analysis (Section 6.2), we select penalties that range from 20 percent to 40 percent of the asset change. On the other side, after retirement, households often must withdraw funds from tax-deferred assets. For example, after age 70, a person with an IRA must start withdrawing funds to avoid paying a penalty. We do not explicitly model that after-retirement penalty because, in our model, it is optimal for a household to start withdrawing assets. It follows that this simplification does not affect the results.

As described above, the penalty function is continuous but not differentiable. Therefore, we approximate it with a Chen-Mangasarian smoothing function that approximates the properties of the penalty, but allows the function to remain continuous, differentiable, and convex:

$$\phi_t(j, \Delta k) = \xi(j)[(-\Delta k - \zeta) + \log(1 + \exp[a(\Delta k - \zeta)])/a] \quad (3)$$

The parameter a defines the curvature of the function around the kink. Higher levels of a push the Chen-Mangasarian smoothing function closer to the value of the true “kinked” function. The point at which the penalty begins to take effect is defined by ζ . Because many retirement accounts and other tax-deferred assets appreciate in value but cannot be withdrawn, we set the value of ζ to be slightly positive.

5.2 Parameterization

The model’s parameters are chosen to mirror the main features of the U.S. economy, and they are consistent with the related literature (Table 13). In the benchmark experiment, the subjective discount factor β equals 0.957, the coefficient of relative risk aversion γ equals 2, and the minimum net asset position equals 0. The parameters for the age-polynomial that describes the life-cycle component of the income shock are taken from Cocco et al. and are obtained using a

sample of the population with a high school diploma. The persistence and variance of the persistent shock to the income process is taken from Storesletten et al. and subsequently employed by Fernandez-Villaverde and Krueger (2011). The after-retirement income is assumed equal to the product between the average income in the last working period and the replacement ratio. The replacement ratio is assumed to equal 0.6 (Bernheim, Skinner, and Weinberg, 2001).

In the model, we assume that the pretax real gross return on assets is 5 percent. However, corporations and businesses pay taxes on net profit, before the household receives any investment income. To take into account the taxes on corporate and noncorporate business income, we parameterize the interest rate r to about 4.3 percent, reflecting an average effective business tax rate of 13.4 percent in 2001. That rate is obtained as the weighted average of the tax rates on corporate and noncorporate business income, under the simplification that all capital income from businesses can be placed into tax-deferred accounts. It follows that the household receives about a 4.3 percent return on its investments. The return on capital income is slightly below the interest rate calculated by taking the reciprocal of the time value of money, $\frac{1}{\beta}$. The tax rate on the taxable assets, τ^{kt} , is chosen such that, together with the interest rate r , net of corporate taxes, the implied total tax burden on taxable assets is 35.2 of capital income.⁹

In the following simulations we compare the model with taxable and tax-deferred assets against a model with a single liquid asset. In the latter model, the tax rate on liquid assets has been computed to match the total savings in the two-asset model.

6 Simulation Results

To conduct our tax rebate experiment, in the initial period we simulated a sample of two cohorts of households that differ by Social Security number. One cohort receives the rebate in period three; the other receives it in period four. In period one, each household believes with certainty that current economic conditions will persist. In period two, each household updates its beliefs to include the timing of the receipt of the tax rebate; its amount; and any future economic conditions, such as reduced wages, increased likelihood of negative income shocks, and reduced

⁹Source: Staff estimates from the Congressional Budget Office.

return on capital.

We generated a random sample of households using the 2001 U.S. population distribution by age and survival probabilities, so that our experiment is comparable with that performed by Johnson et al. We estimate the following equation:

$$\Delta c_{i,t} = \gamma_0 + \gamma_1 \text{rebate}_{i,t} + \zeta_{i,t} \quad (4)$$

where $\text{rebate}_{i,t}$ is the amount received by household i during period t . The dependent variable is the change in consumption $\Delta c_{i,t}$ for household i during period t . We use this equation to identify the simulated consumption response to the tax rebate, which corresponds to the real consumption response to the 2001 tax rebate as identified by Johnson et al. We also can separately identify other effects, such as an identification effect, by constructing a counterfactual simulation in which households do not receive a rebate.

We modeled the tax rebate as a temporary increase in lump-sum transfers, the timing of which depends on the household state s . Specifically, a tax rebate is distributed to households with state $s = 0$ in period t and to households with state $s = 1$ in period $t + 1$. We chose the benchmark value of the simulated tax rebate to be approximately the same as the 2001 tax rebate. Specifically, about 0.5 percent of total annual personal consumption expenditure in 2001. In the next section we present the benchmark experiment. In the sensitivity analysis, we show how the results are modified after changing the benchmark assumptions and the model's parameterizations.

6.1 Benchmark Case

We evaluated our simulated model against two features of household behavior over the life cycle.¹⁰ Figure 1 shows a simulated profile of consumption, labor income, and total assets of

¹⁰We considered five grid points for the persistent shock and three grid points for the transitory shock. The former were selected using the Tauchen (1986) approach, the latter using a quadrature. The household's value function is interpolated using 49-degree, 50-node Chebychev polynomials in each dimension. The number of nodes was selected so that the MPC from the simulated tax rebate was largely invariant to a change in grid points. The household policy functions are interpolated using bilinear interpolation. To solve the household's problem, we employed Fortran optimization software, *dpsol*, written by Yongyang Cai and Kenneth L. Judd.

an average household, without income shocks. For young households, between the ages of 20 and 30, consumption moves in tandem with current income, and little precautionary savings are accumulated to buffer against negative income shocks.¹¹ The rate of income growth slows as the household approaches 30, and the household begins to accumulate assets. Total assets peak a few years before retirement. As the household retires at age 65, income drops, but consumption does not fall suddenly at retirement; households use funds accumulated in retirement accounts and other tax-deferred assets to finance living expenses. Under our assumption that the household lives until 90 and does not leave any bequest, the household's assets are drawn down to zero before dying.

To illustrate how different households behave when they receive a tax rebate, we studied the consumption responses of selected households. Specifically, we charted the difference between the simulated consumption of a household that receives a tax rebate in period three and the simulated consumption of a household that does not receive a rebate (Figure 2). The difference is expressed as a share of the tax rebate.

The upper-left panel of Figure 2 shows the different responses of three 25-year-old households. The solid line shows consumption of a household with no financial assets. That household expects fairly steep income growth in the near future and its borrowing is constrained; it spends nearly all of the rebate upon receipt. The long-dashed line shows consumption of a household with tax-deferred assets but no liquid taxable assets. In contrast to the first household, the second household exhibits a small anticipatory response; namely, spending some of the rebate before receiving it. Moreover, the fraction of the rebate spent upon receipt is smaller than that spent by the first household. The dotted line shows consumption of a household with taxable and tax-deferred assets. That household has a small consumption response, which includes a small increase in consumption before the rebate arrived—the anticipation effect—and a small increase in the year after receiving the rebate to smooth consumption. Clearly, liquidity constraints make a significant difference in a young household's willingness to spend its rebate.

The upper-right panel of Figure 2 shows the different consumption responses of three 30-

¹¹Labor income, consumption, labor taxes, and capital taxes account for the wedge between income and consumption on the figure.

year-old households with financial characteristics similar to those described for the upper-left panel. Their consumption follows similar patterns; however, in contrast to the younger group, all households display some anticipatory response. Moreover, the fraction spent upon receipt is smaller. The first household, with no assets by the time the tax rebate is received but with some accumulated savings in previous years, chooses to spend almost all of the rebate upon receipt. The small anticipatory effect in period 2 in the graph is from a very small amount of money that would otherwise have been invested but that is now spent in anticipation of having the rebate in period three. The third household, with taxable and tax-deferred assets, spends a tiny amount of the rebate upon receipt. Instead, it increases consumption, mostly before receiving the rebate.

The bottom-left panel of Figure 2 shows consumption among 50-year-old households; the bottom-right panel shows consumption among 70-year-old households. The panels display similar behavior. The tax rebate makes a small positive impact on household consumption. Nearly all households choose to spend a small fraction of it, either before or after receipt. For those households, the anticipated growth in income is small, and they choose to smooth consumption over time. The penalty for selling tax-deferred assets drops substantially after age 65, as the obstacles to closing retirement accounts and other illiquid assets fall, and the distinction between households with taxable and tax-deferred assets narrows.

The first panel in Table 14 presents the estimation results for γ_1 obtained from Equation 4. In the benchmark experiment, households receive a rebate either in period three or in period four, and no change in wages, interest rates, or any other variables affects the household decision. The coefficient γ_1 can be interpreted as the average share of the tax rebate that households spend on consumption upon receipt. We compared the results obtained from a standard life-cycle model with only one liquid asset against those from a model with two assets, one taxable and one tax-deferred, when households face a penalty to withdraw their funds invested in the taxable asset.

In the standard life-cycle model with a single liquid asset, the coefficient γ_1 implies that households spend contemporaneously about 16.9 percent of the tax rebate on consumption. When households can invest in tax-deferred and taxable assets and also confront a penalty for

withdrawing funds from tax-deferred account, that share increases to about 20.7 percent. The increase of nearly 4 percentage points puts the simulated consumption response in line with that documented by Hamilton, Johnson et al., and Parker et al. However, the simulated consumption response is highly responsive to the model’s parameterizations.

In the context of our model, we can compare household consumption against a benchmark scenario in which there is no tax rebate. The difference between identical households in the two models can provide an alternative estimate of the MPC in the period of the rebate. Many households spend the tax rebate smoothly. They take that transitory income and, particularly if they are older and not liquidity constrained, spend the money over several periods. Comparing households in the tax rebate scenario against this benchmark, the no-rebate scenario allows us to identify additional spending generated by the tax rebate that may be generated but that is otherwise not identified by the regression. The total consumption response in the period of the tax rebate for a household receiving it compared with a comparable household that does not jumps 4 percentage points to about 20.7 percent. The regression identifies a household’s propensity to consume out of a contemporaneous transitory income shock, but as a policy, a tax rebate can generate a still larger increase in household consumption.

