

NBER WORKING PAPER SERIES

STIMULUS THROUGH INSURANCE:
THE MARGINAL PROPENSITY TO REPAY DEBT

Gizem Koşar
Davide Melcangi
Laura Pilossoph
David G. Wiczer

Working Paper 34399
<http://www.nber.org/papers/w34399>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2025

Gizem Koşar, Davide Melcangi, Laura Pilossoph, and David Wiczer have not received any source of research support nor do they have any financial relationships or potential conflict of interests to report in conducting this research. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Banks of New York or Atlanta, the Federal Reserve System, or the National Bureau of Economic Research. We thank Aaron Jalca, Aidan Toner-Rodgers, Chris Gresh, Michael C Gannon, Srinidhi Narayanan, and Yasin Simsek for excellent research assistance. For helpful comments, we also thank Orazio Attanasio, Adrien Auclert, David Berger, Corina Boar, Gideon Bornstein, Michael Boutros (discussant), Keshav Dogra, Carlo Galli, Felicia Ionescu, Eva Janssens (discussant), Gregor Jarosch, Hyeyoon Jung, Greg Kaplan, Donghoon Lee, Igor Livshits (discussant), Daniel Mangrum, Kurt Mitman (discussant), Ben Moll, Monika Piazzesi, Victor Ríos-Rull, Joelle Scally, and conference and seminar participants at the NBER Summer Institute 2022, AEA Meeting 2022, BSE Summer Forum 2023, Columbia University, Queens University, UC3M, University of Copenhagen, FRB San Francisco, Bank of Canada, Bank of Mexico, Danmarks Nationalbank, Rutgers University, XAmsterdam macro workshop, Esade Macro meetings, and SED 2024.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w34399>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Gizem Koşar, Davide Melcangi, Laura Pilossoph, and David G. Wiczer. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Stimulus through Insurance: the Marginal Propensity to Repay Debt
Gizem Koşar, Davide Melcangi, Laura Pilossoph, and David G. Wiczer
NBER Working Paper No. 34399
October 2025
JEL No. E21, E62

ABSTRACT

Using detailed microdata, we document that households often use “stimulus” checks to pay down debt, especially those with low net wealth-to-income ratios. To rationalize these patterns, we introduce an empirically plausible borrowing price schedule into an otherwise standard incomplete markets model. Because interest rates rise with debt, borrowers have increasingly larger incentives to use an additional dollar to reduce debt service payments rather than consume. Using our calibrated model, we then study whether and how this marginal propensity to repay debt (MPRD) alters the aggregate implications of fiscal transfers. We uncover a trade-off between stimulus and insurance, as high-debt individuals gain considerably from transfers, but consume relatively little immediately. This mechanism lowers the immediate stimulus effect of fiscal transfers, but sustains aggregate consumption for longer.

Gizem Koşar
Federal Reserve Bank of New York
gizem.kosar@gmail.com

Davide Melcangi
Federal Reserve Bank of New York
Davide.Melcangi@ny.frb.org

Laura Pilossoph
Duke University
and NBER
pilossoph@gmail.com

David G. Wiczer
Federal Reserve Bank of Atlanta
and IZA
wiczerd@gmail.com

1 Introduction

Households frequently use stimulus checks to pay down existing debt. For example, after receiving the 2008 rebates as part of the Economic Stimulus Act, 52% of households reported that they used the money to mostly pay down debt, while only 20% reported that they mostly spent it (Sahm et al. (2010)). However, despite the disproportionate use of these checks for debt repayment, both academic and public discussions have instead focused on their role in stabilizing aggregate demand through their prompt impact on spending.¹ Indeed, this is why these transfers have come to be called “stimulus checks,” and why their success is often measured by the extent to which they are immediately spent.

In this paper, we study the transmission of fiscal transfers through household debt repayments. We provide novel empirical evidence demonstrating that it is households with low net-liquid-wealth-to-income ratios that are more likely to use cash windfalls to pay down debt, rather than spending immediately to consume. We show that a standard consumption-savings model rationalizes these patterns if interest rates rise with debt, so that debt reduction incentives lean against the typical consumption smoothing ones, and provide new evidence illustrating that this debt-price schedule is empirically plausible. In this environment, a disconnect between stimulus and insurance motives for cash transfers arises—both across households and over time—as borrowers use transfers to pay down debt, pushing consumption into the future. While an initial tension exists between distributing funds for maximum welfare versus maximum consumption, the stimulus effect eventually materializes *through* the insurance mechanism, as households who had the largest welfare gain eventually consume more and for longer after paying down debt. We illustrate that this tension can change both the evaluation and the design of expansionary fiscal policy.

Our main empirical evidence on the use of stimulus checks comes from microdata collected during the COVID-19 pandemic 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 sparse, most likely due to data limitations.² We make progress by eliciting responses directly from surveyed households in the SCE.³ We document three main facts. First, households use a third of their transfers to pay down debt, a response which is as large as the average marginal propensity to consume (MPC) that usually takes center stage. This is consistent with an

¹For example, Bush described the 2008 checks as a “booster shot,” stating “It’s clear our economy has slowed. But the good news is we anticipated this and took decisive action to bolster the economy by passing a growth package that will put money into the hands of American workers and businesses.”

²One notable exception is Agarwal et al. (2007), who estimate debt responses to 2001 tax rebates in the United States. Consistent with our mechanism, they find that consumers initially used the checks to pay down debt, which stimulated spending in the medium-run.

³Coibion et al. (2020) adopt a similar approach, using another survey of U.S. households. We discuss in Section 2 how our results relate to theirs.

under-emphasized finding in existing studies that measure debt repayment responses to fiscal stimulus.⁴ Second, leveraging the unique aspect of our data which links repayment propensities to measures of liquid wealth and unsecured debt, we show that households with low net liquid wealth-to-income ratios are *more* likely to pay down debt and more likely to improve their net asset positions. Third, and relatedly, households with lower net liquid wealth-to-income ratios have lower MPCs. These facts are not specific to the pandemic and hold through a battery of robustness checks. While prior work has examined MPC heterogeneity using gross liquid wealth or net total wealth,⁵ ours is unique in its use of net liquid wealth and its focus on net borrowers, which is the most direct way to link empirical MPCs and marginal propensities to pay down debt (MPRDs) to models of household interactions with credit markets. By presenting an extensive documentation of the heterogeneity in MPRDs in this population, we offer a novel perspective on the usage of fiscal transfers across households.

Standard incomplete markets models have difficulty generating the cross-sectional relationships we document. In these frameworks, households with lower net liquid wealth-to-income ratios are more constrained, and thus have higher MPCs (Kaplan and Violante (2022)). We show that the simple introduction of a debt price schedule—which we later illustrate is empirically plausible—can reconcile the model with the facts we document. In a tractable two-period framework, we first show how this change modifies the Euler Equation and introduces a debt-service reduction motive, akin to the Generalized Euler Equation that commonly appears in endogenous default models (Arellano et al. (2023)). When households’ net asset position is negative and the interest rate is rising in debt, the consumption function can become convex in assets, rather than concave as is standard in incomplete market models (Carroll and Kimball (1996)). This means that the consumption function can become flatter the more indebted a household is. As a consequence, MPCs in that region are increasing in assets. We formalize the conditions that sustain this result. We also discuss how alternative deviations from the constant debt-price assumption, including a borrowing wedge (Achdou et al. (2022)) and a price schedule generated through endogenous default, alter the shape of the consumption function, both qualitatively and quantitatively.

We then calibrate a quantitative infinite-horizon version of the model to match our empirical evidence. The debt pricing schedule creates a strong savings motive, even among those with negative net assets. Rather than facing a hard borrowing constraint, raising MPCs towards one, high levels of debt generate a high marginal propensity to repay debt. There are fewer households at the highest levels of debt, and those debt levels are associated with lower MPCs.

The pricing schedule required by our model to generate these patterns is empirically plausible. To illustrate this, we introduce new facts about the dynamics of credit card borrowing rates

⁴For a comprehensive list of papers documenting this fact, see Table 1.

⁵For a full discussion, see Section 2.

at the individual level. Using the New York Fed’s Consumer Credit Panel/Equifax (CCP), we track borrowers over time, measuring how their effective interest rates change with debt levels. We find that debt reductions are associated with changes in individuals’ effective interest rates in ways that are consistent with our calibrated pricing schedule.

With our calibrated quantitative model, we next ask how the introduction of a debt-price schedule, and thus realistic debt-repayment responses to fiscal stimulus, matters for their evaluation. First, the welfare value of transfers becomes divorced from the instantaneous consumption response they generate. In a standard incomplete markets model, those who have the highest MPC out of the rebate also value it the most. If, instead, borrowing costs rise with debt, some households have high welfare gains from transfers despite having a small MPC upon their receipt; in other words, the insurance value of transfers no longer perfectly aligns with their instantaneous stimulus value, and these two benefits may conflict. This decoupling also arises when debt prices are determined in equilibrium in a model with endogenous default.

Second, and relatedly, when the interest rate is non-constant, the inter-temporal MPC is quite different, both in the aggregate and across the distribution of households. Households with high MPRDs delay the consumption effect of the stimulus, flattening and extending the impulse response. At the same time, there is an increase in the upon-impact spending of households with little debt, relative to the workhorse model with exogenous liquidity constraints. We find that non-constant interest rates are associated with more persistent aggregate spending effects.

Against this backdrop, our mechanism implies that policymakers may face non-trivial trade-offs when targeting transfers, at least in partial equilibrium. For instance, they may wish to allocate transfers differently when favoring welfare over initial stimulus, or long-run rather than short-run fiscal multipliers. We show this formally with an optimal policy exercise that allows the planner to target by either income or debt service. Using our calibrated model, we conclude that the stimulus of 2020 generated larger welfare gains than would be seen through the lens of a constant-interest-rate incomplete markets model. Households with the highest marginal utility of consumption realize disproportionately large benefits by reducing their debt service payments rather than consuming.

