Stimulus through Insurance: The Marginal Propensity to Repay Debt*

Preliminary and Incomplete

Gizem Kosar — NY Fed

Davide Melcangi — NY Fed

Laura Pilossoph — NY Fed

David Wiczer — Stony Brook

May 31, 2022

Abstract

We empirically document that households often use "stimulus" checks to pay down debt, especially those with low income and high debt balances. We build a heterogeneous-agent model that formalizes the concept of the marginal propensity to repay debt (MPRD), highlights the interaction between debt delinquency, endogenous borrowing constraints, and consumption, and is consistent with our key empirical facts. The MPRD substantially alters the aggregate implications of fiscal transfers. We show that there exists an apparent trade-off between stimulus and insurance, as well as between short- and long-run fiscal multipliers.

*The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

1 Introduction

Many households use their stimulus checks to pay down their existing debt. For example, after receiving their 2008 rebates as part of the Economic Stimulus Act, 52% of households reported that they used the money to mostly pay down debt, while 20% reported that they mostly spent it (Sahm, Shapiro and Slemrod (2010)). However, despite the disproportionate use of these checks for debt repayment, both academic and public discussions have instead focused on their ability to boost consumption through increased spending. Indeed, this is why the checks have come to be called ”stimulus checks,” and why their success is often measured by whether or not they are spent.

In this paper, we argue that the degree to which households use their checks for debt repayment can be as important as the degree to which they spend them. While immediate spending boosts consumption closer to the timing of receipt of the check, those who pay down their debt avoid delinquency, maintain better terms of credit, and are thus able to consume more in the future. Therefore, in evaluating such programs, both the marginal propensity to consume (MPC) and the marginal propensity to repay debt (MPRD) are important yardsticks for measuring their success.

We begin by documenting how households used their stimulus checks during the COVID-19 pandemic using data that was collected as part of the New York Fed's Survey of Consumer Expectations (SCE). While there is an extensive empirical literature estimating marginal propensities to consume out of transitory income shocks, empirical evidence on debt responses is much more limited, most likely due to data limitations.[1] We circumvent this issue by eliciting responses directly from surveyed households in the SCE.[2] We document three main facts. First, on average, households used a third of their checks to pay down debt. Such a share is higher than the average marginal propensity to consume (MPC). Second, households with high debt-to-income ratios are more likely to pay down debt. High-MPRD households typically have both low income and a larger stock of debt. Third, there is a negative cross-sectional relationship between MPCs and MPRDs. In particular, households with larger debt-to-income ratios have lower MPCs.

We then build a heterogeneous-agent, consumption-savings life-cycle model that is consistent with these facts. We extend the framework of Kaplan and Violante (2010) to allow for the fact that households can be delinquent on any share of their outstanding debt. Modeling default as a continuous choice has been recently done in the context of sovereign debt (see Arellano, Mateos-Planas and Ríos-Rull (2019)), as well as when investigating the relationship between unemployment and credit access (see Herkenhoff (2019)). We show that introducing the possibility of delinquency and, in turn, partial repayment of debt, substantially alters the effects of fiscal stimulus.

In line with the literature, our model features two main delinquency costs. First, there is an exogenous utility cost. Second, the price of new debt endogenously depends on households' saving decisions. This channel effectively operates as an endogenous borrowing constraint.

[1] One notable exception is Agarwal, Liu and Souleles (2007), who estimate debt responses to 2001 tax rebates in the United States. They find that consumers initially used the checks to pay down debt, and this stimulated spending in the medium-run, consistent with the channel we uncover in our model.

[2] Coibion, Gorodnichenko and Weber (2020) adopt a similar approach, using another survey of U.S. households. We discuss in Section 2 how our results relate to theirs.

Our model features a novel static trade-off between debt repayment and consumption, as well as a dynamic trade-off between the same objects, because indebted households can save either by repaying part of their delinquent stock of debt, or by borrowing less. Paying down debt lowers the utility cost of delinquency, and delivers a better price for new debt for a given level of new debt, allowing households to borrow more today. However, all else equal, it implies that households have fewer resources to consume.

This trade-off shows up when investigating households' responses to lump-sum transfers (i.e., "stimulus" checks). Indebted households would like to use part of the rebate to pay down debt, but this may come at the cost of foregoing resources away from consumption. However, a better price schedule can increase households' resources through new borrowing. As a result, the model delivers rich heterogeneity in the MPC and MPRD, which depends on a household's net asset position.

Our model is able to replicate the three empirical facts we document in the SCE. In addition, we show that alternative versions of the model typically fail to do so. First, the average MPRD is high and, in line with the data, is even slightly higher than the average MPC. While mechanically being 0 in a standard model without partial default, MPRDs are also much lower if the price of new debt is assumed to be constant. Hence, our model underscores the importance of endogenous borrowing constraints and debt price responses to fiscal transfers. This mechanism is broadly consistent with the empirical observation that credit scores improved during the pandemic, at least partly thanks to households using stimulus checks to pay down debt. The model also uncovers a large degree of MPRD heterogeneity, with half of the debtors not using the check at all to pay down debt, and the remaining half having very different marginal propensities to repay.

Second, our baseline model predicts a positive correlation between MPRD and debt-to-income ratios. Consistent with the data, households with larger debt balances are more likely to pay down debt, and this effect is particularly strong at low debt levels. While MPRDs are increasing with debt-to-income ratios also when debt prices are constant, this relationship is much stronger and more concave in the full model, bringing it closer to the data.

Third, MPCs fall with debt-to-income ratios, a finding that is also consistent with our empirical evidence. In the standard model with no delinquency and a natural borrowing limit, MPCs fall with net worth, a well-known finding in the literature. Therefore, when zeroing in on indebted households, this model is not able to replicate the empirically observed relationship. Such correlation, instead, is strongly negative in our full model with delinquency. MPCs peak among households with little assets or relatively small debt-to-income ratios. The latter are households that have a little incentive to be delinquent, and especially to use the check to pay down debt. In contrast, they may have a relatively high marginal utility of consumption, due to the effect exerted by endogenous borrowing constraints arising from the debt price schedule. As such, they have a relatively high MPC. As debt-to-income ratios increase, households find it preferable to pay down debt. Hence, as shown before, the MPRD increases, and the MPC falls, since paying down debt diverts resources away from consumption in the near term.