Figure 3 presents the percentage of households in each age group that contemporaneously respond to a tax rebate over the life-cycle, in the one-asset and in the two-asset model. We identify such households with the indicator function:

$$I((c_{1,t} > c_{2,t} + \epsilon) \bigcap (c_{1,t+1} + \epsilon < c_{2,t+1})) \tag{5}$$

The first subscript of the household’s consumption c is the Social Security number and the second is the period of time. We typically set ϵ to be about 5 percent of the tax rebate to screen out small positive responses that may result from the approximation methods used to solve the model. The solid bars in the figure indicate the percentage of each cohort of households that respond contemporaneously to tax rebates by increasing their consumption, as indicated in Equation 5, in the one-asset model. The open bars display the same information but for the two-asset model when the penalty to withdraw tax-deferred assets is positive.

The early-life and late-life behavior among households is fairly similar, both in the one-asset and in the two-asset model. Most households at the beginning of their life cycles spend nearly all of the rebate. At age 20, between 65 and 70 percent of households increase consumption upon receiving the rebate. Those households are liquidity constrained and can look ahead to steep improvements in their projected earnings.

By age 23, however, the household response to the tax rebate under the two models begins to diverge.

In the one-asset model, households begin to accumulate assets. Those holdings, although initially quite small, imply that households are no longer income constrained. Furthermore, the slope of the expected future earnings curve begins to decline, indicating that households have a smaller incentive to spend a large share of the tax rebate. By age 30, about 10 percent of the households respond contemporaneously to a tax rebate. That proportion drops nearly to zero by age 40, when almost none of the households contemporaneously respond to the tax rebate.

By contrast, in the two-asset model, even middle-age households are willing to spend some of the tax rebate. The open bars show that approximately 50 percent to 70 percent of young households (between 23 and 30 years old) and approximately 10 percent to 30 percent of middle-aged households (between 30 and 60 years old) spend some fraction of the tax rebate upon receipt.

Young households that have begun to invest in tax-deferred assets and have accumulated small amounts of wealth are likely to be liquidity constrained if a negative income shock occurs. Young households spend a significant fraction of the tax rebate also in the one-asset model, as their accumulated wealth is low and they may be liquidity constrained if a negative event occurs. However, in the one-asset model, the liquid assets accumulated are always larger relative to the two-asset model, implying that a smaller fraction of the rebate is spent to buffer consumption against negative income shocks.

Middle-aged households generally have accumulated more assets than the younger households have. However, some middle-aged households that have invested a large share of savings in tax-deferred accounts may find themselves liquidity constrained if hit by a negative income shock. Therefore, those households might find it optimal to consume some of a tax rebate

upon receipt rather than withdrawing funds from tax-deferred accounts and paying associated penalties. Middle-aged households' behavior in the two-asset model is very different from their behavior in the one-asset model (see Figure 3). In the one-asset model, the middle-aged households have accumulated enough assets that can be withdrawn without penalty. The households are almost never liquidity constrained, even when negative shocks occur.

After retirement, as the surviving households come closer to being liquidity constrained, they begin to spend part of the rebate. Although the onset of a positive consumption response differs slightly, the qualitative features of household behavior are similar in the two models.

Figure 4 shows the percentage of sample households that respond contemporaneously to a tax rebate. The sample is constructed to match the U.S. household distribution by age in 2001. For example, the 20-year-old households that respond to the tax rebate represent about 1.2 percent of the total population. Although a large share of older households spend the tax rebates, that group is not the main one that accounts for the increase in the observed γ_1 . Most of the increase can be accounted for by households in the 20- to 60-year-old group.

The household response to the tax rebate, conditional on holding liquid taxable financial assets, is illustrated by Figure 5. The households that spend a share of the rebate are predominantly those without liquid financial assets. This result mirrors the empirical findings by Agarwal et al., Hamilton, and Johnson et al. The liquidity constraints are a driver of consumption dynamics. In the sample population, about 60 percent of households with zero liquid assets spend the rebate upon receipt, whereas only about 10 percent of households with some liquid assets do so.

Households with no financial assets display a high response to the tax rebate—nearly 100 percent spend the rebate (see Figure 6). The fraction spent declines slowly with the increase in holdings of financial assets. The wealthiest households spend almost none of the rebate, although in any decile there are households that respond to the rebate by increasing consumption. This finding stands in sharp contrast to the results obtainable from the one-asset model.

Household consumption changes as income changes (see Figure 7), with most of the response occurring in the first five quantiles. The lowest quantiles include the youngest households, who are liquidity constrained and face an upward-sloping income curve. The consumption response

is high among those households. As income increases, the share of the tax rebate spent on consumption decreases. Heterogeneity in current and expected labor income explains the non-monotonic consumption response for the higher quantiles; nonetheless, for the higher quantiles the consumption response is always smaller than for the lower quantiles.

6.2 Sensitivity Analysis

To identify possible responses of households to the receipt of a tax rebate, we compared three scenarios: a one-asset model, a two-asset model without a withdrawal penalty, and a two asset-model with a withdrawal penalty.

The results obtained in the one-asset model and in the two-asset model without withdrawal penalty are similar, for two main reasons. First, in each model, the withdrawal penalty equals zero. Second, in the one-asset model, the tax rate on liquid assets, even though slightly different in the various experiments, is overall just a little higher than that applied to tax-deferred assets in the two-asset model without a penalty. That rate is computed to match the total savings obtained in the two-asset model with a withdrawal penalty. Therefore, in the two-asset model without a penalty, households optimally choose to invest only in the tax-advantaged assets, as those assets will result in lower tax assessments.

In the benchmark two-model asset with a positive withdrawal penalty, households consume, on average, about 21 percent of the tax rebate upon receipt; additionally, their consumption out of the tax rebate is about 4 percent larger than in the other models. In this section, we modify some of the baseline assumptions and parametrization choices to evaluate the sensitivity of our benchmark results and of the spread (see Table 14).

A household's behavior is relatively insensitive to the size of the penalty; the MPC drops slightly as the penalty increases, encouraging some households to substitute holdings in taxable assets from the less-liquid tax-deferred accounts. The existence of a penalty above a threshold is enough to generate the household's response. In addition, the household response is almost unchanged even if the 10 percent early-withdrawal penalty must be paid until age 50 instead of age 60. This follows from the more active consumption exhibited by the young households

relative to the close-to-retirement households.

Instead, changes in the capital income tax rate make a larger difference in households' willingness to hold taxable assets. Lowering the tax rate reduces the payoff for holding assets in tax-deferred accounts. Households are more willing to hold savings in a liquid taxable form, which allows them to smooth their consumption out of transitory shocks to their incomes more easily. The consumption response to the tax rebate is less than in the benchmark case. Conversely, higher tax rates give a greater incentive to households to trade risk for tax-deferred status, and both the households' holding of tax-deferred assets and the consumption response to the tax rebate are larger. If tax rebates are distributed in a period in which the Bush-era tax cuts are not yet expired, the lower tax rate on capital income implies a smaller consumption response.

The amount of the rebate affects the consumption response in the no-penalty and penalty scenarios. In each case, increases in the rebate cause households to spend a smaller share upon receipt. The desire to smooth consumption induces households to save part of a larger transfer and spread consumption across several periods.

Changes in the discount factor β affect the penalty scenario and the no-penalty scenario. When β increases to 0.98, without changing the interest rate, households are more patient and prefer to hold more tax-deferred assets. They choose to consume less of the rebate in any model. Without a penalty for withdrawing tax-deferred assets, the more patient households consume only a negligible fraction of the rebate. In contrast, in the face of a withdrawal penalty, those households consume a larger fraction of the rebate. The spread between the two scenarios is significant, as the marginal propensity to contemporaneously consume out of a tax rebate almost doubles. When β decreases to 0.94, households are more impatient, tilt their portfolios optimally toward taxable assets, and choose to consume a significant fraction of the rebate, with or without a penalty. The spread between the penalty and no-penalty scenarios is significant, even though any additional response generated by the introduction of the penalty is mitigated somewhat by the household's impatience.

A larger replacement ratio implies that after-retirement income is a larger fraction of the last-working-year income, compared with the benchmark case. Therefore, households would have to

save less for retirement and consume a larger share of the rebate upon receipt. On the other side, a smaller replacement ratio implies a lower marginal propensity for consumption among working-age households, which must save more to support consumption in retirement.

To test the effect of economic conditions on households' MPC, we modeled a short and a long recession as negative productivity shocks, causing a two-year or six-year decline in wages and interest rates. In each experiment, the marginal propensity to consume decreased with respect to the benchmark scenario. The longer the recession, the larger the decrease in the MPC. Consistent with previous experiments, there is typically a spread of 3 to 4 percentage points between the rate of spending out of those tax rebates, with and without penalties for withdrawal.

6.3 Employee Matching Contributions

We evaluated one developing feature of employer-based retirement plans—the employer matching contribution. We extended our penalty function to approximate this feature and identified the difference between the standard penalty function and a hypothetical one reparameterized so that positive contributions to tax-deferred accounts are matched to a small extent by the employer (see Figure 8). In the benchmark scenario, the penalty paid by households is positive when Δk is negative—when households withdraw assets from the tax-deferred accounts—and zero otherwise. Instead, with matching contributions, the households receive a benefit from further contributions to the tax-deferred accounts, up to some age-specific limit that we define as $m(j)$. In the figure, the upper dashed line reflects the originally formulated penalty function. The vertical dashed lines delimit the approximate area in the scenario with employer matching in which the household has access to the match. The matching-contribution limit is modeled by allowing the penalty function to decrease up to some positive value for $-\Delta k$, and then the function keeps a constant value.

The penalty function in Equation 3 is reformulated to include the maximum allowable match $m(j)$:

$$\phi_t(j, \Delta k) = \xi(j)[(-\Delta k - (\zeta + m(j))) + \log(1 + \exp[a(\Delta k - (\zeta + m(j)))])/a - m(j)]$$

(shown in Figure 8 as the solid line). In this function, $m(j)$ is the maximum matching contribution, and the match depends on age.¹² In the scenario with a 30 percent penalty, for every dollar that goes into the matching-fund account, 30 cents is refunded to the household. Therefore, every additional dollar in the tax-deferred account costs a net 70 cents. The effective employer match would be $\frac{\xi(j)}{1-\xi(j)}$, or about 43 percent of the household contribution up to the matching-contribution limit.

Table 14 shows the effects of progressively larger values of $m(j)$. The matching contributions have two competing effects: They increase savings, and they increase the proportion of assets held in tax-deferred accounts. From the smallest, $m(j) = 0.0025$, to the largest, $m(j) = 0.03$, overall household savings increase by approximately 50 percent. The returns on investment in tax-deferred accounts also increase substantially. The result is a decrease in the consumption response to a tax rebate as the value of $m(j)$ increases, because overall savings are so much greater for the same tax and interest rates. However, savings in taxable accounts actually decrease between the scenarios, suggesting that households invest an even larger fraction of their money in tax-deferred accounts than would otherwise be invested in taxable assets.