The paper is organized as follows. In Section 2, we present empirical facts on the MPRD and the MPC using household data from the U.S.. Section 3 shows how debt-sensitive interest rates can rationalize these facts with an illustrative two-period model. In Section 4, we extend this to a calibrated, quantitative model and investigate aggregate and distributional effects of stimulus payments. Finally, we conclude.

2 Empirical facts on households’ responses to stimulus checks

In this section, we document three main facts regarding the way households responded to the lump-sum transfers that were given as part of the CARES Act. The CARES Act was a large stimulus package passed by the U.S. Federal government on March 27th, 2020. As part of this package, all qualifying adults received a one-time transfer of up to \$1200, with \$500 per additional child. We show the external validity of the facts we document using results from other transfer episodes in the U.S., other data sets, and responses to questions about hypothetical income windfalls.

2.1 Data

Our data come from a special survey module fielded in June 2020 as part of the New York Fed’s Survey of Consumer Expectations (SCE). The SCE is a monthly internet-based survey of a rotating panel of approximately 1300 heads of household from across the United States. 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 may also be asked to participate in additional modules every month and receive extra incentives when they do. The special module used in the main analysis that follows was fielded to the respondents who rotated out of the SCE. However, we also use the main panel of respondents to show that our results are not specific to the pandemic, and also hold in a hypothetical setting, that is, using survey responses to hypothetical questions about how individuals would spend cash windfalls.⁶ The characteristics of both samples can be found in Tables B.1 and B.2 in Appendix B.

Our analysis primarily focuses on questions about the receipt and usage of the stimulus checks – formally called Economic Impact Payments. 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. Around 89% of the respondents in our sample received the stimulus payments by the time of their interview, with the average (median) payment received being \$2080 (\$2400) (see Appendix B). The respondents who reported receipt of the stimulus checks were then asked a question regarding the allocation of this payment in the following form:

Please indicate what share of the government payment you have already used to or expect to use

⁶See Manski (2004) for details on stated-choice experiments and how they might map to actual decisions, Stantcheva (2022) for a recent review of different design considerations for hypothetical scenarios, and Fuster, Kaplan and Zafar (2021) for an example of using hypothetical scenarios to elicit MPCs.

to ...

<i>Save or invest</i>	_____	<i>percent</i>
<i>Spend or donate</i>	_____	<i>percent</i>
<i>Pay down debts</i>	_____	<i>percent</i>
<i>Total</i>	100	<i>percent</i>

where the responses add up to 100.⁷ In the rest of this section, we refer to the “Save or invest” allocation as the marginal propensity to save (MPS), the “Spend or donate” allocation as the marginal propensity to consume (MPC),⁸ and the “Pay down debt” allocation as the marginal propensity to repay debt (MPRD). We also define the marginal propensity to adjust the net asset position (MPA) as $MPA = MPRD + MPS$. We return to a detailed discussion of how these concepts line up with our model measures once we have presented the model. Finally, in what follows, we focus only on those who already reported receiving a payment at the time of the survey and consider their allocation of the payment (1423 observations).


One concern regarding reported allocations of stimulus checks might be whether these align well with revealed preference estimates. Recent papers in the literature have shown that reported MPCs indeed align well with “revealed-preference” MPCs (those estimated from data on spending). For example, Parker and Souleles (2019) find that average propensities to spend from both methods are similar. In a more recent paper, Kotsogiannis and Sakellaris (2025) compare reported and revealed-preference MPCs for the same households and find that they align closely. They also find that the distribution of MPCs is very similar when comparing reported MPCs out of lottery winnings and hypothetical ones for observationally equivalent non-winners. Similarly, in concurrent work, Colarieti et al. (2024) compares responses to various hypothetical scenarios with those estimated from realized actions in other studies and find that participants’ reported behavior under the hypotheticals closely aligns with estimated real-life behaviors.

The timing of the survey is important to keep in mind when interpreting answers to this question. By fielding the survey in June 2020, shortly after the checks were sent, we are eliciting the households’ immediate response. Saving the transfer or using it to pay down debt in the short-run will still eventually lead to consumption in the future: the dynamics of this mechanism are captured in the inter-temporal MPCs highlighted in Auclert et al. (2024) and will be central to our model discussed in the following sections.

⁷The 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.

⁸The follow-up questions ask the respondents to split the “Spend or donate” allocation into separate “Spend” and “Donate” allocations. All of our results hold when we define the MPC using only the “Spend” allocation.

Figure 1: Histograms of MPRD and MPC



Notes. Panel (a) shows the histogram of self-reported MPRDs (the share of government payment used to pay down debts) for the full sample of check recipients. Panel (b) shows the histogram of self-reported MPCs (the share of government payment used to spend or donate). Appendix Figure B.4 repeats this figure for the MPS and MPA. The black dashed line in each panel corresponds to the mean, while the red dash-dotted line corresponds to the median.

Fact 1: The average MPRD across households (32%) is as large as the average MPC across households (30%).

Based on the survey responses, we find that, on average, households used a third of their checks to pay down their debt. 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, making the median MPRD zero, while around 19% of the respondents report using all of their stimulus checks to repay their debt, leading to a bimodal distribution. Restricting attention to households with negative net liquid wealth (Appendix Figures B.3a-B.3b), the average MPRD increases to 48 cents per stimulus dollar, while the average MPC falls to 24 cents.

While the MPC has been the focus of a very large literature on consumption-saving choices (Kaplan and Violante (2022)), the MPRD has received relatively less attention, even though its magnitude is as large as the MPC. Importantly, its similarity in magnitude is evident in different stimulus episodes and in different data sets. Table 1 compiles a comprehensive list of empirical papers that measure MPCs and an MPRDs in episodes of fiscal stimulus in the U.S..⁹ listing alongside them their stimulus episode and data source. Fact 1 above is evident in all

⁹The listed papers with an asterisk use a question wording which corresponds less directly to an MPC or MPRD since they do not give quantitative shares. Instead, the question wording is “[will] you mostly use the check to increase spending, mostly to increase saving, or mostly to pay off debt?” However, Parker and Souleles (2017) show that both “reported preference” types of measures are informative of the average propensities estimated from the “revealed preference” approach which estimates MPCs directly from consumption data. Moreover, Coibion et al. (2020) show that the qualitative and quantitative answers in a single survey are consistent. For a discussion relating micro-level revealed-preference estimates to semi-structural approaches, see Commault (2022).

Table 1: MPCs and MPRDs in Various Studies

Study	Date	Source	MPC	MPRD
This paper (Koşar et al. (2025))	6/20	SCE	.30	.32
Armantier et al. (2020, 2021)	6/20, 7/20, 3/21	SCE	.25-.29	.34-.37
Coibion et al. (2020)*	7/20	Nielsen	.15	.52
Coibion et al. (2020)	7/20	Nielsen	.42	.30
Sahm et al. (2010)*	11/08-12/08	Michigan	.22	.55
Shapiro and Slemrod (2009)*	2/08-6/08	Michigan	.20	.48
Hisnanick and Kern (2018)*	9/08-12/08	SIPP	.28	.53
Boutros (2020)*	07/20	Pulse	.75	.14
Parker et al. (2022)*	6/20-7/20	CEX	.56	.18
Shapiro and Slemrod (2003)*	8/01-10/01	Michigan	.22	.46

Notes. Table reports different papers (column 1), the stimulus episode they study (column 2), the data source (column 3), and statistics related to consumption (column 4) as well as debt repayment (column 5). Papers with an * refer to surveys that use population shares that report “mostly using the checks for” consumption/debt repayment. Papers without an * correspond to MPCs/MPRDs.

other stimulus episodes and is not a unique feature of the checks sent out as part of the CARES Act.

Our findings are also consistent with a similar empirical investigation by Coibion et al. (2020), who instead 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 stimulus money was used to pay down debt. However, they find a higher average MPC than we do in the SCE.¹⁰ Using the qualitative question in the Consumer Expenditure Survey (CEX), Parker et al. (2022) also finds that a non-trivial fraction reports mostly using their checks for debt repayment.

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 reporting that they “mostly” used the rebate for either of the three options. Nonetheless, a key takeaway from these studies is that a large fraction of households self-report using the rebate to pay down debt. Sahm et al. (2010) find that 55% of households report mostly paying down debt with the 2008 stimulus checks, and Shapiro and Slemrod (2009) find that 48% mostly paid down their debt using an earlier sample covering the same stimulus episode. Hisnanick and Kern (2018) use data from the Survey of Income and Program Participation (SIPP) covering the 2008 payment and also find that more than 50% report mostly paying down debt. Taken together the evidence suggests that paying down debt with government stimulus/rebate checks is a general finding that is not specific to the checks sent

¹⁰This higher MPC may reflect differences in sample composition and question wording, because of the different timing of the surveys used. Specifically, the data we use is from the June SCE, relatively soon after households received their payments, while Coibion et al. (2020) use survey responses from July. For this reason, the higher average MPC suggests that some of these payments were temporarily held as savings – or debt repayments – in June, and subsequently used for consumption later on.

out during COVID.¹¹

Fact 2: The MPRD is higher for those with lower net wealth-to-income ratios.

As shown in Figure 2a, MPRDs decrease with net liquid wealth-to-income ratios, hereafter called net wealth-to-income. In the figure, we control for various household characteristics, but raw correlations still show the same qualitative relationship. We focus on households with negative net liquid wealth balances because they will be mostly affected by our theoretical mechanism outlined in the following sections. However, Fact 2 holds for the entire population of households, as we show in Table B.3. Further, we use the net liquid wealth-to-income ratio as our measure of indebtedness because scaling by income removes permanent income differences that scale consumption and wealth, but again, the results do not depend on this.¹² In Appendix Table B.3, we report a regression which formalizes the negative relationship between the MPRD and net wealth-to-income ratios; a one standard deviation increase in net wealth-to-income is associated with an 8.3% decline in the MPRD (from an average of 0.48 to 0.44), keeping everything else constant.