The paper is organized as follows. In Section 2 we present three empirical facts on the MPRD and the MPC using consumption data from the U.S.. Section 3 describes our heterogenous-agent life-cycle consumption-saving model with delinquency and endogenous debt prices. In Section 4 we show that a calibrated version of the model is consistent with the empirical facts. Section 5 discusses how accounting for the MPRD substantially alters the aggregate implications of fiscal transfers, and Section 6 concludes.

2 Empirical facts on households' responses to stimulus checks

In this section we document three main facts on how households respond to lump-sum transfers. To do so, we use data from a special module fielded in June 2021 as part of the New York Fed's Survey of Consumer Expectations (SCE), with a focus on the transfers to individuals that were part of the CARES Act.[3] The SCE is a monthly internet-based survey of a rotating panel of approximately 1,300 heads of household from across the US. As the name of the survey indicates, its goal is to elicit expectations about a variety of economic variables, such as inflation and labor market conditions. Respondents participate in the panel for up to twelve months, with a roughly equal number rotating in and out of the panel each month. Respondents are also asked to participate in additional modules every month and receive extra incentives when they do. Our data come from the June 2021 special survey that was separately administered from the core module and that invited respondents who rotated out of the survey as well, allowing us to reach a sample size of 1,608 observations.

Our analysis primarily focuses on questions about the receipt and usage of stimulus checks. The survey first asks whether the respondent's household has received a stimulus payment (either by direct deposit or via check) and, if so, how much in total they received. We find that around 89% of the respondents in our sample received the stimulus payments, and the average (median) payment received is $2,080 ($2,400). The respondents were then asked the question on the allocation of the stimulus payment in the following form:

Please indicate what share of this government payment you have already used to or expect to use to...
Save or invest __________%
Spend or donate[4] __________%
Pay down debts __________%

where the responses add up to 100.[5] In the rest of this section, we will refer to the "Save or invest" allocation as the marginal propensity to save (MPS), "Spend" allocation as the MPC, and the "Pay down debt" allocation as the MPRD. We will also focus only on those who reported receiving a payment at the time of the survey and consider their allocation of the payment.

[3]The CARES Act was a large stimulus packaged passed by the U.S. Federal government on March 27^{th} 2020. As part of this package, all qualifying adults received a one-time transfer of up to $1200, with $500 per additional child. For further details on the episode, we defer to Coibion et al. (2020).

[4]The follow-up questions ask the respondents to split the "Spend or donate" allocation into separate "Spend" and "Donate" allocations.

[5]The respondents see the running total of their answers and receive an error message if they try to move on to the next question before the total is equal to 100.

Figure 1: Histograms of MPRD and MPC

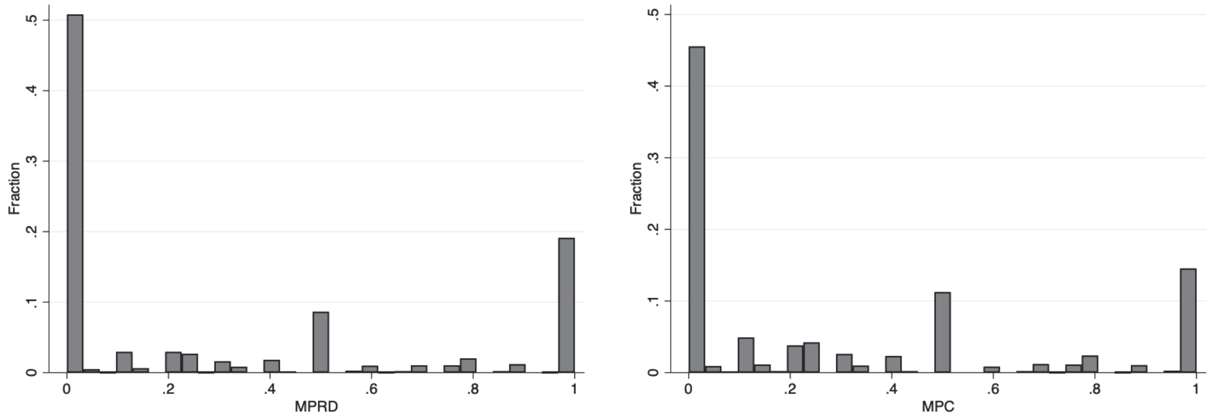

(a) MPRD

(b) MPC

Fact 1: The average MPRD across households (0.32) is as large as the average MPC across households (0.30).

Based on the survey responses, we find that, on average, households used one-third of their checks to pay down their debt, with a median MPRD of 0. The histogram for the MPRD presented in Figure 1a shows that around 50% of the respondents do not use their checks to pay down any debt, while around 19% of the respondents report using all of their stimulus checks to repay their debt, leading to a bimodal distribution. Focusing on the spending allocations, we observe a similarly-shaped distribution for the MPC in Figure 1b, but this time with a lower share of respondents at the two extremes, leading to a median MPC of 0.1. Restricting the attention to households with positive debt, the average MPRD not surprisingly increases, to 39 cents per rebate dollar, while the average MPC falls to 28 cents, in line with Fact 3 shown later.

Our findings are consistent with a similar empirical investigation by Coibion et al. (2020), who ask consumers in the Nielsen Homescan panel how they used the same fiscal stimulus payments we study. They also find that, on average, 30 percent of the rebate was used to pay down debt. However, the average MPC they find based on the responses to their survey is around 0.42, which is higher compared to our findings. This higher MPC may reflect differences in sample composition and question wording, as well as a difference in the timing of the surveys used. The data we use is from the June SCE, a time relatively soon after households received their payments, while Coibion et al. (2020) use survey responses from July. For this reason, the higher average MPC may suggest that some of these payments that were temporarily held as savings – or debt repayments – in June may have been used subsequently for consumption.

Earlier studies on the 2001 and 2008 tax rebates, instead, do not generally estimate the share of transfers used by households to spend, save, or pay down debt, but rather estimate what fraction of households mostly used the rebate for either of the three options. A key takeaway from these studies is that a large fraction of households use the rebate to pay down debt.[6] Agarwal et al. (2007) use panel data of credit card accounts to estimate consumption and debt responses to 2001 tax rebates in the U.S.. They find that households receiving the rebate, on average, reduced their credit card debt by about $27 more than those that did not receive the rebate, in the first quarter upon receipt. While they do not directly estimate an MPRD, their estimated debt response is 8% of the average payment in their sample. Besides limitations in this back-of-the-envelope calculation, the estimated sensitivity is lower than what we estimate in the SCE likely because credit card debt comprises a limited fraction of household debt. We come back to this quantitative discussion when analyzing the model results.[7]

Fact 2: The MPRD is higher for those with high debt balances.