6.4 Catastrophic Income Shocks

Amromin (2008) studied the household portfolio choice over taxable and tax-deferred assets using SCF data. Following Amromin (2003, 2008) and to provide an additional sensitivity test, we introduced some amount of catastrophic income shock, which we modeled as a transitory income shock, occurring with a 1 percent probability and generating a 90 percent drop in income.

When the shock is introduced, household MPC and household behavior change notably. Specifically, the household MPC drops significantly, as households prefer to increase precautionary savings (see Table 14). In this scenario, the precautionary-savings motive drives the dynamics of the model. Households face a trade-off between the wish to hold most savings in tax-deferred accounts, to benefit from the tax advantage status, and their understanding of having to make a costly withdrawal if a catastrophic income shock occurs. The precautionary-

¹²The match is set to zero during retirement because the employer's matching ceases when the employee retires.

savings motive affects household behavior, especially at a young age. This group of households prefers to save more of the rebate, instead of consuming it.

Similar to the benchmark scenario, our one-asset and two-asset models generate a measure of young and middle-aged households that respond to tax rebates in the face of a catastrophic shock (see Figure 9). In the one-asset model with catastrophic shocks, households between the ages of 20 and 60 display almost no response, but in the model with taxable and tax-deferred assets, those households demonstrate a modest response to the rebate. The largest response, however, is exhibited by retirees, consistent with households surveyed by Shapiro and Slemrod (2009). Retired households constitute the largest fraction of households in the population that consume at least some of the tax rebate upon receipt (see Figure 10).

7 Conclusion

We developed a model of household behavior to explain the observed effects of the 2001 tax rebate on household consumption. One component of tax reform in 2001 was a transfer that accounted for about 0.6 percent of yearly personal consumption. We modeled this transfer, disbursed through an exogenously staggered process, in a life-cycle model of portfolio choice with liquid taxable and tax-deferred assets, labor earning risk, and liquidity constraints.

Our model is built on the distinction between tax-deferred and taxable assets. Tax-deferred assets allow households to enjoy higher-than-normal returns in exchange for lower financial flexibility because of the early-withdrawal penalty. Households may find it optimal to constrain themselves by holding few or no liquid taxable assets. When the tax rebates arrive, those households—some holding substantial but few liquid assets—behave more like constrained households in the standard life-cycle model. As their liquid wealth is small enough that it cannot be used to respond to adverse income shocks, those liquidity-constrained households spend more of the rebate on consumption.

In contrast to a standard one-asset life-cycle model, which does not generate any significant consumption response among middle-aged households, our model with two assets and a withdrawal cost generates a consumption response for a wider range of age groups. Households

with a larger share of taxable assets are more likely to use those assets to smooth consumption and not to spend a larger share of the tax rebate when received, behaving in accordance with the standard life-cycle model. Given the overall higher consumption response for younger and middle-aged households, relative to the one-asset model, our theoretical model generates a larger marginal propensity to consume out of tax rebates. Specifically, households contemporaneously consume part of their rebate and, on average, they spend almost 21 percent of it upon receipt. Moreover, our theoretical model qualitatively replicates empirical choices about household consumption and holdings of assets, as well as the share of total wealth invested in taxable assets over the life cycle.

This model can be used to help evaluate the effectiveness of policies designed to stimulate household consumption under different scenarios. Specifically, we consider the impact of a tax rebate during a recession, as well as in the presence of changes in interest rates and capital income tax rates. Overall, our results indicate that transfer policies might be adequate to provide economic stimulus, but targeted transfers may boost the MPC of a tax rebate more than across-the-board transfers. To determine the aggregate economic effects of such a policy, however, we would need to embed this model in a general equilibrium framework in which the effectiveness of tax rebates in stimulating economic output can be evaluated under a variety of economic scenarios.

References

- Agarwal, Sumit, Chunlin Liu, and Nicholas S. 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. Rao. 1994. “Uninsured idiosyncratic risk and aggregate saving.” *Quarterly Journal of Economics* 109 (3): 659–684.
- Amromin, Gene. 2003. “Household Portfolio Choices in Taxable and Tax-Deferred Accounts: Another Puzzle?” *Review of Finance* 7 (3): 547–582.
- . 2008, June. “Precautionary Savings Motives and Tax Efficiency of Household Portfolios: An Empirical Analysis.” In *Tax Policy and the Economy*, Volume 22 of *NBER Chapters*, 5–41. National Bureau of Economic Research.
- Bergstresser, Daniel, and James Poterba. 2004. “Asset Allocation and Asset Location: Household Evidence from the Survey of Consumer Finances.” *Journal of Public Economics* 88 (9-10): 1893–1915 (August).
- Bernheim, B. Douglas, Jonathan Skinner, and Steven Weinberg. 2001. “What accounts for the variation in retirement wealth among U.S. households?” *American Economic Review* 91 (4): 832–857 (September).
- Campbell, John Y., and N. Gregory Mankiw. 1990, June. “Consumption, Income and Interest Rates: Reinterpreting the Time Series Evidence.” In *NBER Macroeconomics Annual 1989*, Volume 4 of *NBER Chapters*, 185–246. National Bureau of Economic Research.
- Carroll, Christopher. 2001. “A theory of the consumption function, with and without liquidity constraints.” *Journal of Economic Perspectives* 15 (3): 23–45 (Summer).
- Cocco, Joao, Francisco Gomes, and Pascal Maenhout. 2005. “Consumption and portfolio choice over the life cycle.” *Review of Financial Studies* 18 (2): 491.
- Dammon, Robert, Chester Spatt, and Harold Zhang. 2004. “Optimal Asset Location and Allocation with Taxable and Tax-Deferred Investing.” *The Journal of Finance* LIX, no. 3 (June).

- Fernandez-Villaverde, Jesus, and Dirk Krueger. 2011. "Consumption and Saving over the Life Cycle: How Important are Consumer Durables?" *Macroeconomic Dynamics*, vol. forthcoming.
- Gomes, Francisco, Alexander Michaelides, and Valery Polkovnichenko. 2009. "Optimal savings with taxable and tax-deferred accounts." *Review of Economic Dynamics* 12 (4): 718–735.
- Hamilton, Douglas. 2008. "A reexamination of Johnson, Parker, and Souleles 2001 tax rebate estimate." Mimeo, Congressional Budget Office.
- Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles. 2006. "Household Expenditure and the Income Tax Rebates of 2001." *American Economic Review* 96 (5): 1589–1610.
- Li, Geng, and Paul Smith. 2008. "Borrowing from Yourself: 401(k) Loans and Household Balance Sheets." *Finance and Economics Discussion Series*, vol. 2008-42, Divisions of Research and Statistics and Monetary Affairs, Federal Reserve Board, Washington, D.C.
- Longstaff, Francis. 2009. "Portfolio claustrophobia: asset pricing in markets with illiquid assets." *American Economic Review* 99 (4): 1119–1144.
- Lusardi, Annamaria. 1996. "Permanent income, current income, and consumption: Evidence from two panel data sets." *Journal of Business & Economic Statistics* 14 (1): 81–90.
- Modigliani, Franco, and Charles Steindel. 1977. "Is a Tax Rebate an Effective Tool for Stabilization Policy?" *Brookings Papers on Economic Activity* 8 (1): 175–210.
- Parker, Jonathan A. 1999. "The Reaction of Household Consumption to Predictable Changes in Social Security Taxes." *American Economic Review* 89 (4): 959–973.
- Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland. 2011, January. "Consumer spending and the economic stimulus payments of 2008." Working paper 16684, National Bureau of Economic Research.
- Shapiro, Matthew D., and Joel 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.
- Shoven, John B., and Clemens Sialm. 2004. “Asset location in tax-deferred and conventional savings accounts.” *Journal of Public Economics* 88 (1-2): 23–38 (January).
- Souleles, Nicholas S. 1999. “The Response of Household Consumption to Income Tax Refunds.” *American Economic Review* 89 (4): 947–958.
- . 2002. “Consumer response to the Reagan tax cuts.” *Journal of Public Economics* 85:99–120.
- Stephens, Jr., Melvin. 2003. ““3rd of tha Month”: Do Social Security Recipients Smooth Consumption between Checks?” *American Economic Review* 93 (1): 406–422.
- Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron. 2004. “Cyclical Dynamics in Idiosyncratic Labor Market Risk.” *Journal of Political Economy* 112 (3): 695–717 (June).
- Tauchen, G. 1986. “Finite state Markov-chain approximations to univariate and vector autoregressions.” *Economics Letters* 20 (2): 177–181.

Table 1: Empirical Household Responses to Tax Rebates in Consumption of Nondurable Goods

Experiment	Contemporaneous Response	Cumulative Two-Quarter Response
Johnson et al. (2006)		
All households	0.373 (0.135) ¹	0.691 (0.260) ²
Only households getting rebate	0.247 (0.214) ³	n/a
Low income (< \$34,298)	0.756 (0.223) ⁴	1.380 (0.428)
Low liquid assets (< \$1,000)	0.633 (0.208) ⁴	1.256 (0.425)
Drop Outliers; Ordinary Least Squares		
All households ⁵	0.226 (0.105)	0.454 (0.213)
Only households getting rebate	0.093 (0.191)	n/a
Low income (< \$3,4298)	0.751 (0.220)	1.405 (0.420)
Low liquid assets (< \$1,000)	0.634 (0.208)	1.316 (0.417)
Median Regression		
All households ⁵	0.126 (0.048)	0.292 (0.120)
Only households getting rebate	0.002 (0.076)	0.075 (0.195)

¹ From Johnson et al., Table 2.

² From Johnson et al., Table 4.

³ From Johnson et al., Table 3.

⁴ From Johnson et al., Table 5. Standard errors for the contemporaneous responses not published in original paper; authors estimated.