A significant fraction of households report both paying down debt and saving at the same time, resulting in 45% of respondents using the entirety of their check to improve their asset positions. The same negative correlation is also present when we look at the relationship between MPAs and net wealth-to-income ratios. Importantly, this relationship is primarily driven by the MPRD rather than the MPS, as can be seen in Figure B.5. While MPRDs decrease with our net wealth-to-income measure, MPSs in fact have a weakly positive association with net wealth-to-income; even still, the negative relationship between the MPA and net wealth-to-income holds.¹³


In our analysis, we prefer to use net liquid wealth-to-income as the main measure because, as we discuss later, it has a more direct mapping to the canonical model and mechanisms in the heterogeneous-agent literature. In models with an interest rate schedule, this is also generally the correct measure because it captures the lender’s risk, as we show in our endogenous default extension in Appendix D. That being said, our results also hold for gross unsecured debt and net liquid wealth, as we show in the Appendix.

¹¹Outside of the U.S., Crossley et al. (2021) elicit an average MPC of 11% out of a £500 unanticipated, hypothetical shock in the UK, while 22% of the sample report they would pay down debt with the unspent portion of the payments. Fagereng et al. (2021) estimate that, within one year of winning the lottery, Norwegian households used 7% of the prize for debt repayments.

¹²Net liquid wealth is defined as the sum of savings and investments (such as checking and savings accounts, CDs, stocks, bonds, mutual funds, Treasury bonds, excluding retirement) and additional savings or assets (such as cash value in a life insurance policy, a valuable collection for investment purposes, or rights in a trust or estate) less all outstanding debt excluding housing debt. Our results are robust to using different measures of household balance sheets, such as dropping “additional savings”, and using levels of net liquid wealth, gross unsecured debt, or debt-to-income ratios as shown in B.4.

¹³In fact, the relationship is stronger for those on the negative end of the liquid wealth-to-income distribution, as shown in Table B.3 in Appendix B.

Figure 2: MPRDs, MPCs, and net liquid wealth-to-income ratio



Notes. Panel (a) shows a bin scatter of the self-reported, residualized MPRD against the residualized net-liquid wealth to income ratio. Panel (b) shows a bin scatter of the self-reported, residualized MPC by residualized net liquid wealth-to-income ratio. The controls include having a child under age 6, having a child under age 18, marital status, gender, race, and age group of the household head, whether the household head lost their job between March 2020 and June 2020, and whether the household experienced a decline in income between March 2020 and June 2020. Figures B.5a-B.5b in Appendix B repeat this analysis for the MPS and MPA. Appendix Table B.3 shows the analysis for all households, with and without residualization.

A limited set of papers have touched on the relationship between the MPRDs and alternative measures of household resources, generally with outcomes consistent with our findings. Coibion et al. (2020) find that individuals self-identified as liquidity constrained are significantly more likely to pay down debt, although their data has no information on the stock of wealth or debt. Hisnanick and Kern (2018) find that households in the bottom *income* quintile were more likely to use the rebate to mostly pay down debt than richer households. Fagereng et al. (2021) find that Norwegian households in the top quartile of the distribution of *gross liquid assets* use a smaller fraction of the lottery prize to repay debt than illiquid households. Relative to these papers, we instead systematically document the relationship between the MPRD and measures of *net* liquid wealth among US households, and show that it is robust to a host of balance sheet measures, datasets, and episodes. In concurrent work, Colarieti et al. (2024) also find in a novel survey that MPRDs out of hypothetical income windfalls decline with indirect measures of net liquid wealth.

Fact 3: The MPC is lower for those with lower net wealth-to-income ratios.

As the MPC is the mirror image of the MPA based on the design of the survey, it follows that we find a *positive* and statistically significant relationship between the MPC and net wealth-to-income ratio (Figure 2b), among those with negative net liquid wealth. Appendix Table B.3 shows that, when controlling for a host of observable characteristics, a one standard deviation increase in net wealth-to-income is associated with a 12.7% *increase* in the MPC.

Existing empirical evidence on MPC heterogeneity differs in its focus on heterogeneity by

either gross liquid or total net wealth. Fagereng et al. (2021) find no significant relationship between MPCs and *gross debt levels* when controlling for other household characteristics. Crawley and Kuchler (2023) show that MPCs increase with *net wealth* (total, including housing) for Danish households with negative balances, thus broadly in line with our findings for U.S. households using *net liquid wealth*, and then decrease for positive net wealth.

Among net liquid savers, we find no statistically significant correlation between the MPC and net wealth-to-income ratios. Other papers have looked at the relationship between MPCs and different measures of wealth among savers. The lack of statistical association among this group is not surprising in light of that evidence. For example, Sahm et al. (2010), Kueng (2018), Christelis et al. (2019), Fuster et al. (2021), Kotsogiannis and Sakellaris (2025)) find no relationship, while others find a negative and significant relationship (e.g., R. Baker et al. (2023), Ganong et al. (2020), Fagereng et al. (2021)).¹⁴ Several recent papers have highlighted that even liquid households have high MPCs, and offered various explanations.¹⁵ We see this issue as beyond the scope of our paper, which is instead focused on responses to fiscal transfers, and heterogeneity thereof, among net borrowers.

2.2 External validity and robustness of survey responses

MPRDs for hypothetical shocks The facts we have presented thus far are not specific to the COVID-19 episode. The SCE panel also contains pre-COVID information on the response to a hypothetical transitory income shock. Specifically, we use the Household Spending module that has been collected every 4 months since August 2015. As part of this module, households were asked to imagine the hypothetical scenario of finding themselves with 10% higher income and to report which fraction of it they would use to spend, save, or pay down debt. In this setting too, with data from before 2020, Appendix Figures B.7 and B.8 show that (i) MPRDs are large, (ii) MPRDs are increasing in gross unsecured debt-to-income ratios, and (iii) MPCs are decreasing with them. Augmenting this module with household wealth data from the SCE’s annual Housing Module between 2014 and 2019, we further show in Appendix Table B.6 that (iv) MPRDs are declining with net liquid wealth and (v) MPCs are increasing with them, outside of the COVID-19 episode as well. The main difference in behavior between the hypothetical and our 2020 evidence is that the MPRD is slightly larger in the former.

We also consider the approach used in Fuster et al. (2021) to elicit spending responses

¹⁴We find a negative but insignificant relationship between MPC and the log of net liquid wealth for those with positive net balances. We also find a positive and statistically insignificant relationship between the MPC and household income when we control for other observable household characteristics. This aligns with findings by Sahm et al. (2010) and Shapiro and Slemrod (2009), for the 2008 tax rebate. For the same episode, Lewis, Melcangi and Pilossoph (2021) estimate a statistically significant positive correlation between income and estimated spending propensities, as does Kueng (2018) using payments from the Alaska Permanent Fund.

¹⁵See, for instance, Holm et al. (2021), Graham and McDowall (2024), and Ilut and Valchev (2023).

to hypothetical scenarios. These questions were fielded in SCE special modules in 2016 and 2017: in Appendix B.4 we report the results in detail and discuss how the question’s framing differs from the other approaches. These results confirm that MPRDs are large and decline with net wealth-to-income ratios, although this correlation is statistically insignificant. This is consistent with little systematic heterogeneity in elicited spending responses to income gains in their survey, as discussed by Fuster et al. (2021), and may also be traced to the smaller sample size. Interestingly, we find that households who do not have at least two months of liquid funds to face an unexpected income loss have a statistically higher MPRD, consistent with the mechanisms shown in our model.

MPRDs from the Survey of Income and Program Participation Our empirical findings are also not specific to the SCE. To show this, we use the 2020 and 2021 waves of the Survey of Income and Program Participation (SIPP), which provide information on households’ debts and assets. Different from the SCE, the SIPP asks about the predominant use of the checks, rather than the fraction put towards various purposes.

Nevertheless, the SIPP shows the same patterns as the SCE. Overall, the MPRD is quite large: about 38% of respondents reported it as their primary use, while 33% reported primarily consuming their payments. We also validate our second and third facts, that the MPRD rises with debt and falls with net liquid wealth, and associated MPC patterns, as we show in Appendix B.3.

MPRDs with heterogeneous preferences Preference heterogeneity is potentially important for consumption behavior and MPCs (see, for instance, Carroll et al. (2017) and Aguiar et al. (2024) or in the context of mortgage refinancing (Berger et al. (2024))). Using data collected as part of the SCE by Fuster et al. (2021), we find that MPRDs are not systematically driven by heterogeneity in discount factors or risk aversion, as we report in Appendix B.4. Given this null result that preference heterogeneity is not the main driver of our findings, the next section shows how the empirical patterns can be generated by a simple change to agents’ borrowing rates in the standard incomplete markets model.

Robustness of survey responses We validate survey responses and check the correct interpretation of the questions, with open-ended response questions fielded in the SCE in April 2025. Appendix B.5 provides details on this analysis. In a nutshell, we find that SCE respondents correctly interpret debt repayments as paying down credit card debt, car loans, or personal loans, while their saving interpretation includes high-yield savings accounts, buying stocks, bonds, and CDs, and in some cases retirement accounts. They therefore understand that debt repayment is different from saving and from spending. We also do not think that our results are affected by misreporting due to social desirability bias Bursztyn et al. (2025). The SCE modules with the

hypothetical questions do not include any questions about the respondents’ household balance sheets, so there is no explicit reason to conform with socially acceptable behavior. In additional analysis, we find that households who cited government stimulus payments as a reason for having paid down credit card debt in June 2020 were also those with significantly higher reported MPRDs. These findings complement a growing literature, referenced earlier in the paper, that documents that reported behavior aligns with actual behavior.