While many household characteristics are individually correlated with the MPRD in a statistically significant way, only few maintain a significant relationship when tested jointly with other observables. Female respondents, those between the age of 40 and 65, those with higher debt stocks – defined as gross unsecured debt –, and those with lower annual household incomes report higher MPRDs. We find that a 1% increase in the debt stock is associated with an 8% increase in the MPRD (from an average of 0.32 to 0.346) and a 1% increase in the annual household income is associated with a 25% decline in the MPRD (from an average of 0.32 to 0.24), keeping everything else constant. MPRDs increase with debt-to-income ratios, as shown as a binscatter in Figure 2. Even when controlling for households' demographics, a 10 percentage-point increase in the debt-to-income ratios increases the MPRD by 1.2 cents, and the correlation is statistically significant at the 1% level. The figure displays an increasing, concave relationship between the MPRD and the debt-to-income ratio, showing that the MPRD is lower for those with debt-to-income ratios close to 0; it rapidly increases with the debt-to-income ratio; and it remains more or less stable at or above 0.4 for levels of debt-to-income ratio above 0.5. This shape, and the statistically significant positive relationship, remains when restricting attention to households with positive debt-to-income ratios.

While a few existing papers have explored what drives MPRD heterogeneity, most of them do not explicitly look at debt balances. For example, Sahm et al. (2010) and Coibion et al. (2020) find that low-income households are more likely to be those who reported having used the rebate to mostly pay down debt, but neither study estimates the relationship with debt. Fagereng, Holm and Natvik (2021) find that Norwegian households in the top quartile of the distribution of liquid assets use a smaller fraction of the lottery prize to repay debt than illiquid households, but the study does not look at heterogeneity by debt balances.

Fact 3: The MPC is lower for those with high debt balances.

Contrary to what was previously shown for the MPRD, we find a positive but statistically insignificant relationship between the MPC and annual household income when we control for other observable household characteristics. This is in line with Sahm et al. (2010) and Shapiro and Slemrod (2009), who find a weakly positive and insignificant relationship between income

[6]See Sahm et al. (2010), Shapiro and Slemrod (2009) and Hisnanick and Kern (2018).

[7]Outside of the U.S., Fagereng et al. (2021) estimate that, within one year of winning the lottery, Norwegian households used 7% of the prize for debt repayments.

Figure 2: MPRD and debt-to-income ratio

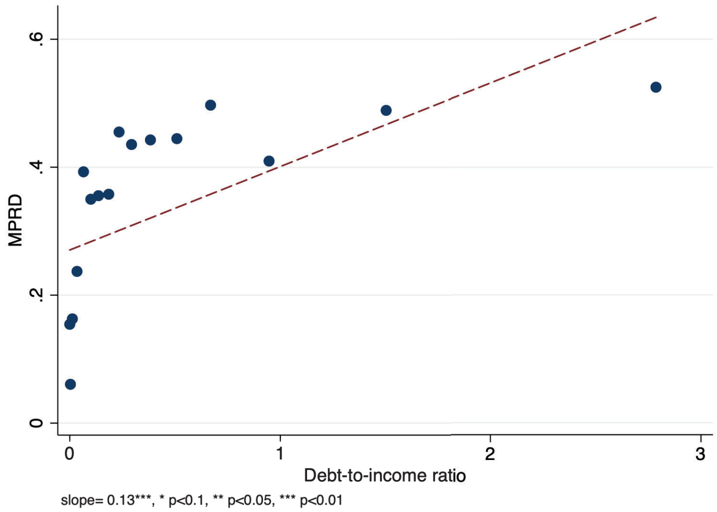

and the likelihood of mostly spend the rebate. For the same 2008 tax rebate episode in the U.S., Lewis, Melcangi and Pilossoph (2021) and Kueng (2018) estimate a statistically significant positive correlation between income and estimated spending propensities.

We document a negative relationship between the MPC and debt, as well as with debt-to-income ratios. Both correlations are statistically significant, even when controlling for other household characteristics. However, the MPC changes less with the debt stock than does the MPRD. Keeping everything else constant, we find that a 1% increase in the debt stock is associated with a 3.2% decline in the MPC. Figure 3 plots the negative relationship between MPCs and debt-to-income ratios. Conversely to what was shown for the MPRD, MPCs fall quite sharply when debt is small relative to income, and then this negative association gets weaker as debt balances grow. The negative and convex relationship, as well as its statistical significance, are maintained when restricting the sample to households with positive debt.

While not as sparse as for the MPRD, the empirical evidence on the relationship between MPCs and debt is also much less common than for other household characteristics. Various papers have focused on mortgage debt (see, for instance, Misra and Surico (2014) and Hedlund, Karahan, Mitman and Ozkan (2017)), while we are particularly interested in unsecured debt. Fagereng et al. (2021) find no significant relationship between MPCs and debt levels when controlling for other household characteristics.

Finally, all the three facts we have showed so far are not specific to the COVID-19 episode. Indeed, we analyze additional data from the spending module of the SCE that has been collected every 4 months since August 2015. As part of this module, households were asked to think about an hypothetical 10% increase in income, and to report which fraction of it they would use to spend, save, or pay down debt. Even in this setting, using data from before 2020, MPRDs were high, and were increasing and concave in debt-to-income ratios, while MPCs were decreasing and convex. In the next section we introduce a model that is consistent with these facts.