⁵ From Hamilton. We extended his analysis to every other category we reported in this group.

n/a = not applicable

Table 2: Taxable and Tax-Deferred Assets

Asset	Example	Tax Advantage	Transaction Cost
Taxable Assets	Bonds, stocks, CDs mutual funds	None; interest and dividends subject to income tax	Very low to zero; must pay capital gains (if any) on sales of mutual funds, stocks, bonds
Defined-Contribution Retirement Accounts	IRAs, 401(k)s, SEP-IRAs, Keoghs, thrift-type acc'ts	Tax-deferred, contribution may be tax-deductable	Typically 10 percent penalty if distribution occurs before age 59.5, state and federal income tax applies
Defined-Benefit Retirement Accounts	Pensions	Tax-deferred growth	Limited withdrawals, some loans permitted; can be terminated and rolled into an IRA (see above)
Life Insurance		Tax-deferred growth	Some have 10 percent penalty if distribution occurs before age 59.5, state and federal income tax applies; other penalties vary by policy type
Education Savings Accounts (529 plans)	State-run education saving accounts	Tax-deferred; contribution may be tax-deductable	Typically 10 percent penalty (if not used for education); state and federal income tax applies

Table 3: Tax-Deferred Assets for Households That Hold Tax-Deferred Assets Only, by Age Group

(Dollars)										
Percentile	Household Age Group									
	20-24	25-29	30-34	35-39	40-44	45-49	50-54	55-59	60-64	65+
25th	650	1,400	1,500	2,000	2,000	2,600	2,000	8,000	2,500	2,000
50th	4,100	3,200	8,400	6,300	5,400	10,000	14,100	20,500	7,000	10,000
75th	9,500	5,000	28,000	28,000	21,000	47,000	69,000	29,000	15,000	30,000

Table 4: Tax-Deferred Assets for Households That Hold Tax-Deferred Assets Only, by Total Assets

(Dollars)					
Percentile	Total Financial Assets				
	\$ 0-10,000	\$ 10,000-50,000	\$ 50,000-100,000	\$ 100,000-300,000	\$ 300,000 +
25th	900	15,200	64,000	129,000	385,500
50th	2,100	24,520	76,500	140,000	595,000
75th	5,000	35,000	85,000	164,000	832,000

Table 5: Taxable Assets for Households That Hold Taxable Assets Only, by Age Group

(Dollars)										
Percentile	Household Age Group									
	20-24	25-29	30-34	35-39	40-44	45-49	50-54	55-59	60-64	65+
25th	200	200	200	200	240	250	350	210	300	3,000
50th	400	750	2,000	1,200	1,000	1,700	1,500	2,900	4,700	16,420
75th	1,470	4,000	11,800	5,000	5,600	8,800	10,000	15,500	29,000	68,500

Table 6: Taxable Assets for Households That Hold Taxable Assets Only, by Total Assets

(Dollars)					
Percentile	Total Financial Assets				
	\$ 0-10,000	\$ 10,000-50,000	\$ 50,000-100,000	\$ 100,000-300,000	\$ 300,000 +
25th	200	15,400	61,200	122,800	341,000
50th	780	24,500	73,000	156,000	550,000
75th	3,100	32,800	86,040	205,000	1,050,000

Table 7: Tax-Deferred Assets for Households That Hold Tax-Deferred and Taxable Assets, by Age Group

(Dollars)										
Percentile	Household Age Group									
	20-24	25-29	30-34	35-39	40-44	45-49	50-54	55-59	60-64	65+
25th	500	2,000	4,200	7,300	11,700	16,200	20,000	21,000	17,000	11,000
50th	1,800	8,000	16,000	29,000	40,000	54,000	68,500	80,500	61,000	52,000
75th	13,500	28,000	56,000	70,000	108,700	166,000	185,000	237,000	267,830	152,000

Table 8: Tax-Deferred Assets for Households That Hold Tax-Deferred and Taxable Assets, by Total Assets

(Dollars)					
Percentile	Total Financial Assets				
	\$ 0-10,000	\$ 10,000-50,000	\$ 50,000-100,000	\$ 100,000-300,000	\$ 300,000 +
25th	800	8,000	25,300	61,000	166,000
50th	2,000	15,100	50,000	103,200	321,000
75th	5,000	26,000	63,000	150,000	552,000

Table 9: Taxable Assets for Households That Hold Tax-Deferred and Taxable Assets, by Age Group

(Dollars)										
Percentile	Household Age Group									
	20-24	25-29	30-34	35-39	40-44	45-49	50-54	55-59	60-64	65+
25th	500	810	1,200	1,500	2,000	3,000	5,000	7,600	3,500	15,000
50th	1,500	4,000	6,000	6,320	11,030	16,960	17,300	28,500	33,000	60,000
75th	7,050	13,000	22,800	32,200	52,000	82,460	68,300	132,800	160,800	211,000

Table 10: Taxable Assets for Households That Hold Tax-Deferred and Taxable Assets, by Total Assets

(Dollars)					
Percentile	Total Financial Assets				
	\$ 0-10,000	\$ 10,000-50,000	\$ 50,000-100,000	\$ 100,000-300,000	\$ 300,000 +
25th	280	1,860	6,100	15,100	90,000
50th	1,000	6,300	20,000	53,500	266,000
75th	2,500	14,300	44,000	113,000	550,000

Table 11: Mean Portfolio Share, by Age Group

Asset	Household Age Group									
	20-24	25-29	30-34	35-39	40-44	45-49	50-54	55-59	60-64	65+
Saving, Money Market, Call Deposits	0.40	0.36	0.22	0.26	0.22	0.24	0.19	0.24	0.23	0.29
Certificates of Deposit	0.03	0.02	0.03	0.02	0.02	0.04	0.06	0.05	0.03	0.16
Directly Held Mutual Funds	0.06	0.05	0.06	0.06	0.06	0.06	0.06	0.08	0.08	0.08
Stocks	0.09	0.08	0.08	0.07	0.07	0.05	0.04	0.06	0.10	0.08
Bonds	0.00	0.00	0.00	0.00	0.00	0.01	0.01	0.01	0.01	0.01
Savings Bonds	0.05	0.02	0.02	0.03	0.03	0.01	0.02	0.01	0.01	0.02
Quasi-Taxable Retirement Accounts	0.19	0.32	0.40	0.41	0.44	0.43	0.43	0.35	0.30	0.17
Life Insurance	0.09	0.08	0.13	0.10	0.11	0.11	0.14	0.12	0.14	0.11
Other Managed Assets	0.01	0.01	0.01	0.01	0.02	0.02	0.03	0.05	0.06	0.05
Other Financial Assets	0.07	0.05	0.05	0.04	0.03	0.04	0.03	0.04	0.03	0.02

Table 12: Mean Portfolio Share, by Financial Assets

	Total Financial Assets				
	\$0-10,000	\$10,000-50,000	\$50,000-100,000	\$100,000-300,000	\$300,000+
Saving, Money Market, Call Deposits	0.48	0.22	0.15	0.14	0.08
Certificates of Deposit	0.04	0.09	0.09	0.06	0.04
Directly Held Mutual Funds	0.02	0.05	0.08	0.10	0.14
Stocks	0.04	0.05	0.07	0.09	0.17
Bonds	0.00	0.00	0.00	0.01	0.03
Savings Bonds	0.03	0.01	0.02	0.01	0.01
Quasi-Taxable Retirement Accounts	0.23	0.38	0.44	0.41	0.37
Life Insurance	0.11	0.15	0.10	0.11	0.06
Other Managed Assets	0.00	0.01	0.02	0.06	0.09
Other Financial Assets	0.06	0.04	0.02	0.02	0.01

Table 13: Parameterization

Parameter	Variable	Value
β	Time rate of preference	0.957
γ	Coefficient of relative risk aversion	2
\underline{b}	Minimum net asset position	0
ρ	Autoregressive parameter of the persistent income component	0.935
σ_η^2	Variance of the innovation to the persistent income component	0.061
σ_ϵ^2	Variance of the transitory income component	0.017
φ	Replacement ratio	0.6
τ^{kt}	Tax rate on taxable income	0.352
τ^{kd}	Tax rate on tax-deferred income	0.134
τ^l	Tax rate on labor income	0.30
τ^c	Tax rate on consumption	0.05
w	Wage rate	1.0

Table 14: Household Contemporaneous Consumption of Tax Rebate Under Alternative Scenarios

(Percent)	One Asset	Two Assets, No Penalty	Two Assets, Penalty
Benchmark Experiment	16.88	16.45	20.71
Changes in Penalty			
40 percent nonretirement; 25 percent retirement	n/a	n/a	19.24
30 percent before age 50; 20 percent after age 50	n/a	n/a	20.62
Alternative Tax Rates on Taxable Income			
$\tau^{kt} = 0.25$	n/a	n/a	18.46
$\tau^{kt} = 0.45$	n/a	n/a	22.11
$\tau^{kt} = 0.299$ (Bush)	n/a	n/a	19.30
Size of the Tax Rebate			
Rebate = 0.3 percent of personal consumption	17.02	16.53	20.73
Rebate = 6.0 percent of personal consumption	13.19	12.80	16.40
Changes in discount factor			
$\beta = 0.98$	5.06	4.87	7.64
$\beta = 0.94$	28.13	26.84	31.93
Replacement Ratio			
$\varphi = 0.7$	21.10	20.57	24.30
$\varphi = 0.5$	13.46	13.10	16.47
Recession			
Short (w, r fall by 5 percent for 2 years)	14.82	14.32	17.12
Long (w, r fall by 5 percent for 6 years)	14.77	11.48	15.18
Matching Contributions (30 Percent Match) Up To			
0.0025	n/a	n/a	19.98
0.005	n/a	n/a	18.94
0.01	n/a	n/a	17.42
0.03	n/a	n/a	16.41
Catastrophic Shocks			
1 percent chance of 90 percent drop in income	8.79	8.29	11.36

n/a = not applicable

[Table revised on July 22, 2011, from the original version, published on July 14, 2011]