3 Explaining the empirical facts

In this section, we present a simple yet natural twist on a standard consumption-savings model to illustrate one intuitive mechanism that can generate the facts from the previous section. In short, borrowing constraints and income uncertainty typically work against the empirical evidence shown in Facts 1-3, making the consumption function more concave and giving low-asset (more constrained) households a higher MPC than high-asset (unconstrained) ones. We show that a simple and empirically relevant alteration to this model via a non-constant price schedule of debt can overturn these channels and make the model consistent with the empirical findings. We then use the framework to explore the implications for fiscal policy.

To begin, we fix ideas using an illustrative two-period model without income risk to clarify when the consumption function becomes convex. Then, we present a richer model that we use for our quantitative exploration and counterfactuals. In both, we assume there is a single, non-defaultable asset so that paying down debt and saving are not distinct actions.¹⁶ Therefore, the main empirical counterpart to the model is the MPA, rather than the MPRD and MPS separately. However, our empirical facts have already been shown to carry over to the MPA, and are driven mainly by the MPRD.

3.1 An illustrative two-period model

Households live two periods, $t = 1, 2$, and use a single asset, a_2 , to move consumption between periods. We assume that households choose a terminal asset level a_3 equal to zero and are born with $a_1 = 0$ assets. Income is exogenous and equal to y_1 in the first period and y_2 in the second. We abstract from income risk to cleanly highlight our main mechanism, but reintroduce income risk in the full quantitative model in the next section, to discuss how our channel co-exists with precautionary saving behavior studied by Carroll and Kimball (1996) and many others. Households derive utility from consumption, with a utility function given by $u(\cdot)$, with a positive first derivative and a negative second derivative. They discount the future at a factor β and face a debt price schedule $q(a_2)$, which we assume to be exogenous and continuously differentiable.

¹⁶Hence, given that our empirical results hold for the MPA as well, our model abstracts from the so-called credit card debt puzzle. Gorbachev and Luengo-Prado (2019) provides a recent review of the topic.

Specifically, $q(a_2) = \frac{1}{1+r}$ for all $a_2 \geq 0$, while for $a_2 < 0$ we will have $q(a_2) \leq \frac{1}{1+r}$ and $\frac{\partial q(a_2)}{\partial a_2} \geq 0$. When the latter inequalities are strict, the pricing function represents a borrowing premium that changes with debt as in any model with endogenous default risk.

The problem the household solves is then:

$$\max_{a_2} u(y_1 - q(a_2)a_2) + \beta u(y_2 + a_2) \quad (1)$$

where we have made use of the first period budget constraint $c_1 + q(a_2)a_2 = y_1$ and second period consumption is just available resources $y_2 + a_2$.

Constant Prices. When interest rates are constant, $q(a_2) = q = \frac{1}{1+r}$ everywhere, and therefore $\frac{\partial q(a_2)}{\partial a_2} = 0$. The optimal choice of consumption and savings is determined by the standard Euler Equation, which boils down to:


$$\frac{\partial u(c_1)}{\partial c_1} = \beta \frac{\frac{\partial u(c_2)}{\partial c_2}}{q(a_2)} = \beta (1+r) \frac{\partial u(c_2)}{\partial c_2}$$

in which the household is indifferent between consuming another unit today - valued at the marginal utility of consumption today - and buying a bond at price $\frac{1}{1+r}$ to consume another unit tomorrow, valued at the marginal utility of consumption tomorrow (appropriately discounted). In this risk-free framework, it is a standard result that a_2 will be linear in cash-on-hand y_1 (see Friedman (1957) and Appendix A). Therefore, $\frac{\partial q(a_2)}{\partial a_2} = 0$ implies that both $\frac{\partial a_2}{\partial y_1}$ and MPCs ($\frac{\partial c_1}{\partial y_1}$) are constant in cash-on-hand.

The left-hand panel of Figure 3 depicts this in $\{c_1, c_2\}$ space graphically for some endowments y_1, y_2 and prices q . The intertemporal budget constraint is represented by a straight line with slope $-(1+r)$. Starting from some initial y_1 , the optimal consumption bundle is the point of tangency between the indifference curve and the budget set.

As we increase y_1 to $\bar{y}_1 = y_1 + \tau$, as depicted in the right-hand panel, the horizontal intercept shifts by τ , while the vertical intercept shifts by $\tau(1+r)$. Importantly, the slope of the budget line is unchanged. The income effect implies that consumption in both periods rises, while there is no substitution effect because the interest rate is unchanged. Since consumption goes up in the second period, this implies that savings increase, so that some of the increase in income is used for increased savings/ less borrowing. When preferences are HARA, the MPC is independent of the initial cash-on-hand and is equal to $\frac{1}{1+\beta}$.

Figure 3: Optimal Consumption in Constant q Model



Notes. The solid blue line in the left panel depicts the intertemporal budget constraint, while the curved black line depicts the indifference curve. In the right panel, an increase in y_1 shifts the budget constraint and the indifference curve right from the dash-dotted lines to the solid.

Non-Constant Prices. Allowing for non-constant prices can change how the MPC varies with y_1 . When $\frac{\partial q(a_2)}{\partial a_2} > 0$, the Euler equation becomes:


$$\frac{\partial u(c_1)}{\partial c_1} = \beta \frac{\frac{\partial u(c_2)}{\partial c_2}}{q(a_2) + \frac{\partial q(a_2)}{\partial a_2} a_2} \quad (2)$$

Now, the household is indifferent between consuming another unit today, again valued at the marginal utility of consumption today, and buying a bond at price $q(a_2) + \frac{\partial q(a_2)}{\partial a_2} a_2$ to consume another unit tomorrow, valued at the marginal utility of consumption tomorrow (appropriately discounted). The key difference is that the price of the bond changes with the level of borrowing, and the new effective price takes that change into account.

Allowing the bond price to change with the level of borrowing changes both the optimal level of consumption and borrowing, but also its shape over y_1 space. Specifically, it can make c_1 *convex* in y_1 , and a_2 *concave* in y_1 .¹⁷ As depicted in Figure 4, as we increase y_1 to some $\bar{y} = y_1 + \tau$, two things happen. First, this increases the range of c_1 for which savings are positive (where the slope of the budget constraint is $-(1+r)$). Second, the slope of the budget set where there is borrowing becomes more vertical, indicating it is relatively cheaper to consume tomorrow relative to today. This substitution channel is absent in the constant q environment, so for a given income level, the MPC will be larger in the non-constant $q(\cdot)$ model relative to the constant q model; moreover, it will increase in cash on hand, as that substitution effect increases at higher levels of y_1 . The following proposition states this result formally for the specific bond

¹⁷Notice this is the opposite effect of income uncertainty, which makes the consumption function concave, as shown in Carroll and Kimball (1996).

Figure 4: Optimal Consumption in Non-Constant $q(\cdot)$ Model



Notes. The blue dotted lines show the constant q budget constraints of Figure 3. The dash-dotted blue curve shows the budget constraint for non-constant $q(\cdot)$, which shifts rightward and rotates (solid blue) when y_1 increases. Black lines are indifference curves. Red dots indicate tangency points between indifference curves and budget constraints.

price function that we use in our quantitative analysis.


Proposition 1 *Under log utility, without income uncertainty, and assuming $q(a_2) = \frac{1}{1+r}$ for $a_2 \geq 0$ and $q(a_2) = \frac{1}{1+r} - \phi_1(-a_2)^{\phi_2}$ for $a_2 < 0$, where $\phi_1 > 0$ and $0 < \phi_2 \leq 1$, the $MPC = \frac{\partial c_1}{\partial y_1}$ is increasing in cash-on-hand y_1 .*

Proof. See Appendix A. ■

Non-Linearity in the Price Function The second derivative of $q(\cdot)$ can have an important quantitative effect on the shape of the MPC across y_1 space. Specifically, the change in the steepness of the budget set in response to changes in y_1 depends on the level of income if $q(\cdot)$ is non-linear. Differentiating the budget constraint, we can see that MPCs depend on two terms: how $q(\cdot)$ changes with a_2 (i.e., a price effect) and how a_2 changes with income (i.e., a choice effect).

Ceteris paribus, the MPC function is higher with a stronger price effect ($\frac{\partial q(a_2)}{\partial a_2} > 0$); this is also evident when looking at the Euler equation, where the effective interest rate becomes $(q(a_2) + \frac{\partial q(a_2)}{\partial a_2} a_2)^{-1}$. More importantly, MPCs can be upward sloping in cash on hand, depending on the combined “price” and “choice” effects. In Figure 5, we show that the savings function,

Figure 5: Household choices with different $q(\cdot)$ functions



the a_2 choice, goes from linear to concave when $\frac{\partial q(a_2)}{\partial a_2} > 0$, and thus the consumption function becomes increasingly convex. This is already true if $q(\cdot)$ is linear. With a convex $q(\cdot)$, this effect is even stronger and, as a result, the MPC increases by even more.

Borrowing wedges Recent versions of standard incomplete market models often incorporate a wedge between the interest rate faced by savers and borrowers (see, for instance, Kaplan et al. (2018)). With deterministic income as in the simple model of this section, this feature would not affect MPCs except locally around zero assets. However, it can generate MPCs that are increasing in net assets for borrowers over a larger portion of the negative asset space, as proven by Achdou et al. (2022) (see their Appendix G.3) with income risk. In that setup, income uncertainty makes the wedge work similarly to a soft borrowing constraint, as agents place a non-zero probability on being at the asset kink in the future. Our debt price function generalizes that setup, which allows the model to flexibly fit the empirical pattern of debt and MPCs regardless of the stochastic process for income. Indeed, the strong convexity of our $q(\cdot)$ function resembles a borrowing wedge near zero assets, as it also creates a discontinuity in the derivative of the interest rate.