Figure 3: MPC and debt-to-income ratio

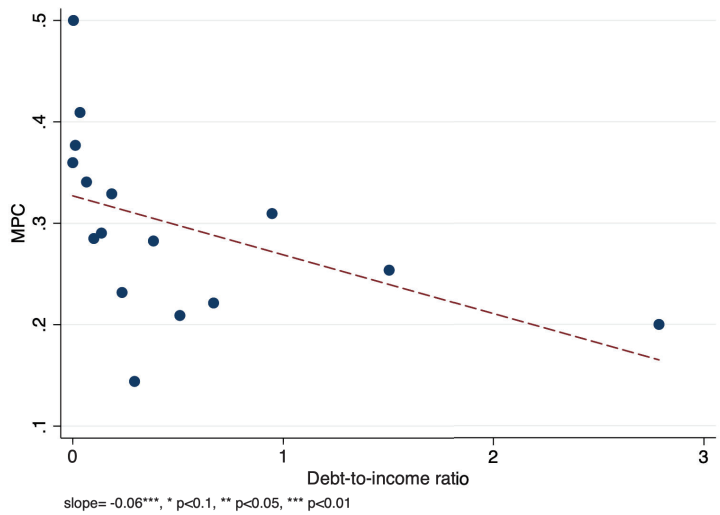

3 Model

We consider a standard life-cycle incomplete-markets model of consumption and savings, building off Kaplan and Violante (2010). We extend this framework to incorporate delinquency in the form of partial default. The economy is populated by a continuum of households indexed by their age t, their net asset holdings a, and their income y.[8] Households can borrow at an endogenous price schedule q, whose determination we describe below.

Income process Our characterization of the evolution of households' income follows Kaplan and Violante (2010) and Blundell, Pistaferri and Preston (2008). Households work until retirement which occurs at age $t = T_{ret}$. While working, (log) labor income y follows:

$$y = \kappa + z + \epsilon \tag{1}$$

where κ is a deterministic experience profile, z is the permanent component of income that follows a random walk, and ϵ is the transitory component of income. Innovations to z, denoted with η, as well as transitory shocks ϵ, have mean zero and are normally distributed with variances σ_η and σ_ϵ, respectively. ν and ϵ are orthogonal to each other and independently distributed over time and across households. Retired households receive social security income, which is a function of the income they earned while working.

Household choices Households start the period with their net asset position a and income y. $a < 0$ denotes an outstanding stock of debt. Every period, households choose the amount of newly purchased net assets, b, and the fraction of outstanding net debt that the household repays, s. As often done in the literature, a fraction κ of delinquent debt $(1 - s)\,a$ becomes future debt obligations. The law of motion of households' net worth is described by:

[8] We omit time subscripts and denote next period variables with $'$.

$$a' = b + (1 - s)\, a\kappa \tag{2}$$

Households always choose $s = 1$ when $a > 0$, since otherwise they would be giving up some of their positive assets for free. However, they can decide to be delinquent on part of their outstanding debt (e.g.: it is possible that $s < 1$ when $a < 0$). b is purchased at a price q, which we describe later. In a nutshell, q can be lower than its risk-free level when the choice of b takes the household into debt ($a' < 0$), due to endogenous default premia.

Households also consume a nondurable good c from which they derive utility every period. The household problem can be summarized as follows:

$$
\begin{aligned}
V(t, a, z, \epsilon) = &\max_{c>0, s\in[0,1], b} u(c) - \phi(s) + \beta E V(t+1, a', z', \epsilon') \\
&\text{subject to :} \\
&\qquad c + q(t, a', z)b - as = e^{z+\epsilon+\kappa_t} + \tau \\
&\qquad z' = z + \eta \\
&\qquad a' = b + (1 - s)a\kappa
\end{aligned}
$$

where $\phi(s)$ denotes a utility cost from delinquency, which is increasing in partial default (i.e., $\frac{\partial \phi(s)}{\partial s} < 0$), and such that $\phi(1) = 0$ and $\phi(0) > 0$. Herkenhoff (2019) also assumes there exists a utility penalty from partial default, in a model that studies the relationship between consumer credit and unemployment. Similar costs are also used in the sovereign debt literature, for instance by Arellano et al. (2019). τ denotes a lump-sum transfer, akin a rebate check, which we revisit later.

(Partial) default premia The price schedule of net asset purchases, q, is determined such that lenders make no profit. As a result, the amount lent by the lender in period j, qb, needs to be equal to the discounted return of what the lender will get in period $j + 1$. The latter is the sum of two components. First, the expected repayment in period $j + 1$:

$$\frac{1}{1+r} \int_{\epsilon'} \sum_{z'} \left[s' b \pi(z, z') \right] dF(\epsilon')$$

Second, the expected repayment of the non-repaid portion, in the following period:

$$\frac{1}{1+r} \int_{\epsilon'} \sum_{z'} [(1 - s')\kappa b q' \pi(z, z')] dF(\epsilon')$$

Combining, we get:

$$
\begin{aligned}
&q(t, z, a')b(t, a, z, \epsilon) = \\
&\quad \frac{1}{1+r} \int_{\epsilon'} \sum_{z'} [s(t+1, a', z', \epsilon')b(t, a, z, \epsilon) \\
&\quad + (1 - s(t+1, a', z', \epsilon'))\kappa b(t, a, z, \epsilon) q(t+1, z', a'')]\pi(z, z'))dF(\epsilon')
\end{aligned}
$$

where a'' is the asset choice evaluated at $(t+1, a', z', \epsilon')$. Simplifying b:

$$\Longleftrightarrow$$

$$q(t,z,a') = \frac{1}{(1+r)}\Big(\int_{\epsilon'} \sum_{z'} [s(t+1,a',z',\epsilon') + (1-s(t+1,a',z',\epsilon'))\kappa q(t+1,z',a'')]\pi(z,z'))dF(\epsilon')\Big) \tag{3}$$

The price q is therefore a function of age, the permanent component of income, and end-of-period net asset position. Positive net assets earn the risk-free interest rate r, because households set $s = 1$ in that case.

3.1 Model mechanisms

Our model departs from the standard framework by allowing for delinquency, in the form of $s < 1$. Delinquency costs are twofold: an exogenous utility cost governed by $\phi(s)$ and an endogenous effect through the price of new debt, q. In this section we outline how delinquency affects household decisions.

First, indebted households face a static trade-off between delinquency, $1 - s$, and consumption, c. This trade-off is described by the following optimality condition:

$$\underbrace{-\frac{\partial \phi(s)}{\partial s}}_{\text{marginal gain from paying down debt}} = \underbrace{u'(c)\, a\,(q_t(z,a')\,\kappa - 1)}_{\text{opportunity cost of paying down debt}} \tag{4}$$

Paying down debt, in the form of increasing s towards 1, delivers a marginal gain equal to the marginal improvement in the disutility from delinquency, as shown in the left hand side of the equation. Such gain equates the opportunity cost of paying down debt, in the form of giving up consumption. Note that this equation holds for indebted households (i.e., $a < 0$), because otherwise $s = 1$. Moreover, $(q_t(z,a')\,\kappa - 1) < 0$. Hence the opportunity cost is larger the larger the debt stock, and the lower the price of new debt.