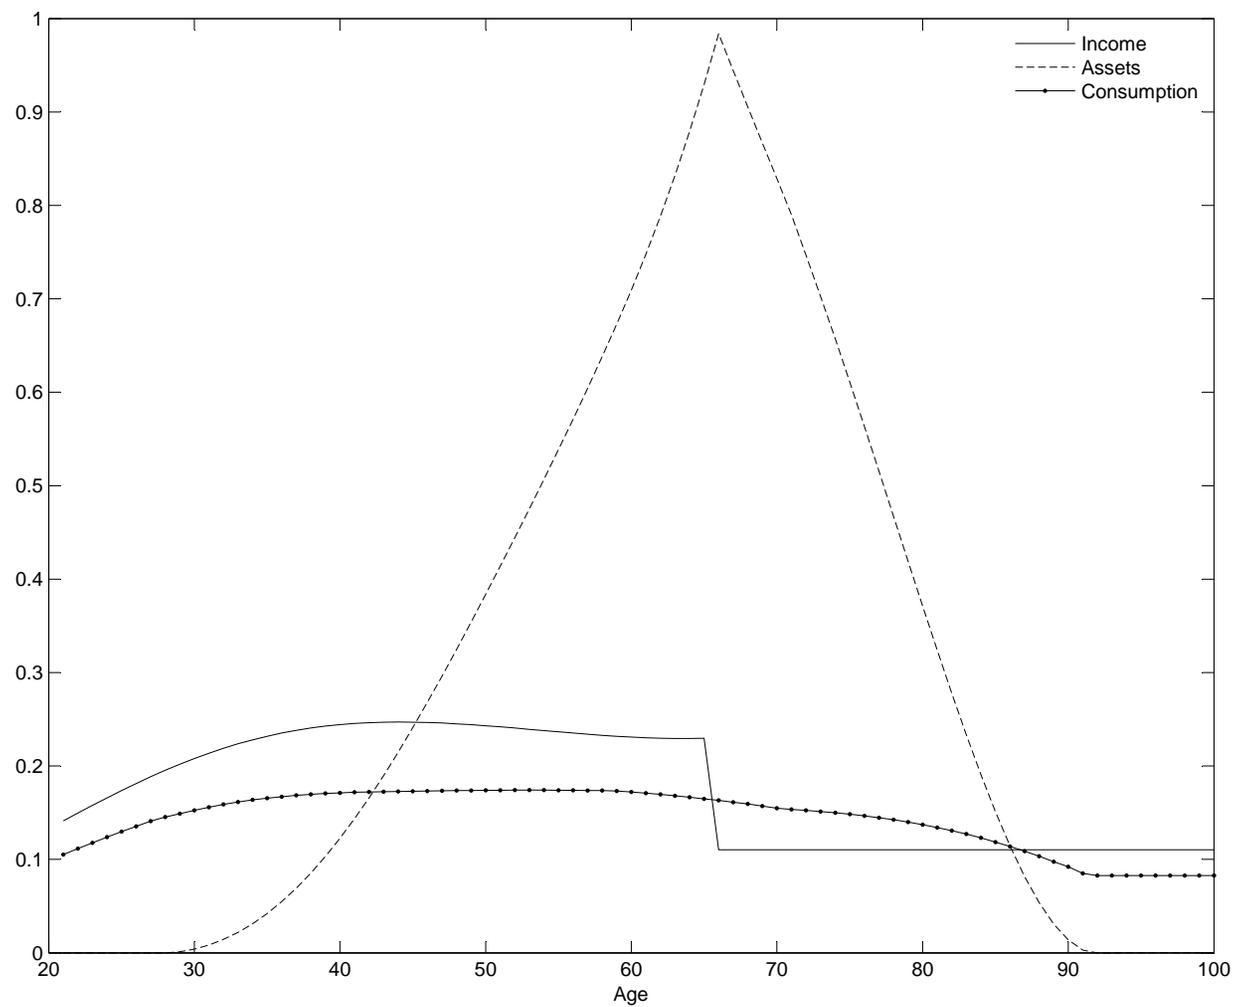
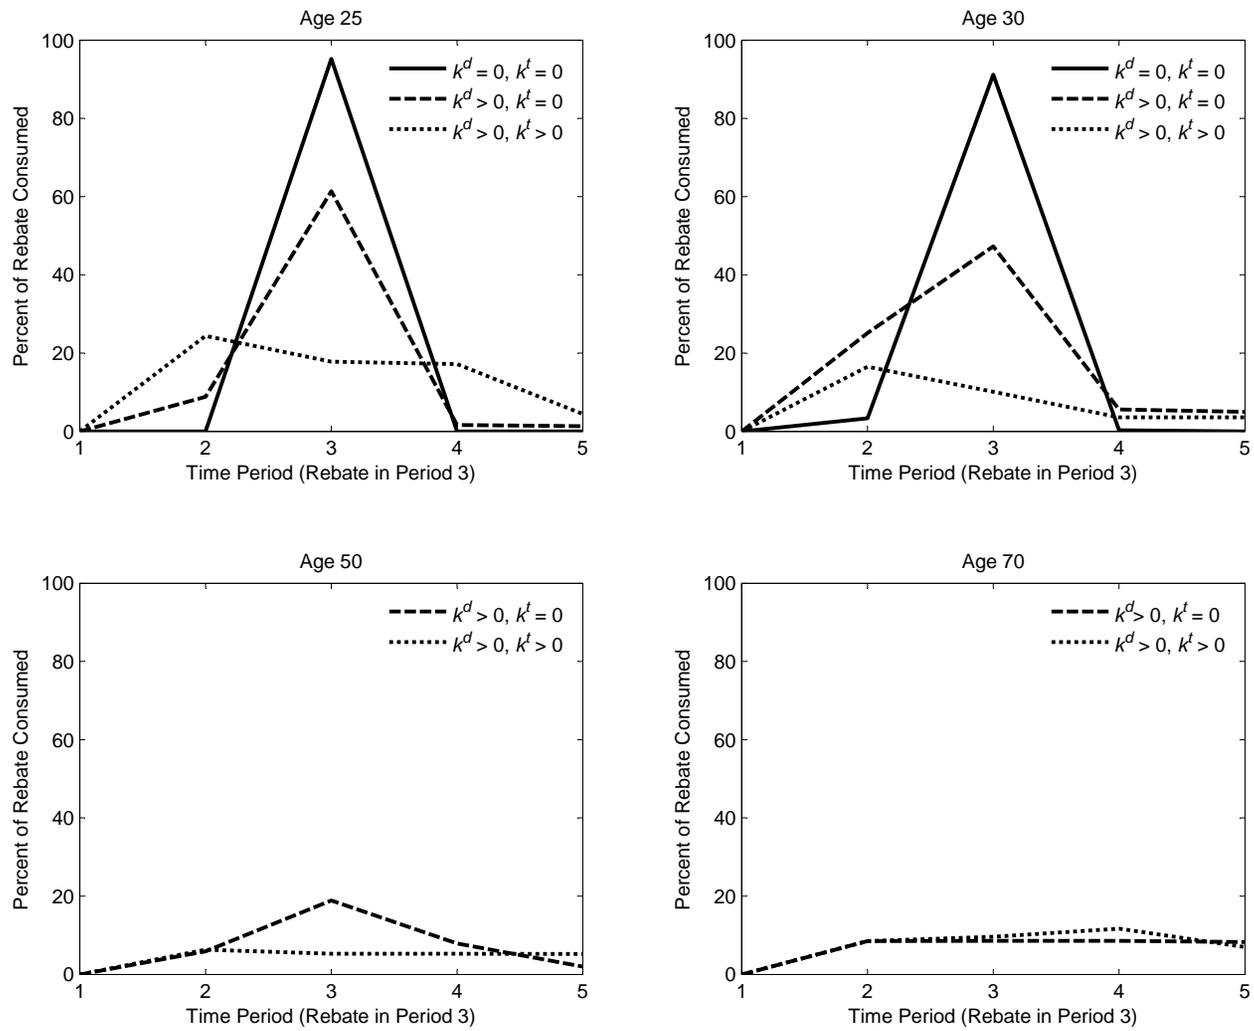


Figure 1: Income, Assets, and Consumption Path Over the Life Cycle



k^d = tax-deferred assets, k^t = taxable assets.