Endogenous default To generate a non-constant debt price, models with endogenous default would vary $q(\cdot)$ with the borrowers' probability of default. In this very large literature, households taking on more debt are less likely to pay it back and therefore are given a higher interest rate from their lenders, making the price non-constant because it incorporates a risk premium.

To bring this idea into our two-period, deterministic model, consider the second period income to be $y_2 + (1 - d_2)a_2$, where d_2 is the amount of debt the household does not pay back. In equilibrium, this will be directly tied to the price of debt $q(a_2) = \frac{1}{1+r}(1 - d_2)$. In Appendix D, we ensure an internal solution to d_2 by introducing a linearly separable and convex utility cost of default. We obtain a slightly different Euler equation than Equation 2:

$$\frac{\partial u(c_1)}{\partial c_1} (1 + \varepsilon_q(a_2)) = \beta(1 + r) \frac{\partial u(c_1)}{\partial c_1}$$

where $\varepsilon_q(a_2)$ is now the elasticity of q with respect to a_2 , $q'(a_2) \frac{a_2}{q(a_2)}$. Still, this gives us all of the same qualitative conclusions for the convexity of the consumption function and the slopes of the MPC with cash on hand. This is shown more completely in Appendix D.2.

3.2 Full Quantitative Model

In this section, we extend the intuition outlined with the two-period model to a standard, infinite-horizon, incomplete-markets model of consumption and savings. We will use this model to assess the quantitative plausibility of our mechanism and to study fiscal stimulus policy.

The economy is populated by a continuum of infinitely lived households indexed by their net asset holdings a and their exogenous income y .¹⁸ Households can borrow at a price schedule $q(\cdot)$, whose features we describe below.

Our characterization of the evolution of households' income follows Krueger et al. (2016) and, more generally, a large empirical literature in labor economics. Log labor income y follows:

$$y = z + \epsilon \tag{3}$$

where z is the persistent component of income and ϵ is the transitory component of income. The persistent component of log income follows $z' = \rho z + \eta$. Innovations to z , denoted with η , as well as transitory shocks ϵ , have mean zero and are normally distributed with variances σ_η^2 and σ_ϵ^2 , respectively. We denote by $F(\cdot)$ the CDF of ϵ and $\pi(z'|z)$ the conditional probability of z' given z . η and ϵ are orthogonal to each other and independently distributed over time and across households. Finally, households can receive a lump-sum transfer τ from the government.

Households begin the period with their net asset position a and income y . Households can have a negative net asset position $a < 0$, which we assume to be bounded below by the natural debt limit a_{ndl} . Every period, households choose consumption of a nondurable good c , from which they derive utility $u(c)$, and next period assets a' . New assets are purchased at a price $q(\cdot)$. As we discussed previously, in the standard incomplete markets model $q = \frac{1}{1+r}$ for all households, where r is the risk-free rate. We nest this case with a generic formulation for the

¹⁸We omit time subscripts and denote next period variables with the superscript $'$.

asset price:

$$q = \begin{cases} \max \left[\frac{1}{1+r} - \phi_1 (-a')^{\phi_2}, 0 \right] & \text{if } a' \leq 0 \\ \frac{1}{1+r} & \text{if } a' > 0 \end{cases} \quad (4)$$

When $\phi_1 > 0$, larger borrowed amounts are associated with worse prices (higher interest rates). The household problem can be summarized as follows:

$$\begin{aligned} V(a, z, \epsilon) &= \max_{c > 0, a' > a_{ndl}} u(c) + \beta E_{z', \epsilon'} V(a', z', \epsilon') \\ \text{subject to :} \\ c + q(a')a' - a &= e^{z+\epsilon} \\ z' &= \rho z + \eta \end{aligned}$$

Household decisions To understand the role played by the price schedule $q(\cdot)$, and how it alters the intertemporal trade-off faced by households, we look at the generalized Euler Equation in the infinite horizon model:

$$\frac{\partial u(c)}{\partial c} \left\{ \frac{\partial q(a')}{\partial a'} a' + q(a') \right\} = \beta E_t \frac{\partial u(c')}{\partial c'} \quad (5)$$

The intuition is the same as in the two-period model. When prices are constant ($\frac{\partial q(\cdot)}{\partial a'} = 0$), 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 an interest rate schedule of debt, this condition is affected in two ways. First, borrowing households will have fewer available units to consume as $q(\cdot)$ declines. Second, savings decisions (i.e., a') affect the pricing schedule and therefore change the amount of available resources.

Put differently, the Euler Equation now includes an additional term, $\frac{\partial u(c)}{\partial c} \frac{\partial q(a')}{\partial a'} a'$, that was not present in the constant interest rate framework. In levels, this term acts like a higher interest rate, but it also changes at different levels of c and a' . Since both of these are monotone in cash-on-hand, as resources increase, marginal utility declines just as a' is also approaching zero; thus, the additional term is smaller. Again, this makes our consumption function convex, with a deviation from the constant $q(\cdot)$ consumption function that narrows with cash-on-hand. Absent our channel, the model would instead deliver a concave consumption function, as in the seminal work by Zeldes (1989) and Carroll and Kimball (1996).

A natural consequence is that a non-constant price schedule of debt will alter the distribution of marginal propensities to consume and, in particular, how they correlate with net asset

positions. Households with a high marginal utility of consumption can face lower prices $q(\cdot)$ and positive $\frac{\partial q(a')}{\partial a'}$. For borrowers (ie, $a' < 0$), both margins increase the household's effective discount factor. As such, they decrease the motive to bring forward consumption. Other things equal, this reduces their incentive to immediately spend transitory transfers (i.e., it lowers their instantaneous marginal propensity to consume).¹⁹

Responses to transitory income shocks We now describe households' responses to a transitory income shock like a tax rebate, τ . In what follows, we consider these innovations τ as local perturbations of income. We differentiate the budget constraint with respect to the τ component of income:

$$\frac{\partial c}{\partial \tau} + \frac{\partial q(a')}{\partial a'} \frac{\partial a'}{\partial \tau} a' + q(a') \frac{\partial a'}{\partial \tau} = 1 \quad (6)$$

The well-known marginal propensity to consume (MPC) is defined as usual:

$$\text{MPC} := \frac{\partial c}{\partial \tau}$$

while the marginal propensity to adjust the net asset position is:

$$\text{MPA} := \frac{\partial q(a')}{\partial a'} \frac{\partial a'}{\partial \tau} a' + q(a') \frac{\partial a'}{\partial \tau}$$

As in the data, the MPC and the MPA must sum to 1 at the household level.

4 Quantitative Evaluation

4.1 Calibration

One model period is one quarter. Table 2 presents the details of the calibration. The parameters governing the income process come from Krueger et al. (2016). Utility from consumption follows a CRRA specification, with risk aversion parameter γ equal to 1 (log utility). The quarterly risk-free interest rate, r , is set to match a 3 percent annual rate.

The three remaining parameters, β , ϕ_1 , and ϕ_2 are chosen by targeting three moments from the SCE: (i) the share of households with negative net liquid assets, (ii) the MPC of households in the bottom quintile of net liquid wealth-to-income, $\frac{a}{y}$, conditional on negative net liquid assets, and (iii) the MPC in the top quintile, also for negative asset holders.²⁰ All of these are

¹⁹Another way to think about this effect is that the non-constant interest rate schedule gives low asset households an additional marginal value of increasing assets, akin to the description of “saving-constrained” households in Miranda-Pinto et al. (2020).

²⁰As in the data, we scale assets by income to remove permanent (or persistent) differences in income that scale both asset and consumption choices.

Table 2: Calibration

Pre-defined parameters		
r	0.0074	
γ	1	
ρ	0.9923	
σ_η	0.0985	
σ_ϵ	0.2285	
Internally calibrated parameters		
β	0.97	
ϕ_1	0.05	
ϕ_2	0.10	
Targeted moments	Data	Model
Share of Households with $a < 0$	0.386	0.378
MPC (bottom quintile of $\frac{a}{y}$, $a < 0$)	0.195	0.189
MPC (top quintile of $\frac{a}{y}$, $a < 0$)	0.266	0.254

taken from our SCE data sample used in Section 2. The first target is somewhat larger than the share of households with negative liquid wealth in Kaplan, Moll and Violante (2018). There are a range of estimates in the literature, partly due to an ongoing debate on whether to use net worth or gross unsecured debt, and whether to use liquid or total assets as the calibration target.²¹ Our liquid wealth calibration implies a relatively low value of β , 0.97 quarterly, because the relevant concept of wealth for our model excludes many illiquid assets and secured debts.

The remaining two targets ensure that our model matches our second and third empirical facts documented in Section 2: the MPC is higher and MPA lower for those with less negative net wealth-to-income ratios. The model also does well along other, non-targeted dimensions. For example, the MPC of the median negative asset holder is 0.25 in the model and 0.23 in the data, whereas the average MPC of debtors is 0.25 in the model and 0.24 in the data.

In Appendix D.3 we incorporate endogenous default into our model, tying $q(\cdot)$ to an equilibrium condition. There, we show that we can match the same moments by choosing parameters of the utility cost of default, a power function, $\kappa_1 d^{\kappa_2}$. In that case, instead of two parameters, ϕ_1 and ϕ_2 , directly dictating the q function, κ_1 and κ_2 dictate default rates, which imply a similarly convex $q(\cdot)$ and lead to all of the same quantitative results.

4.1.1 Calibration-implied $q(\cdot)$ function

Given the calibration, we consider whether the debt price schedule delivered by the model is empirically plausible. This is a somewhat difficult question to answer directly because there is little empirical evidence on the relationship between $q(\cdot)$, debt, and consumption behavior.²²

²¹For instance, 39% of US households in the 2016 SCF carried a balance on their credit card, as reported by Exler and Tertilt (2020).