Delinquency also alters the inter-temporal trade-off faced by households. In particular, the Euler equation for assets is:

$$\frac{\partial u(c)}{\partial c}\left\{\frac{\partial q_t(z,a')}{\partial a'} b + q_t(z,a')\right\} = \beta E_t \frac{\partial u(c')}{\partial c'}\{(1-s')\,\kappa q_{t+1}(z',a'') + s'\} \tag{5}$$

When q prices are constant and it is not possible to be delinquent, we obtain the standard incomplete-markets Euler equation according to which a household equates the marginal utility gain of consuming a dollar today, to the gain of not consuming it, saving, and consuming tomorrow the interest-accrued dollar. With delinquency, this condition is affected in four ways. First, borrowing households have q available marginal units to consume, and this price is endogenous. Second, savings decisions (i.e., a') affect the pricing schedule and therefore change the amount of available resources. Note that a' can be affected either by delinquency, $1-s$, or debt inflows, b. Third, the household may be delinquent tomorrow, and this affects the right hand side of the Euler equation through s'. Finally, the inter-temporal savings decision takes into account expected prices tomorrow, $E_t q_{t+1}(z', a'')$.

3.2 Responses to transitory income shocks

In this section we describe households' responses to a purely transitory income shock (e.g.: tax rebate) τ. In what follows, we consider local perturbations (i.e.: small shocks). Let us first differentiate the budget constraint and define various marginal propensities:

$$\frac{\partial c}{\partial \tau} + \frac{\partial q(t, a', z)}{\partial a'} \left(\frac{\partial b}{\partial \tau} - a \frac{\partial s}{\partial \tau} \kappa \right) b + q(t, a', z) \frac{\partial b}{\partial \tau} - a \frac{\partial s}{\partial \tau} = 1 \tag{6}$$

The well-known marginal propensity to consume is defined as $MPC = \frac{\partial c}{\partial \tau}$. The sensitivity of $\frac{\partial b}{\partial \tau}$ can be thought of as a marginal propensity to save (MPS). Note that, for positive asset holders, $a' = b$. For these households, MPC and MPS sum to $1 + r$, as in the standard consumption-saving model.

Our model expands on the concept of Marginal Propensity to Repay Debt (MPRD), which is 0 when $a > 0$ and is otherwise defined as follows:

$$MPRD = -\frac{\partial s}{\partial \tau} a \tag{7}$$

Indebted households can use the tax rebate τ to improve their net asset position in two ways. First, they can borrow less: according to our definition, this will show up in the MPS. Second, they can use the transfer to repay some of the accumulated debt. In our model, this happens through a reduction in the extent to which households are delinquent on their stock of debt.

Therefore, we extend the standard incomplete markets model in two ways. First, we give a new tool to indebted households, the MPRD. Second, the tax rebate affects the pricing schedule of new debt, via its effect on end-of-period net assets a', and this feeds back into the budget constraint, affecting households' marginal propensities.

In line with this characterization, our quantitative analysis will consider three main models. First, a standard incomplete markets model where $s = 1$ and $q = \frac{1}{1+r}$ for all households. In this framework, the MPRD is always 0. Second, we allow for endogenous s, but we keep q constant. In this setting, if households want to use part of the tax rebate to pay down debt, they have to divert resources away from consumption, as can be seen from Equation (6). This may introduce a negative correlation between MPC and MPRD. This tradeoff, however, also operates through the static optimal condition shown in Equation (4), and the dynamic decision described in Equation (5). We will show later how the MPRD increases with debt, but only slightly. Finally, in our full model we also allow for endogenous prices q. This works as an endogenous borrowing constraint, which limits the extent to which households want to be delinquent. Moreover, it affects households' responses to the tax rebate. On the one hand, households that use the check to pay down debt improve their asset position, their price schedule and hence, in turn, can borrow more, thus potentially pushing up the MPC. On the other hand this channel gives an additional incentive to use the check to pay down debt. As such, the MPRD is higher, and more positively correlated with debt, as we show in the next section.

Table 1: Calibration

Targeted moments	Data	Model
Share of Households with debt	0.15	0.15
Share of delinquent households	0.08	0.08
Internally calibrated parameters		
β		0.9974
ϕ_1		0.0025
Pre-defined parameters		
r		0.0074
γ		2
σ_η		0.0025
σ_ϵ		0.0125
σ_{η_0}		0.0375

4 Quantitative exploration

4.1 Calibration

The model period is one quarter. Our calibration closely follows Kaplan and Violante (2010). Households enter the labor market when they are 25 years old, retire after 35 years and die with certainty at 95. For simplicity, we do not use age-varying survival rates as in Kaplan and Violante (2010), but set them all to 0 except for the last time period. Utility from consumption follows a CRRA specification, with risk aversion parameter γ equal to 2. The risk-free interest rate, r, is 3 percent (annualized). The parameters governing the income process are the variance of permanent shocks, σ_η, the variance of transitory shocks, σ_ϵ, and the cross-sectional initial variance of the permanent component, σ_{η_0}. Each of these is the quarterly translation of what used by Kaplan and Violante (2010). Retirement income is calibrated as in Kaplan and Violante (2010) and mimics the actual US system. Social security transfers are set to 90% of pre-retirement earnings up to a given bend point, beyond which they are 32% up to a second bend point, and 18% thereafter. The two bend points are set at 18% and 110% of the cross-sectional average of labor earnings.

We impose a natural debt limit, such that households do not die in debt. We also assume that the initial wealth of all households is zero. We follow Herkenhoff (2019) and parameterize $\phi(s) = \phi_1 \frac{1-s}{s}$. For simplicity, we also set $\kappa = 1$.