Figure 2: Household Consumption Response to Tax Rebate

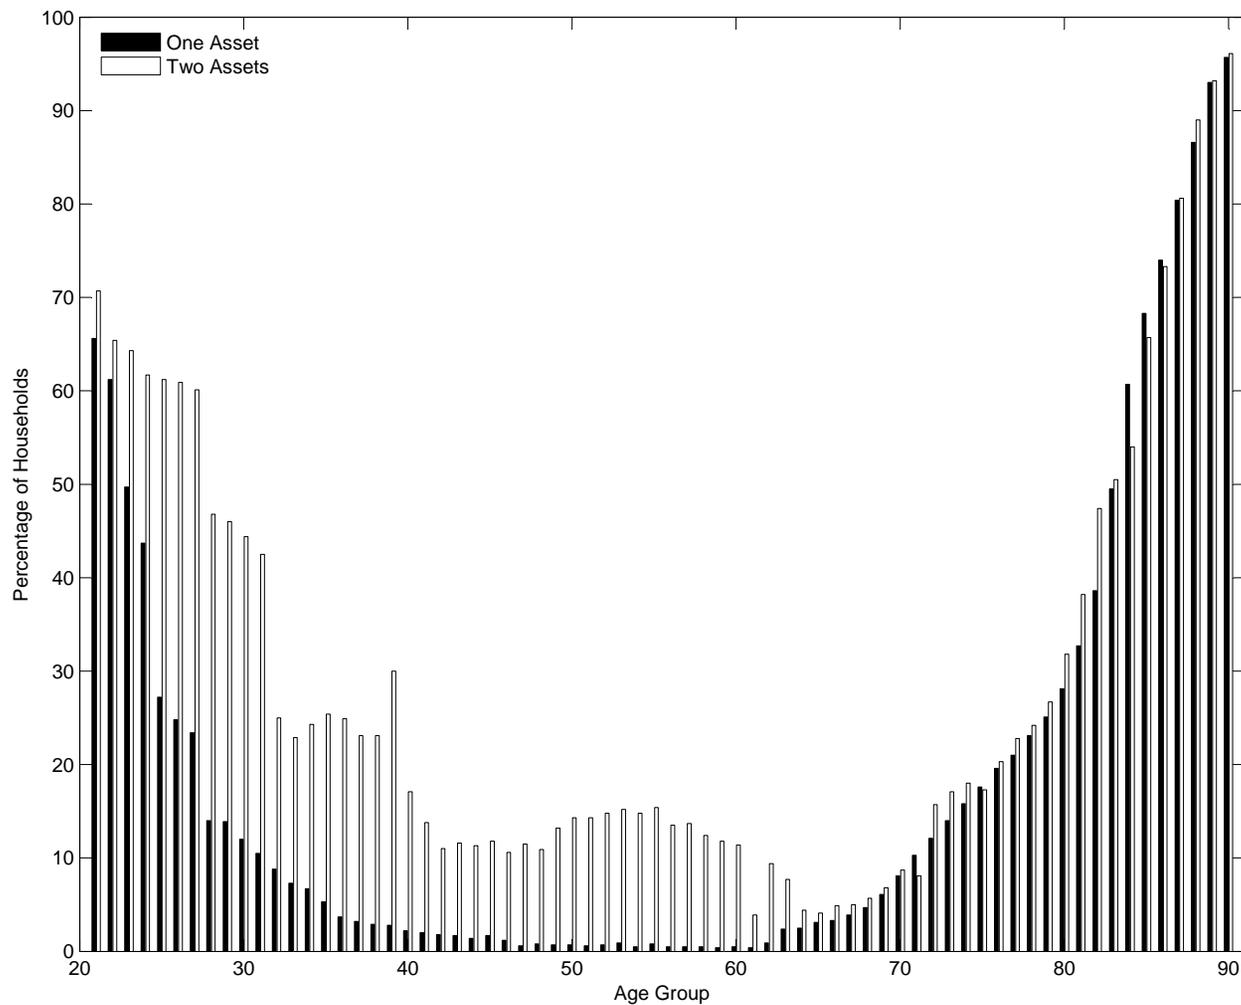


Figure 3: Percentage of Households Responding Contemporaneously to Tax Rebate, by Age Group. The figure reports the household probability of responding to a tax rebate upon receipt, conditional on age. The results are displayed for the model with one asset and for the model with two assets, tax-deferred and taxable, with withdrawal penalty.

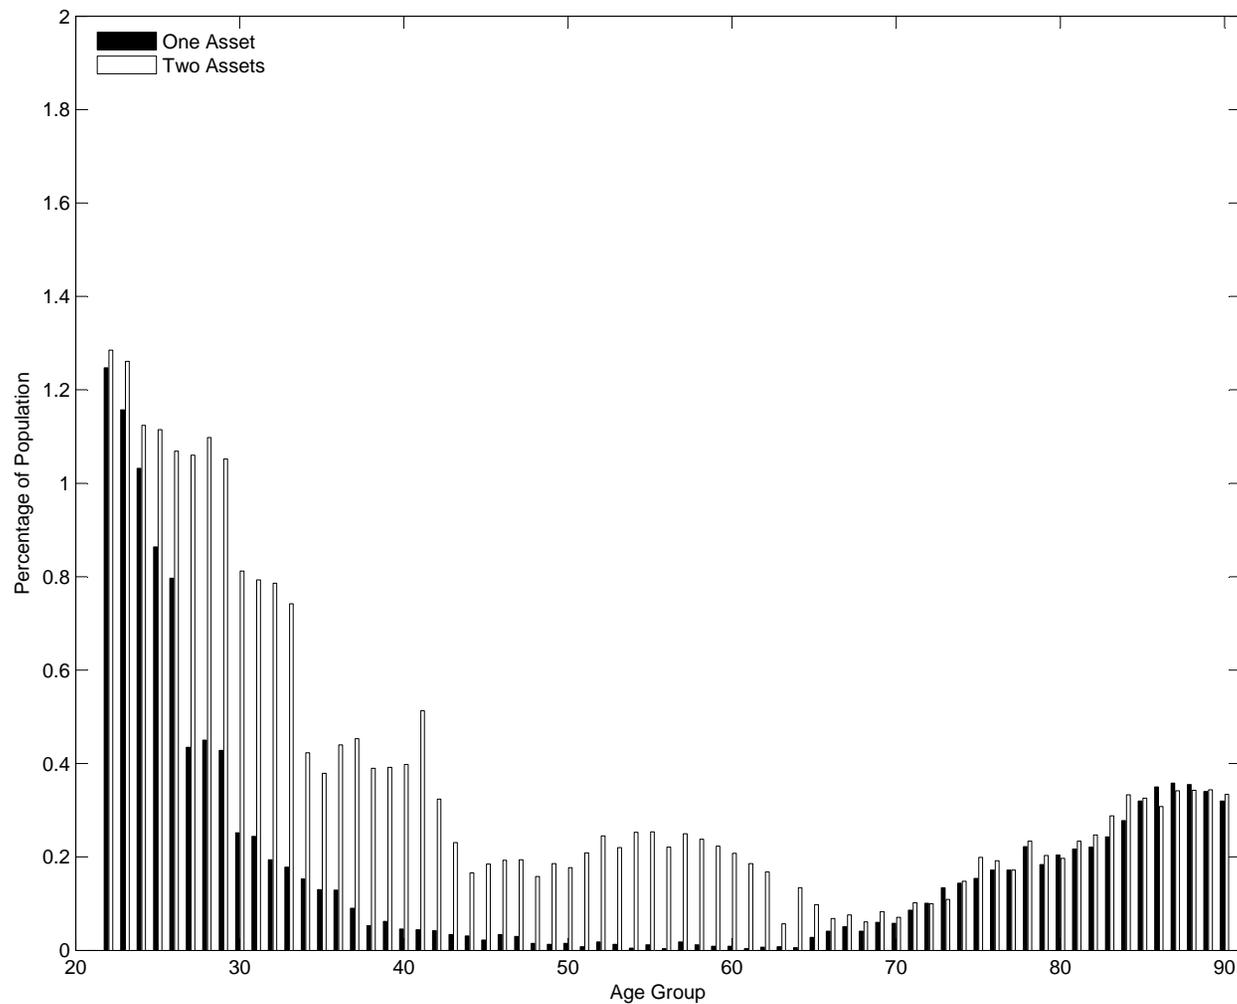


Figure 4: Percentage of Total Population Responding Contemporaneously to the Tax Rebate, by Age Group. The figure reports the household joint probability of belonging to a given age group and responding to a tax rebate upon receipt. The results are displayed for the model with one asset and for the model with two assets, tax-deferred and taxable, with withdrawal penalty.

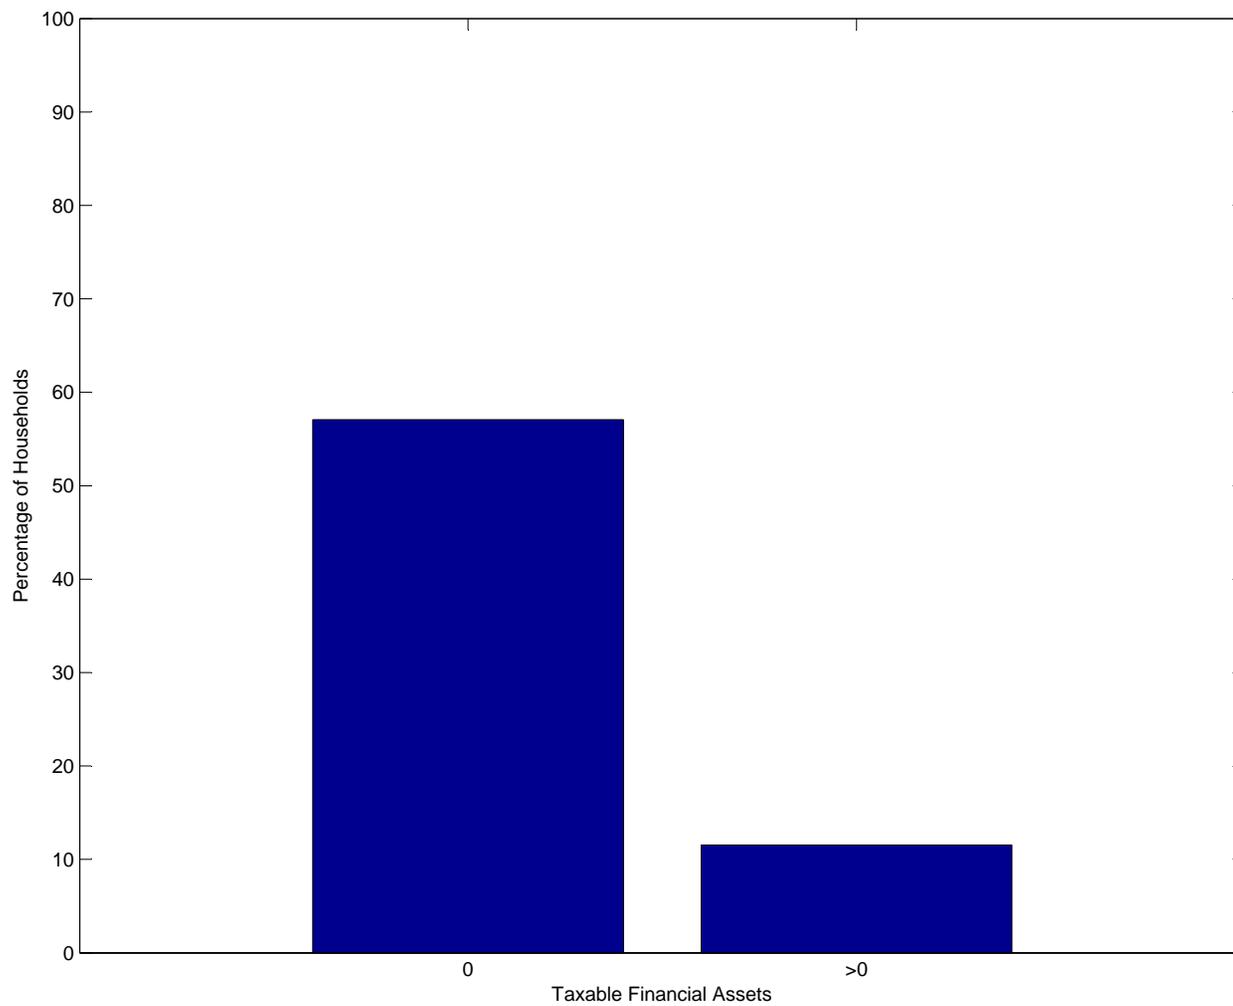


Figure 5: Percentage of Households Responding Contemporaneously to Tax Rebate, by Holdings of Taxable Assets.