²²An important exception is Kreiner et al. (2019), who estimate marginal interest rates in Denmark and show

We proceed in two steps. First, we present novel empirical evidence on debt-price schedules that supports our main model mechanism. Second, we show that our calibration is also consistent with existing evidence on default rates when we assume that the exogenous $q(\cdot)$ function arises from lenders’ pricing default risk.

Novel evidence from the Credit Consumer Panel We turn to credit card data to empirically investigate the nature of the $q(\cdot)$ function, which is central to our theory. We use administrative data from the New York Fed’s Consumer Credit Panel (CCP),²³ and a methodology proposed by Guttman-Kenney and Shahidinejad (2024), to construct account-level financing charges that credit card holders face each quarter between 2017 and 2024. The dataset is ideal for our purposes since we can observe for a single individual over time multiple credit card accounts and the dynamics of payments. We combine this with a method to estimate effective interest rates to construct a $q(\cdot)$ function. Appendix B.6 contains detailed information on the dataset and the empirical strategy. The approach proposed by Guttman-Kenney and Shahidinejad (2024) relies on the fact that credit card lenders use the following deterministic formula to calculate minimum payments m :


$$m = \max\{\mu, \theta b + (1 - \theta)f\} \quad (7)$$

In the data, we observe minimum payments and the statement balance, b , directly in the CCP, for each account at a quarterly frequency. We follow Guttman-Kenney and Shahidinejad (2024) in addition to our own empirical investigation and set μ , the floor dollar amount determined by the creditor, and θ , a percentage determined by the creditor, to values that encompass typical credit card agreements in the US, as discussed in the Appendix. With that in hand, we are then able to estimate financing charges, f , for each credit card held by account holders in the sample, at a quarterly frequency. We also construct a proxy for the effective interest rate faced at each account and time period, $r = \frac{f}{b-f}$, which is the model counterpart of the inverse of our price $q(\cdot)$ function. We use two alternative measures for the effective interest rate, both of which aggregate r at the individual level, for each quarter. For the “average interest rate,” we average r across accounts, at a point in time, for each individual; for the “marginal interest rate,” we take the maximum interest rate across accounts, at a point in time, for each individual. Balances are always the sum across accounts for each individual at each time period. As a validity check, we find that our estimated interest rates fall with credit scores, consistent with recent findings by Drechsler et al. (2025) using Y-14 account-level data.

that they fall with liquid assets to income and increase with the amount of debt service—consistent with our model. But unlike our findings, they see a positive correlation between reported MPCs and marginal interest rates. This could be because of differences in institutional setups, bunching of their reported MPCs, and the type of fiscal policy studied.

²³Note that the CCP does not include any race or ethnicity information.

Figure 6: Interest rates by debt balances




Notes. Source: NY Fed CCP / Equifax. Data is at the individual level, between 2017-q4 and 2024-q4. For each quarter, account-level interest rates are aggregated to the individual level as an average (left panel) and as the maximum (right panel) across the accounts held by an individual in that quarter. Total balance is the sum of credit card balances for a given individual each quarter. The figure is a binned scatterplot: we group observations in 20 quantiles of total balance (horizontal axis) and plot the average interest rate at each quantile (vertical axis). For additional details, see Appendix B.6.

Three sets of findings provide direct support for our mechanism. First, pooling all accounts and time periods, we find that (both average and marginal) interest rates are increasing and concave in credit card balances as shown in Figure 6, thus directly giving empirical support to the calibrated $q(\cdot)$ we use in our model.

Our second and third findings exploit the panel dimension of the CCP data. Since the $q(\cdot)$ function is the debt price function faced by an individual, this variation is preferred over cross-sectional variation, which may be confounded by unobserved factors that correlate with debt balances and make certain households face higher interest rates always. To this end, we turn to how quarterly changes in individuals' credit card balances are associated with changes in interest rates at a quarterly frequency. First, we find that decreases in balances (i.e., debt repayments) are associated with declines in (both average and marginal) interest rates and present these findings in Appendix Table B.8. Specifically, a one standard deviation increase in debt is associated with an increase in average (marginal) interest rates of 1.5 (2.7) percentage points. In the model, the same regression delivers a coefficient of 2.2. We then explore whether the sensitivity of interest rates to changes in debt balances is smaller at larger debt balances, as would be consistent with our convex $q(\cdot)$.²⁴ Figure 7 depicts the estimated coefficients from a regression of changes in effective interest rates on changes in individual credit card balances at different quintiles of debt balances. We find that the sensitivity for both average (panel (a)) and marginal (panel (b)) interest rates is smaller for larger debt balances. In other words, the

²⁴So long as households face interest rate schedules that are the same up to an intercept term (which is differenced out), the interaction with debt quintiles will deliver unbiased estimates for sensitivities in different quintiles.

Figure 7: Interest rate to debt sensitivities by debt quintile



Notes. Source: NY Fed CCP / Equifax. Data is at the individual level, between 2017-q4 and 2024-q4. Interest rates and debt balances defined as in the text and in Figure 6. Debt is standardized dividing it by the overall standard deviation of all debt balances. We group observations by 5 quantiles of debt balances (horizontal axis). For each quintile, we regress within-individual changes in the interest rate on debt changes and a constant, and report the regression coefficient and 95% confidence bands. For additional details, see Appendix B.6.

convexity of the $q(\cdot)$ function also holds using individual-level variation. The sensitivity for the third quintile is consistent with the sensitivity implied by our calibrated model, although we underestimate the dispersion of the regression coefficients.

As we show in Appendix Table B.9, there is a simple intuition for the within-individual sensitivity of interest rates to debt repayments. We estimate that, as credit card balances increase, the number of credit card accounts associated with positive financing charges rises. This suggests that, as debt rises, individuals end up facing higher effective average/marginal rates, because they have to carry balances at higher rate cards and/or new cards come with higher rates.

Existing evidence Households in our model face empirically plausible interest rates. The mean annualized interest rate faced by households with debt in our model is 20.6%, varying from risk-free 3% to 27.9%. The median interest rate, 21.3%, is broadly in line with estimates by Galenianos and Gavazza (2022). Our calibration delivers a dispersion of interest rates slightly lower than what estimated by Galenianos and Gavazza (2022), with ten percent of households facing an interest rate of 24% or more.

We can also construct actuarially fair default probabilities, $q(a) = \frac{1}{1+r}(1 - Pr[\text{default}](a))$, assuming that the $q(\cdot)$ function arises from lenders' pricing default risk. Our calibrated $q(\cdot)$ function has implied quarterly default probabilities as high as 5.3% in the negative asset region. We have an average default probability among negative asset holders of 3.9%, or 1.5% overall. Once annualized, this is higher than the bankruptcy rates considered by Chatterjee et al. (2007) and Athreya et al. (2018), but lower than delinquency rates. Our maximum annualized implied

default probability (19.6%) is basically equivalent to the upper bound considered by Dempsey and Ionescu (2021) in their empirical investigation. We explore this framework more directly in Appendix D by adding endogenous default to our current model. There, we are able to replicate our quantitative results by choosing the default cost so as to generate a convex endogenous $q(\cdot)$ rather than imposing a convex exogenous $q(\cdot)$ directly.

Finally, as noted earlier, recent papers (e.g., Kaplan et al. (2018)) also feature borrowing “wedges” that make MPCs increase in net assets for borrowers in at least part of the asset space. The wedge is a constant increase in the interest rate for borrowers relative to savers and creates a local convexity in the consumption function near zero net assets. Appendix C.3 provides details for how we consider this form of $q(\cdot)$ -menu. In our calibrated model, a wedge of their size (a 6 percentage point wedge in the annual interest rate) generates MPCs that are slightly decreasing in net assets among borrowers and do not quantitatively fit our empirical estimates. If we recalibrate the model targeting the same empirical moments of Table 2 with a wedge, a better fit can be achieved with a wedge of 23 percentage points. In sum, while other mechanisms can qualitatively generate MPC patterns like those we highlight, we see our $q(\cdot)$ function as a usefully flexible generalization of empirically plausible debt price schedules.

4.2 Fiscal stimulus through insurance

In this section, we show how accounting for the documented empirical facts through a non-constant debt price schedule can alter the expected effects of fiscal policy. We perform four exercises. First, we show that, when interest rates are debt-sensitive, stimulus and insurance motives do not align well across households as they typically would in a model with a constant (flat) debt-price schedule. In fact, they are negatively correlated, with indebted households experiencing large welfare gains and displaying relatively small instantaneous consumption responses.

This tradeoff materializes also when comparing short-run and long-run fiscal multipliers: debt-sensitive interest rates lower upon-impact consumption responses of the poorest households, while making them more persistent over time. Conversely, they increase MPCs of households with little debt, but steepen their dynamic profile. In our second exercise, we show how these two effects aggregate up in the economy and unfold over time. The presence of a non-constant debt-price schedule amplifies the consumption effects of fiscal policy.

Third, we discuss optimal policy with debt-sensitive interest rates. Even when targeting only based on income, the allocation of transfers is different from what is implied by the canonical constant-interest rate model, and depends on whether the planner wishes to maximize aggregate (upon-impact) spending or aggregate welfare. We then consider an alternative policy that instead forgives debt service. This policy enhances welfare gains by targeting more appropriately, but at the expense of a smaller short-run boost to aggregate consumption following the logic

previously described.