The two remaining parameters, β and ϕ_1, are internally calibrated by targeting two moments: (i) the share of households with negative net assets (15%), and (ii) the share of delinquent households (8%). For the first target we use the share of households with negative liquid wealth, as in Kaplan, Moll and Violante (2018). We see this moment as an intermediate choice within the possible ranges used in the literature, and in particular as a balanced approach to the ongoing debate on whether to use net worth or gross unsecured debt as a calibration target. For instance, an alternative could be to target the fraction of households with negative worth, which Chatterjee, Corbae, Nakajima and Ríos-Rull (2007) report to be 6.7% in the 2001 Survey of Consumer Finances (SCF). Alternatively, 39% of US households in the 2016 SCF carried a balance on their credit card, as reported by Exler and Tertilt (2020). Our target of 15% stands roughly in the middle of these two values. The discount factor is particularly informative for this moment, as the share of households with negative assets monotonically falls with β, all other things equal. The second target is calculated using data from the Federal Reserve Bank of New York Consumer Credit Panel. In particular, we follow Athreya, Sánchez, Tam and Young (2018) and calculate the share of individuals who are more than 60 days delinquent but not bankrupt. we average this moment over the period 2004-2013. In the model, this moment denotes the households with $s < 1$. Intuitively, the higher ϕ_1, the lower the share of delinquent households, as the disutility cost from delinquency increases.

In the following sections, we also consider two alternative versions of the model. First, one in which $s = 1$ and $q = \frac{1}{1+r}$ for all households, obtained by setting $\phi_1 = 0$. We label this model SIM. Second, a version of our baseline model in which $q = \frac{1}{1+r}$ is constant for all households. We label this model CQ. Recalibrating the parameters of the CQ model is particularly challenging because the disutility from delinquency is typically not enough to offset the fact that all households can borrow at the risk free rate, regardless of their stock of debt and the extent to which they are delinquent. We circumvent this issue by introducing a fixed utility cost, $\phi_3 = 0.02$, which applies to all delinquent households and whose value we keep the same in the baseline and CQ model. For illustrative purposes, we also keep the other parameters unchanged in all models.

4.2 Explaining the three facts on households' responses to stimulus checks

We revisit the three main empirical facts that we showed in Section 2 through the lens of the model, and show how our baseline full model is able to match them all, whereas alternative versions fail along certain dimensions. To do so, we focus on the stationary distribution of households in each model. We evaluate how each household's decisions compare with and without the receipt of a transitory and unexpected lump-sum transfer. Then we compute the marginal propensities as the fraction of the transfer that it used for a given choice. For example, the formal definition of a MPC for a household with a certain age, net worth, permanent and transitory income level is $MPC = \frac{c(t,a,z,\epsilon,\tau)-c(t,a,z,\epsilon,0)}{\tau}$.

4.2.1 The MPRD

We start by looking at whether households in our models use the transfer to pay down debt, and by how much. The MPRD is mechanically 0 for all households in the SIM model without delinquency, as households can only consume or save and always repay their debts in full in every period. Introducing delinquency, but keeping debt prices constant, raises the average MPRD only slightly to 1 cent per transfer dollar, as shown in Table 2. We focus our attention only on households with debt, as households in the data can hold debt and

Table 2: The average MPRD

	SIM	CQ	Baseline
Average MPRD	0	0.01	0.11

positive assets at the same time, while our one-asset model does not feature this possibility.[9] As we show in Figure 4, the vast majority of indebted households in this model do not use any part of the stimulus check to pay down debt. Indeed, doing so only helps them reducing the disutility from delinquency, but diverts resources away from consumption.

The average MPRD goes up substantially in the full model. Allowing for endogenous debt prices gives an additional incentive to indebted households to use the check to pay down debt since, by doing so, they can borrow more and thus consume more, today as well as in the future. This behavior is broadly consistent with the empirical observation that credit scores improved since the onset of the COVID-19 pandemic, as discussed by Kowalik et al. (2021). While other factors, such as forbearance programs and lower credit card utilization, may have played an important role, repayment of outstanding debt is also positively associated with credit scores.[10] As we show in Figure 4, a higher average MPRD is obtained both through an extensive and intensive margin effect. First, the share of indebted households not using the check at all to pay down debt falls to 55%, quite close to the 45% documented in the data. Second, the distribution of positive MPRD shifts to the right, since some households decide to use a relatively large portion of their transfer to pay down debt.

While our full model allows us to get much closer to the data than alternative settings, we are still falling short of the empirical importance of the MPRD. One potential reason is that households surveyed in the SCE report a positive MPRD whenever they pay down any type of debt, including mortgage debt. In contrast, our model is not suited to study collateralized debt, and only features unsecured debt, which in the U.S. accounts for less than a fifth of aggregate debt.[11] Nevertheless, the average MPRD in our full model is higher than the average MPC, consistent with what observed in the data. The average MPC is also lower than in the data. This is most likely because the model only considers notional consumption of nondurable goods, while the SCE reports marginal spending in any expenditure item.[12]

4.2.2 MPRD and debt

Which households have the highest MPRD? In Figure 5 we show that, even in the CQ model with constant debt prices, there is a positive correlation between debt-to-income ratios and MPRD, consistent with the empirical evidence shown in Section 2. This relationship,

[9]As we have shown in Section 2, our three main empirical facts also hold when restricting the sample to households with positive debt.

[10]Some commentators have suggested there may have been a direct link between using the stimulus check to pay down debt and reduction in credit scores. To the best of our knowledge, this link has not been formally established empirically though.

[11]As reported by the New York Fed's Consumer Credit Panel, credit card debt accounted for 5% of total debt balance in the last quarter of 2021, student debt for 10%, auto loans 9% and mortgage debt 70%.

[12]See Laibson, Maxted and Moll (2022) for a recent discussion of the relationship between MPC and marginal propensities to spend.

Figure 4: The MPRD distribution

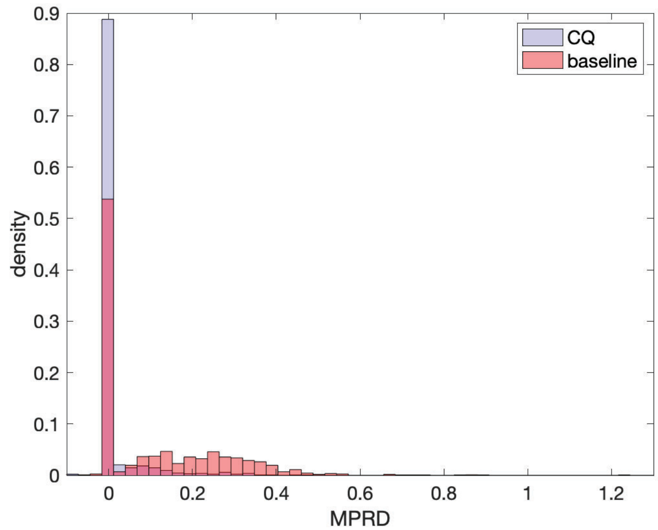

Notes: The figure plots the histogram of MPRD in the full model (baseline) and model with delinquency and constant debt prices (CQ), for all households with positive debt.

however, is much stronger in the full model, with a jump up for small levels of debt, as in the data.