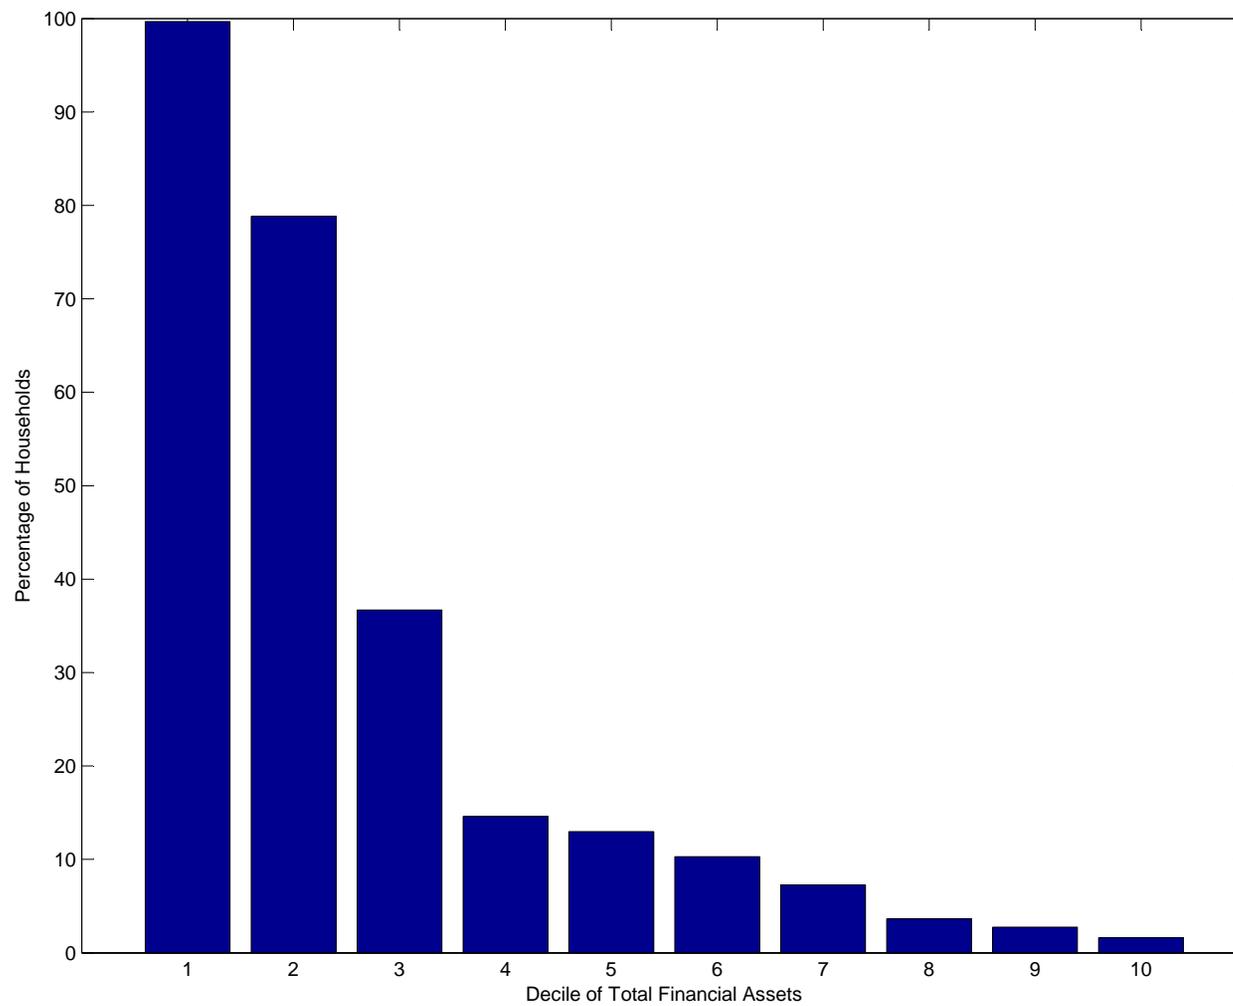


Figure 6: Percentage of Households Responding Contemporaneously to Tax Rebate, by Holdings of Total Assets.

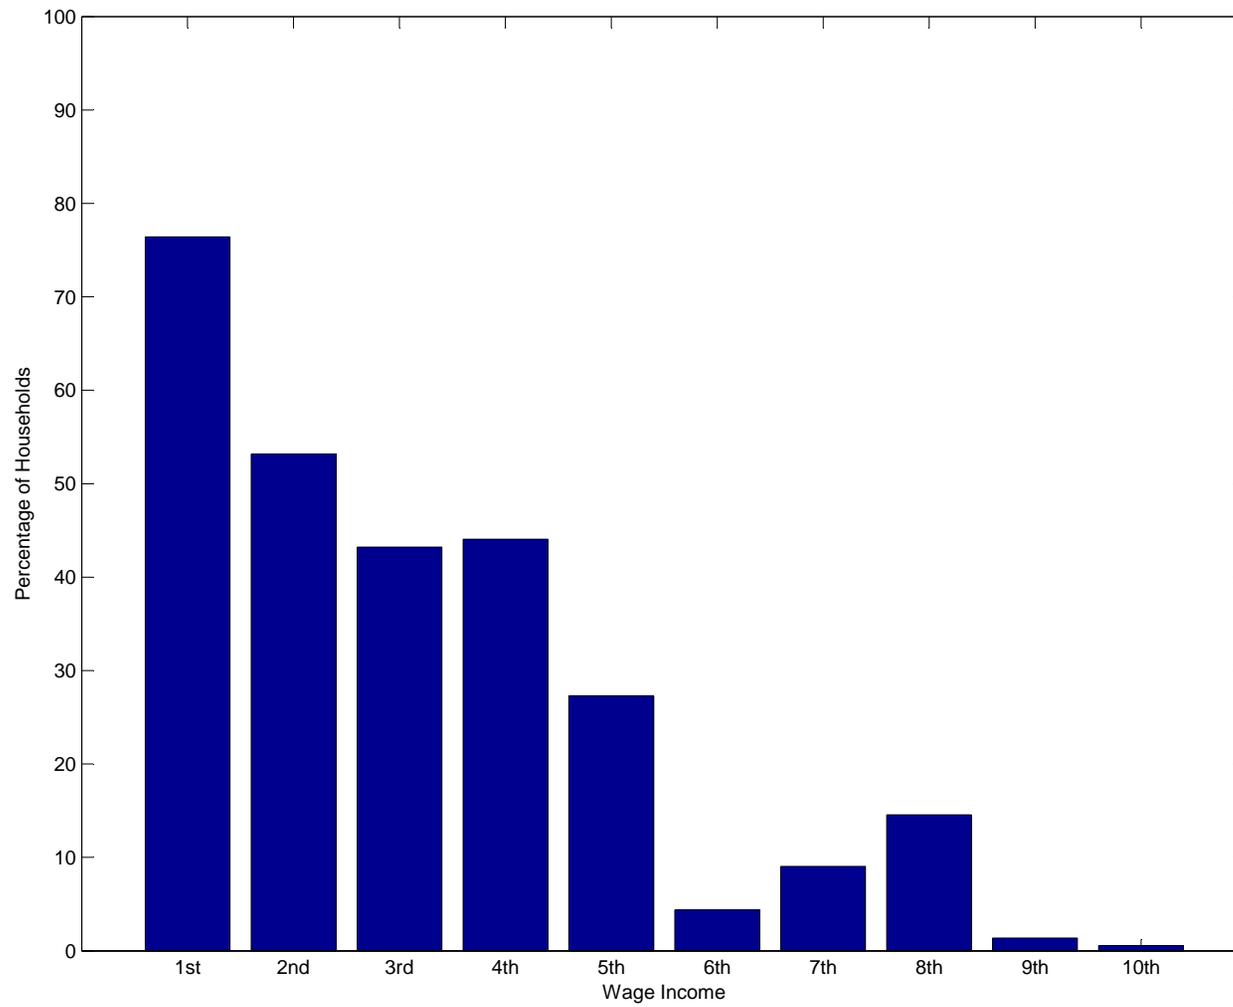


Figure 7: Percentage of Households Responding Contemporaneously to Tax Rebate, by Income Group.