Finally, we compute the average welfare and consumption gains stemming from the Economic Impact Payments (EIP) by allocating the payments in the model according to their allocation in the EIP program. We find that both the aggregate welfare and upon-impact consumption gains are larger in our model than in the alternative with constant interest rates. As in Kaplan and Violante (2014), our analysis is in partial equilibrium, and thus interest rates are exogenous. This is consistent with our liquid wealth calibration, which excludes much of the US capital stock, but focuses on the assets most relevant to the interest rate schedule we infer, and which are most likely to be affected by households' MPRDs. Hence, the aggregate response reflects heterogeneity in household responses towards the bottom of the distribution.²⁵

4.2.1 Stimulus vs insurance across households

To compute welfare in the model, we solve for a consumption variation function $\lambda(\cdot)$ that renders agents indifferent between receiving the rebate and not receiving the rebate τ . This amounts to solving for the $\lambda(a, \varepsilon, z)$ such that:

$$V(a + \tau, \varepsilon, z) = \sum_{t=0}^T \beta^t u(c_t^\tau) = \sum_{t=0}^T \beta^t u((1 + \lambda(a, \varepsilon, z))c_t)$$

where $\{c_t^\tau\}_{t=0}^T$ is the sequence of optimal consumption choices starting from date $t = 0$ and asset level $a + \tau$ and $\{c_t\}_{t=0}^T$ is the sequence of consumption choices starting from a . With log utility this becomes fairly straightforward to compute:

$$\lambda(a, \varepsilon, z) = \exp\{(1 - \beta)(V(a + \tau, \varepsilon, z) - (V(a, \varepsilon, z)))\} - 1$$

We show this graphically in Figure 8, depicting the $\lambda(\cdot)$ (multiplied by 100) in blue. In addition, we also plot the upon-impact spending responses to the rebate in red. Given the focus of our paper, we show only negative asset holders. Throughout this section, we consider a uniform, transitory, and unexpected lump-sum transfer equal to 10% of average quarterly income.²⁶ We bin households by 10 quantiles of their asset-to-income ratios, and, for each quantile, we compute the average welfare gain and spending propensity.


Figure 8 clearly illustrates that welfare and instantaneous spending gains are starkly negatively related.²⁷ As debt decreases, households spend a larger fraction of the rebate today, but

²⁵Our setup also abstracts from how stimulus checks are financed by the government and from the credit sector to whom debt service would be paid, although we discuss the latter in the model with endogenous default of Appendix D.

²⁶This is approximately equivalent to \$1200, which was the payment size received by an individual with no dependent children.

²⁷In Appendix D.3 we show that the same result holds in our model with q generated in equilibrium with endogenous default.

Figure 8: Stimulus vs insurance across households



Notes. We bin the stationary distribution of households' assets-to-income ratios, conditional on $a < 0$, by 10 deciles of equal mass. For each decile, we plot the average welfare gain due to the transfer, in blue, and the dollar-for-dollar spending effect of the transfer upon impact, in red. Portions of the line where there are more points, i.e. deciles are closer together, imply there is more mass.

also have marginally decreasing welfare gains. Binning more finely, across 50 quantiles of negative asset-income ratios, the correlation between welfare gains and short-run spending effects is -0.88 . This underscores how debt-sensitive interest rates create a clear trade-off between stimulus and insurance. If instead we consider an alternative model in which households face a constant interest rate, spending and welfare gains co-move, with a correlation across quantiles of $+0.97$.²⁸ Therefore, debt-sensitive interest rates imply that policymakers may wish to design fiscal policies differently depending on their objective. This consideration also underscores a trade-off between short- and long-run fiscal multipliers, which we investigate next.

4.2.2 The aggregate effects of fiscal policy

While debt-sensitive interest rates make MPCs increase in net assets for debtors, they also alter the dynamics of spending responses. Figure 9 plots intertemporal MPCs – i.e., the dollar-for-dollar consumption response over time to a transitory cash transfer in the first period. In the left panel, we average across households in the bottom first percentile of the asset distribution when the transfer is paid out, while in the right panel we look at the 10% of households with the smallest debt balances.


²⁸For this environment, we assume that households face an exogenous borrowing limit equivalent to the one generated in our baseline model. The equivalent of Figure 8 for this model can be found in Appendix C.1.

We compare our calibrated model with two alternative constant q (CQ hereafter) models. In both, households exogenously face the same borrowing limit, but interest rates do not vary. We consider two different recalibrations of this model, changing β to target different moments (since both models are with a constant interest rate, the only parameter to lever is β). The “industry standard” is to calibrate β to moments of the asset distribution (see Kaplan and Violante (2022)). This is exactly what we do for the model in dashed red (“Constant q - match % in debt”) - we match the share of households with negative net liquid assets in the data. The steady state of this model is thus observationally equivalent to our model in terms of the share of borrowers, but they face a different interest rate schedule. The model represented in dash-dotted black (“Constant q - match average MPC”) instead recalibrates β to match the instantaneous average MPC, which is achieved by allowing for a counterfactually high share of households with negative net liquid wealth, 88%. This second recalibration only serves to show cleanly how our mechanism alters the persistence of aggregate consumption effects, and thus long-run fiscal multipliers, since we force the short-run average consumption response to be the same.

When $q(\cdot)$ is constant, in both recalibrations, the poorest behave as standard liquidity-constrained households, consuming most of the rebate in the very first quarters. In our baseline model, instead, there is an extra incentive to use the transfer to pay down debt. The instantaneous MPC is therefore much lower than in either recalibration of the model with constant interest rates, but these households keep spending more for longer (the blue line is above the others beginning 2 quarters from the transfer date). Indeed, their implied cumulative spending response exceeds the amount of rebate dollar already after three years, thanks to endogenous improvements in the price of debt $q(\cdot)$. The right panel, instead, shows that households with little debt respond more upon impact in our baseline model than in both models with constant interest rates. This is driven by the force shown in Figure 5, in which the very steep interest rate schedule in that region of assets raises the level of the MPC.

The aggregate spending effect of fiscal transfers depends on the combination of these offsetting forces. As shown in Figure 10, although they both feature the same share of households with negative net liquid wealth, an economy with debt-sensitive interest rates has a much higher aggregate MPC than the constant q model depicted in red. Moreover, the baseline economy displays more persistent effects of fiscal policy. To see this, we compare the cumulative aggregate intertemporal MPC in our baseline to the CQ model that starts at the same aggregate MPC in the first quarter (black dashed line). At the beginning of the fourth year, the entire size of the fiscal package has been spent in our baseline model, whereas this happens more than a year later in the CQ model in black, despite the larger share of borrowers. The gap between these two particular models also opens up over time. In our baseline model, improvements in debt positions and the resulting lower interest rates imply that the cumulative aggregate spending

Figure 9: Heterogeneous intertemporal MPCs



Notes. In the Constant q model, $\beta = 0.9745$ in the version plotted with a dash-dotted red line, while the recalibration with $\beta = 0.9875$ is shown with dashed black.

is 8 percentage points higher than the aggregate size of transfers 30 quarters out. In summary, a non-constant debt-price schedule amplifies the consumption effects of fiscal policy relative to an environment with a constant debt-price schedule, both in the short- and in the long- run.²⁹

Auclert et al. (2024) discuss how intertemporal MPCs can be used to discipline heterogeneous agent macro models. Our model generates consumption responses in the first two years that are quantitatively similar to their analysis: in particular, not only a bit more than 50% of the rebate is spent in the first year, but also a fairly large fraction is spent in the following year.³⁰ In addition, our analysis suggests that debt-sensitive prices not only matter for the short-run persistence of intertemporal MPCs, but also for the amplification of fiscal policy several years out. Our mechanism is thus also consistent with empirical evidence by Agarwal et al. (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, stimulating an increase in spending in the medium-run.


4.2.3 Optimal policy and the Economic Impact Payments

In this final section, we consider optimal targeting of fiscal policy, and also measure the effects of the 2020 EIP. Throughout this section, we consider our baseline model and the CQ model in

²⁹In Appendix D.3 we show how this mechanism is present even with endogenous default and thus equilibrium-determined interest rates, albeit quantitatively dampened.

³⁰Our first-year response (63%) is on the upper end of the estimates they report using Italian data and Norwegian estimates by Fagereng et al. (2021), and our second-year response is also slightly higher at 26 cents rather than 16.

Figure 10: Aggregate spending effects of a transfer



Notes. In the constant q model, $\beta = 0.9745$ in the version plotted with a dash-dotted red line, while the recalibration with $\beta = 0.9875$ is shown with dashed black.


which β has been recalibrated to match the share of households with negative net liquid wealth. We set the total aggregate size of the fiscal package to be the same across all exercises and models so that the only choice is how to allocate a fixed amount of fiscal spending.

First, we formally show that, even if the planner is constrained to target transfers based solely on income, she would allocate transfers very differently in the two models and depending on the object she wishes to maximize. For each of the two models, we consider two allocations: the one that maximizes aggregate consumption upon impact, and the one that maximizes aggregate welfare.³¹ Figure 11 shows the corresponding optimal allocations by plotting the mean rebate disbursed in each decile of the income distribution.

Three of the four allocations are similar: the two allocations in the CQ model (corresponding to consumption (blue diamonds) and welfare maximizing (black stars)) and the allocation that maximizes aggregate welfare in the baseline model (green circles) featuring a debt-price schedule. In these three scenarios, it is optimal to give more of the total fiscal transfer to households whose income is below the median. In the baseline model, however, the rebate allocation is very different if the planner instead maximizes the immediate effect on aggregate consumption. In this case (depicted in red squares), lower-income households receive a lower share of the overall rebate, as a larger fraction of higher-income households become rebate recipients. This result demonstrates that the divorce of welfare gains from upon-impact spending depicted in Figure 8 is present even if the planner targets rebates by income, though it is not present in the standard model.

³¹See Appendix C.2 which formalizes both planner problems.

Figure 11: Optimal targeting by income



Notes. Using the optimal allocation of transfers in each of the four exercises, we plot the average rebate (vertical axis) for each decile of income (horizontal axis). In the constant q model, $\beta = 0.9875$ to match the empirical share of households with negative net liquid wealth.