Endogenous debt prices have two effects. First, they add an additional disincentive to hold large debt balances. As a result, the stationary distribution of debt-to-income ratios shifts to the left, with the highest quintile being half of what obtained in the constant prices (CQ) model. Second, there is an additional positive effect of using the transfer to pay down debt, working through debt prices q, as we showed in Equation (6). This effect looks quantitatively stronger the larger debt-to-income ratios, opening an increasing gap between the baseline and CQ models.

4.2.3 MPC and debt

Finally, we show that MPCs fall with debt-to-income ratios, in line with the data. As can be seen from Figure 6, this relationship only holds in our baseline model.

In the SIM model without delinquency and with a natural borrowing limit, MPCs fall with net worth. This well-known finding in the literature is related to the concavity of the consumption function and precautionary savings behavior typical of buffer stock models. If we restrict the attention to indebted households, however, there is not a clear relationship between MPCs and debt-to-income ratios. The same holds true in the CQ model with delinquency and fixed prices.

In the full baseline model, instead, MPCs are highest for households with small but positive debt-to-income ratios. In fact, taking a broader view, MPCs peak for households

Figure 5: The MPRD and debt-to-income ratios

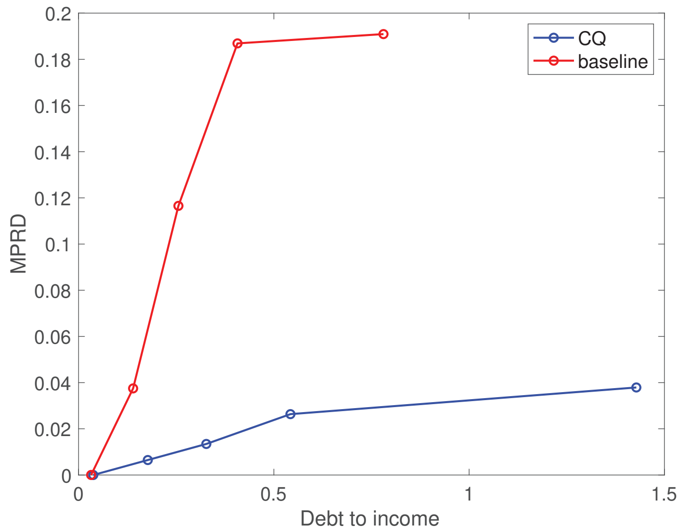

Notes: The figure plots the relationship between MPRD and debt-to-income ratios in the full model (baseline) and model with delinquency and constant debt prices (CQ), for all households with positive debt. Debt-to-income ratios are defined as the stock of household debt in period t divided by 4 times quarterly income in period t. For each quintile of debt-to-income ratios, we plot the average debt-to-income ratio on the horizontal axis, and the average MPRD on the vertical axis.

with small – in absolute value – ratios of net worth to income.[13] This behavior is the mirror image of what shown in the previous section about the MPRD. When debt-to-income ratios are low, households have little incentive to be delinquent, and especially to use the check to pay down debt. In contrast, they have a relatively high marginal utility of consumption, due to the effect exerted by endogenous borrowing constraints arising from the debt price schedule. As such, they have a relatively high MPC. As debt-to-income ratios increase, households find it preferable to pay down debt. Hence, as shown before, the MPRD increases, and the MPC falls, since paying down debt diverts resources away from consumption in the near term.

We can also make use of Equation (4) and the consumption policy functions to further understand why the baseline model works this way. For high levels of debt, and for given income and age, consumption is locally convex in net assets a. This is because the repayment policy, s, is convex in net assets, this effect offsets the mild concavity of q relative to a, and in turn affects the convexity of the consumption function through Equation (4). A convex consumption function implies that, for indebted households, MPCs increase with net assets (decrease with debt). As debt gets smaller, the repayment policy function s becomes concave in net assets; moreover, households typically repay their debts in their near entirety. As such, the consumption function becomes concave, and remains so as net assets become positive, as is standard in the SIM model. As a result, MPCs peak for low debt, and fall as debt gets

[13] As such, our results are broadly consistent with Kaplan et al. (2018), who find that MPCs are higher for households with little or no liquid wealth than for households with more negative liquid wealth.

Figure 6: The MPC and debt-to-income ratios

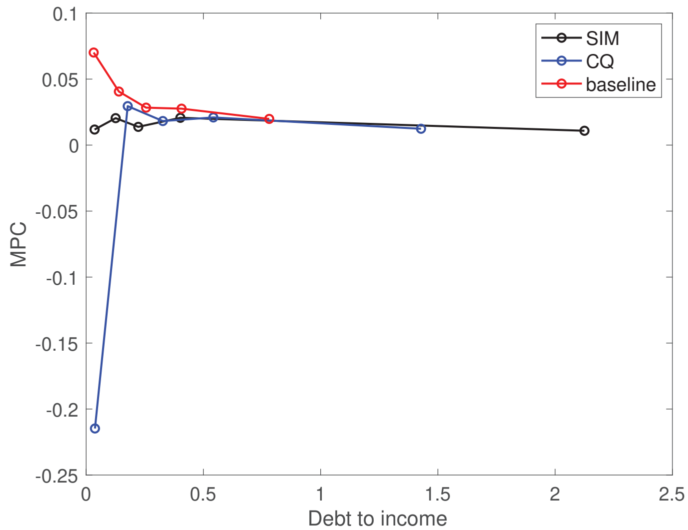

Notes: The figure plots the relationship between MPC and debt-to-income ratios in the full model (baseline) and model with delinquency and constant debt prices (CQ), for all households with positive debt. Debt-to-income ratios are defined as the stock of household debt in period t divided by 4 times quarterly income in period t. For each quintile of debt-to-income ratios, we plot the average debt-to-income ratio on the horizontal axis, and the average MPC on the vertical axis.

larger or net assets turn positive.

5 Stimulus vs insurance

TO BE COMPLETED

6 Conclusions

TO BE COMPLETED

References

Agarwal, Sumit, Chunlin Liu, and Nicholas S Souleles, "The reaction of consumer spending and debt to tax rebates—evidence from consumer credit data," *Journal of political Economy*, 2007, *115* (6), 986–1019.