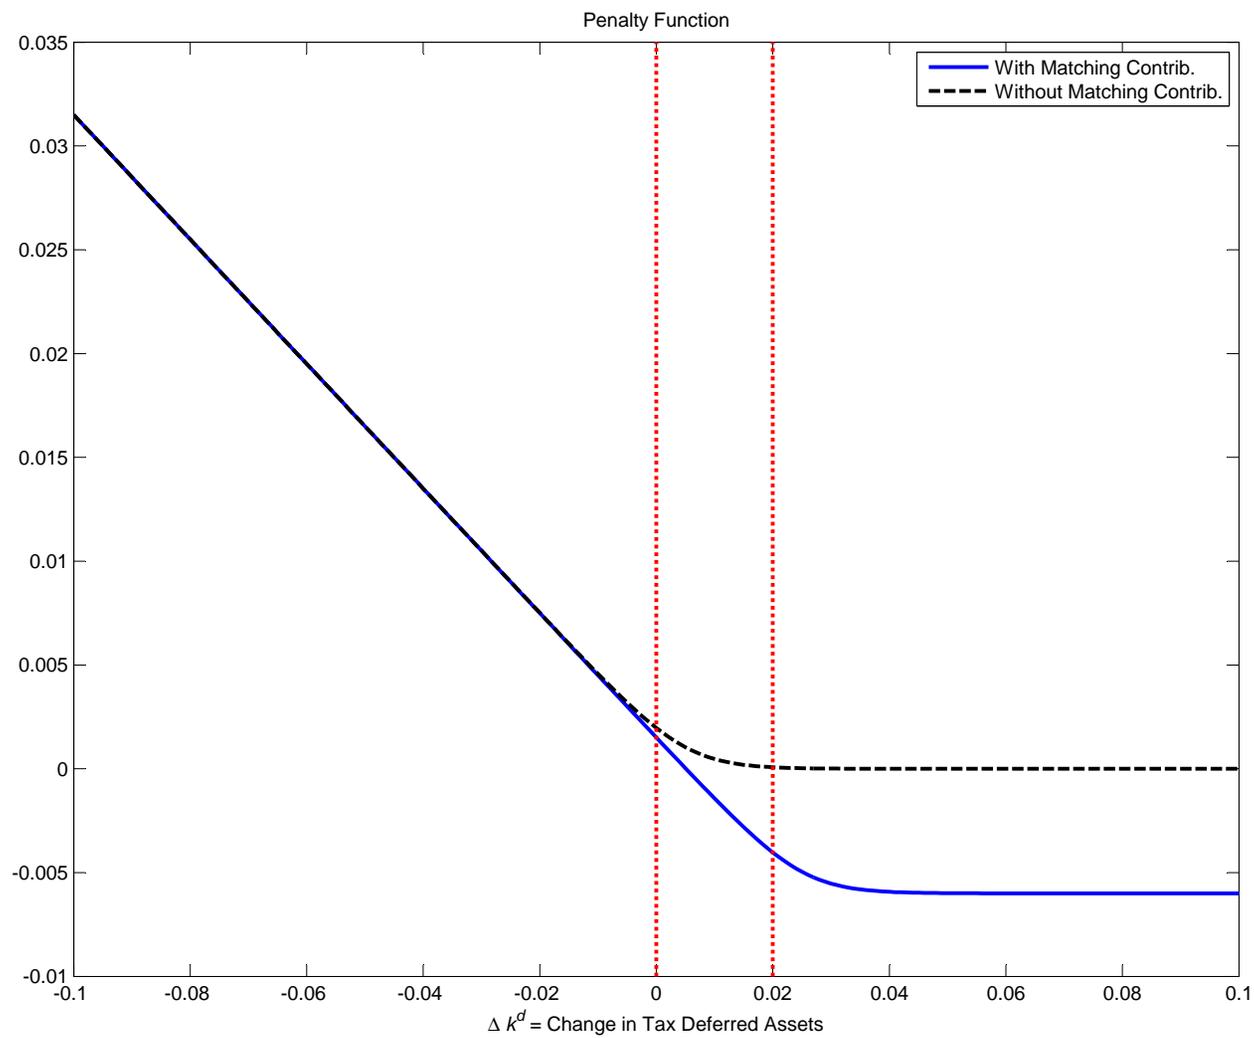


Figure 8: Cost of Changing Stock of Tax-Deferred Assets.

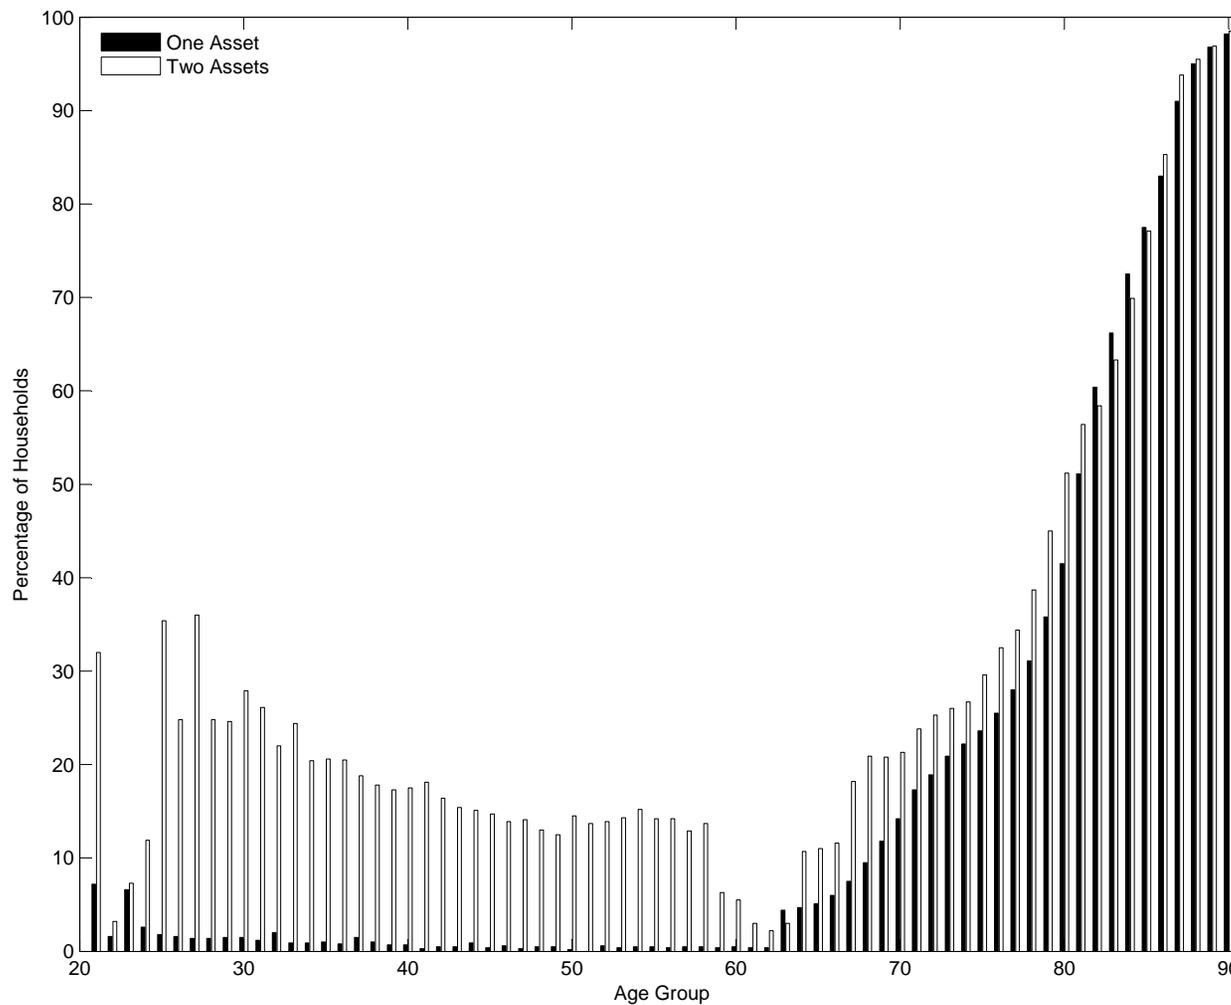


Figure 9: Percentage of Percentage of Households Responding Contemporaneously to Tax Rebate; in a Model with Catastrophic Shocks, by Age Group. The figure reports the household probability of responding to a tax rebate upon receipt, conditional on age. The results are displayed for the model with one asset and for the model with two assets, tax-deferred and taxable, with withdrawal penalty.

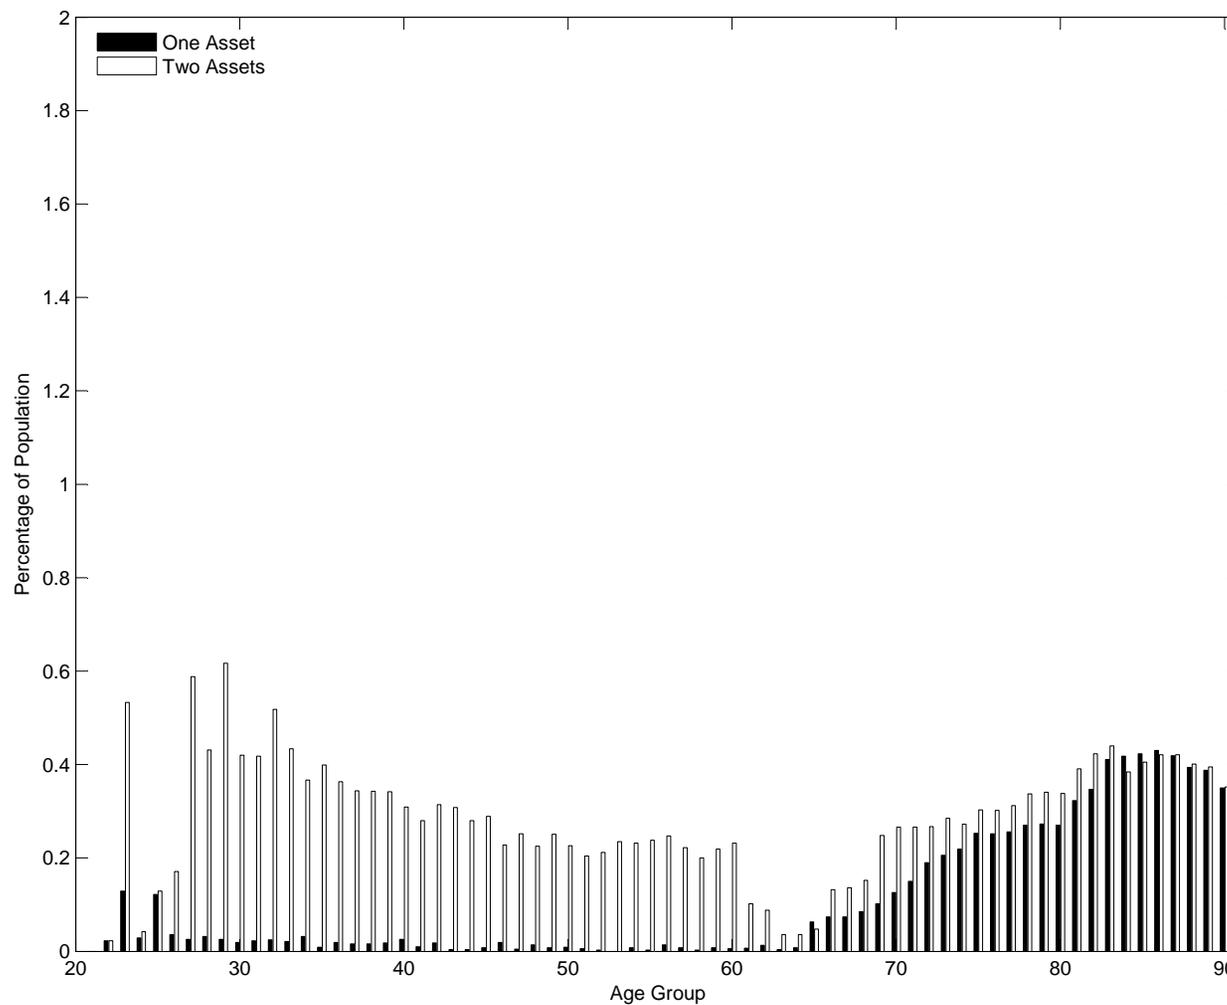


Figure 10: Percentage of Total Population Responding Contemporaneously to Tax Rebate; in a Model with Catastrophic Shocks, by Age Group. The figure reports the household joint probability of belonging to a given age group and responding to a tax rebate upon receipt. The results are displayed for the model with one asset and for the model with two assets, tax-deferred and taxable, with withdrawal penalty.