Arellano, Cristina, Xavier Mateos-Planas, and José-Víctor Ríos-Rull, "Partial default," Technical Report, National Bureau of Economic Research 2019.

Athreya, Kartik, Juan M Sánchez, Xuan S Tam, and Eric R Young, "Bankruptcy and delinquency in a model of unsecured debt," *International Economic Review*, 2018, *59* (2), 593–623.

Blundell, Richard, Luigi Pistaferri, and Ian Preston, "Consumption inequality and partial insurance," *American Economic Review*, 2008, *98* (5), 1887–1921.

Chatterjee, Satyajit, Dean Corbae, Makoto Nakajima, and José-Víctor Ríos-Rull, "A quantitative theory of unsecured consumer credit with risk of default," *Econometrica*, 2007, *75* (6), 1525–1589.

Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber, "How did US consumers use their stimulus payments?," Technical Report, National Bureau of Economic Research 2020.

Exler, Florian and Michèle Tertilt, "Consumer Debt and Default: A Macro Perspective," 2020.

Fagereng, Andreas, Martin B Holm, and Gisle J Natvik, "MPC heterogeneity and household balance sheets," *American Economic Journal: Macroeconomics*, 2021, *13* (4), 1–54.

Hedlund, Aaron, Fatih Karahan, Kurt Mitman, and Serdar Ozkan, "Monetary policy, heterogeneity, and the housing channel," in "2017 Meeting Papers," Vol. 1610 2017.

Herkenhoff, Kyle F, "The impact of consumer credit access on unemployment," *The Review of Economic Studies*, 2019, *86* (6), 2605–2642.

Hisnanick, John J and Andreas Kern, "The 2008 tax rebate and US household debt," *Applied Economics Letters*, 2018, *25* (9), 592–596.

Kaplan, Greg and Giovanni L Violante, "How much consumption insurance beyond self-insurance?," *American Economic Journal: Macroeconomics*, 2010, *2* (4), 53–87.

_ , Benjamin Moll, and Giovanni L Violante, "Monetary policy according to HANK," *American Economic Review*, 2018, *108* (3), 697–743.

Kowalik, Michael, Lily Y. Liu, and Xiyu Wang, "Credit Scores Since the COVID-19 Outbreak," Technical Report, Federal Reserve Bank of Boston 2021.

Kueng, Lorenz, "Excess sensitivity of high-income consumers," *The Quarterly Journal of Economics*, 2018, *133* (4), 1693–1751.

Laibson, David, Peter Maxted, and Benjamin Moll, "A Simple Mapping from MPCs to MPXs," Technical Report, National Bureau of Economic Research 2022.

Lewis, Daniel J, Davide Melcangi, and Laura Pilossoph, "Latent heterogeneity in the marginal propensity to consume," *FRB of New York Staff Report*, 2021, (902).

Misra, Kanishka and Paolo Surico, "Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs," *American Economic Journal: Macroeconomics*, 2014, *6* (4), 84–106.

Sahm, Claudia R., Matthew D. Shapiro, and Joel Slemrod, "Household Response to the 2008 Tax Rebate: Survey Evidence and Aggregate Implications," in "Tax Policy and the Economy, Volume 24" NBER Chapters, National Bureau of Economic Research, Inc, 2010, pp. 69–110.

Shapiro, Matthew D and Joel Slemrod, "Did the 2008 tax rebates stimulate spending?," *American Economic Review*, 2009, *99* (2), 374–79.

A Data appendix

In Figure 7 we repeat the analysis of Section 2, but restricting the sample to households with positive gross unsecured debt. Each dot represents quintiles of debt-to-income ratios, with the average within-quintile ratio on the horizontal axis, and the average MPRD (MPC) on the vertical. As such, we plot the relationship in the exact same way as we do in the model in Section 4.2.2 and 4.2.3. As in the baseline model, MPRDs are increasing and concave in debt-to-income ratios, while MPCs are decreasing and convex.

Figure 7: MPRD, MPC and debt-to-income ratio: quintiles

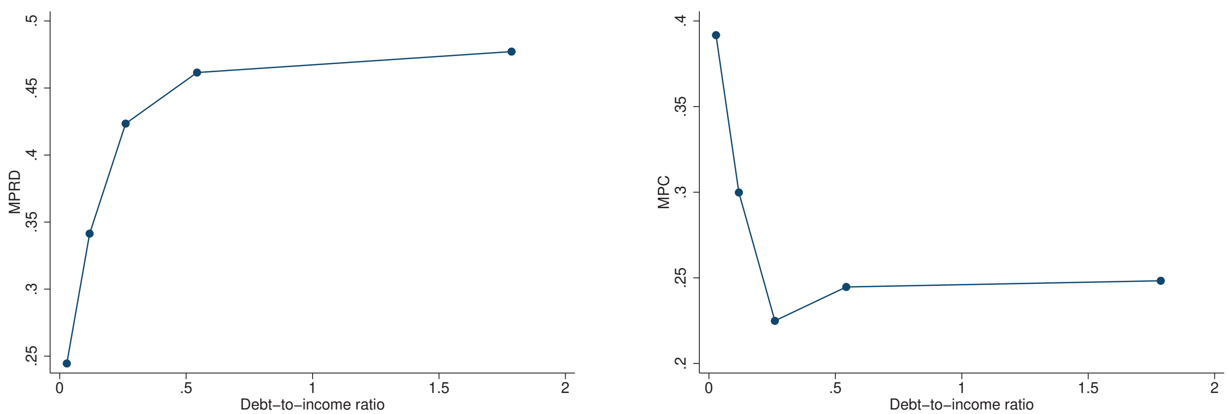

(a) MPRD (b) MPC

Notes: SCE empirical sample restricted to households with positive gross unsecured debt.

B Computational appendix

B.1 Solution method

1. start with a guess for the function $q(t, \cdot, \cdot, \cdot)$.
2. for a given guess of $q(\cdot)$, start with period T, implied consumption is known, an we impose $s^*(a) = 1 \; \forall a$ (we could do otherwise), calculate $E_t V(t+1, a', b', z', \epsilon' | t, a, b, z, \epsilon)$
3. move to $t = T - 1$. For given a today, solve for optimal $s(a)$ and $c(a)$ today using bisection and find intersection
4. continue to $t = T - 2$, all the way until $t = 0$.
5. update q according to Equation (3).
6. repeat steps 2-5 until convergence.