The previous exercises considered how to target rebates when the planner can only base their choices on income. However, the model formalizes that those with the largest debts will benefit the most from transfers, and suggests that larger gain can be had with alternative policies which can more directly target those with larger debts. We therefore next compare the rebate policy above, which in practice can only target households by income, with an alternative policy that allows for debt service forgiveness (which we label debt service targeting). This policy can be thought of as analogous to the mortgage interest deduction, and demonstrates what happens when we explicitly target debt service relief. In both policies, we consider a planner making a binary choice: whether or not to give the household a \$1200 rebate in column (I) or whether to forgive 100% of the debt service in column (II) of Table 3.³² In column (III), we consider transfers that resemble the 2020 EIP episode (see Appendix C.2 for details): we give agents a \$1200 check if they are in the bottom 88% of income and phase out the payment as income increases.

For each of these three policies, the first two rows of Table 3 report the aggregate welfare gains in the baseline and CQ models, respectively. In the first two columns, the planner aims

³²See Appendix C.2 for the full planner's problem we solve. Implementing optimal debt service targeting, in which the planner chooses what share of debt service to forgive, rather than a constant share, is computationally not feasible. Our analysis is also consistent with Kaplan and Violante (2014), who also consider a binary policy for income targeting. This is without loss of generality, since, when we repeat Figure 11 using the solution constrained by a binary policy, the qualitative results are the same, but are naturally more muted since targeting is less precise due to the discrete nature of the planner's decision.

Table 3: Aggregate Effects of Fiscal Policy

	I	II	III
Aggregate welfare effect (%)	Income targeting	Debt service targeting	2020 EIP
Baseline model	22.4	24.7	22.5
Constant q	1.6	2.9	1.6
Aggregate consumption effect (%)	Income targeting	Debt service targeting	2020 EIP
Baseline model	2.0	1.8	2.0
Constant q	0.5	1.4	0.5

Notes. The top two rows of the table report the aggregate welfare gain, that is, the percent increase in aggregate welfare for an equally-sized aggregate transfer. For Columns (I) and (II), the planner maximizes aggregate welfare. The bottom two rows report the aggregate upon-impact consumption gain, that is, the percent increase in consumption for an equally-sized aggregate transfer. For Columns (I) and (II), the planner maximizes aggregate consumption in the first quarter of the policy. In the constant q model, $\beta = 0.9875$ to match the empirical share of households with negative net liquid wealth.

to maximize aggregate welfare. The third and fourth rows report the aggregate consumption gains for the baseline and CQ models. In this case, in Columns (I) and (II), the planner aims to maximize aggregate consumption.

Our model delivers much larger aggregate gains relative to the CQ model, in both the welfare- and consumption-maximizing allocations. Aggregate consumption goes up by 2.0%, four times what we see in the CQ model. The gap is even bigger in terms of aggregate welfare, jumping 22.4% in our model, a 14-fold increase compared to the canonical model. In terms of consumption equivalent units, as computed with $\lambda(\cdot)$ in Section 4.2.1, the average consumption increase each period would be 0.51% in the baseline model and only 0.19% in the constant q model. The strong welfare effects are the result of the debt-repayment mechanism: by using the rebates to repay their debts, households persistently improve their net asset position, face lower interest rates, and sustain persistently higher consumption.

These results are very similar to our estimates of the aggregate welfare and consumption gains of the 2020 EIP policy, as shown in Column (III). Indeed, Economic Impact Payments were phased out in a similar way to what is implied by a binary policy that targets a rebate by income.³³

Finally, in column (II), we show the effects of debt service targeting. Welfare increases even more than under income-targeted transfers in both models, but by relatively more in our baseline model (+2.3pp in our model, +1.3pp in the CQ model over what income-targeting can deliver in terms of welfare). However, aggregate consumption gains under debt-service targeting are relatively lower in our model compared to income targeting (1.8% versus 2.0%, respectively). This once again shows the divorce of welfare gains and on-impact spending gains: this time not

³³Note that the 2020 EIP policy delivers a slight improvement on the binary rebate policy of column (I) because it featured a gradual phase-out. Aggregate welfare would have increased much more if the rebate size was freely chosen for each income level, as in Figure 11.

just in the cross-section, but also in the aggregate. In our model, high-debt-service households, targeted by this policy, display large welfare gains, but not necessarily high MPCs.

5 Conclusions

We provide new empirical evidence for an underappreciated fact: most households use fiscal transfers to pay down debt. We document that these MPRDs are largest among households with more negative net wealth-to-income ratios. Conversely, these households spend relatively little of their rebate checks upon impact.

A standard heterogeneous-agent, consumption-savings model can be consistent with these facts if borrowing rates increase with household debt. We analytically characterize the conditions underlying this result. A full quantitative model replicates the empirical facts with empirically plausible interest rate schedules. The model is used as a laboratory to show how our novel mechanism alters the effects of fiscal stimulus. In particular, a tradeoff arises between immediate spending stimulus and longer-run welfare gains. As such, policymakers can maximize the former or the latter, targeting different sets of households.

References

- Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll, “Income and wealth distribution in macroeconomics: A continuous-time approach,” *The review of economic studies*, 2022, 89 (1), 45–86.
- 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.
- Aguiar, Mark, Mark Bils, and Corina Boar, “Who Are the Hand-to-Mouth?,” *The Review of Economic Studies*, 05 2024, 92 (3), 1293–1340.
- Arellano, Cristina, Xavier Mateos-Planas, and José-Víctor Ríos-Rull, “Partial Default,” *Journal of Political Economy*, 2023, 131 (6), 1385–1439.
- Armantier, Olivier, Leo Goldman, Gizem Koşar, and Wilbert Van der Klaauw, “An update on how households are using stimulus checks,” Technical Report, Liberty Street Economics 2021.
- , –, –, Jessica Lu, Rachel Pomerantz, and Wilbert Van der Klaauw, “How have households used their stimulus payments and how would they spend the next?,” Technical Report, Liberty Street Economics 2020.
- 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.
- Auclert, Adrien, Matthew Rognlie, and Ludwig Straub, “The Intertemporal Keynesian Cross,” *Journal of Political Economy*, 2024, 132 (12), 4068–4121.
- Baker, Scott R., Robert A Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine Yannelis, “Income, Liquidity, and the Consumption Response to the 2020 Economic Stimulus Payments*,” *Review of Finance*, 03 2023, 27 (6), 2271–2304.
- Berger, David W, Konstantin Milbradt, Fabrice Tourre, and Joseph S Vavra, “Refinancing frictions, mortgage pricing and redistribution,” Working Paper, National Bureau of Economic Research 2024.
- Boutros, Michael, “Evaluating the Impact of Economic Impact Payments,” Technical Report December 2020.

- Bursztyn, Leonardo, Ingar K Haaland, Nicolas Röver, and Christopher Roth**, “The Social Desirability Atlas,” Working Paper, National Bureau of Economic Research 2025.
- Carroll, Christopher D and Miles S Kimball**, “On the concavity of the consumption function,” *Econometrica: Journal of the Econometric Society*, 1996, pp. 981–992.
- Carroll, Christopher, Jiri Slacalek, Kiichi Tokuoka, and Matthew N White**, “The distribution of wealth and the marginal propensity to consume,” *Quantitative Economics*, 2017, 8 (3), 977–1020.
- 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.
- Christelis, Dimitris, Dimitris Georgarakos, Tullio Jappelli, Luigi Pistaferri, and Maarten Van Rooij**, “Asymmetric consumption effects of transitory income shocks,” *The Economic Journal*, 2019, 129 (622), 2322–2341.
- Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber**, “How did US consumers use their stimulus payments?,” Working Paper, National Bureau of Economic Research 2020.
- Colarieti, Roberto, Pierfrancesco Mei, and Stefanie Stantcheva**, “The How and Why of Household Reactions to Income Shocks,” NBER Working Papers 32191, National Bureau of Economic Research, Inc Mar 2024.
- Commault, Jeanne**, “Does Consumption Respond to Transitory Shocks? Reconciling Natural Experiments and Semistructural Methods,” *American Economic Journal: Macroeconomics*, April 2022, 14 (2), 96–122.
- Crawley, Edmund and Andreas Kuchler**, “Consumption heterogeneity: Micro drivers and macro implications,” *American Economic Journal: Macroeconomics*, 2023, 15 (1), 314–41.
- Crossley, Thomas F, Paul Fisher, Peter Levell, and Hamish Low**, “MPCs in an economic crisis: Spending, saving and private transfers,” *Journal of Public Economics Plus*, 2021, 2, 100005.
- Dempsey, Kyle and Felicia Ionescu**, “Lending Standards and Borrowing Premia in Unsecured Credit Markets,” Finance and Economics Discussion Series 2021-039, Board of Governors of the Federal Reserve System (U.S.) June 2021.
- Drechsler, Itamar, Hyeyoon Jung, Weiyu Peng, Dominik Supera, and Guanyu Zhou**, “Credit card banking,” Technical Report, Staff Reports 2025.

- 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.
- Friedman, Milton**, “Theory of the consumption function,” in “Theory of the Consumption Function,” Princeton university press, 1957.
- Fuster, Andreas, Greg Kaplan, and Basit Zafar**, “What Would You Do with \$500? Spending Responses to Gains, Losses, News, and Loans,” *The Review of Economic Studies*, 11 2021, *88* (4), 1760–1795.
- Galenianos, Manolis and Alessandro Gavazza**, “Regulatory interventions in consumer financial markets: The case of credit cards,” *Journal of the European Economic Association*, 2022, *20* (5), 1897–1932.
- Ganong, Peter, Damon Jones, Pascal J Noel, Fiona E Greig, Diana Farrell, and Chris Wheat**, “Wealth, race, and consumption smoothing of typical income shocks,” Working Paper, National Bureau of Economic Research 2020.
- Gorbachev, Olga and María José Luengo-Prado**, “The credit card debt puzzle: The role of preferences, credit access risk, and financial literacy,” *Review of Economics and Statistics*, 2019, *101* (2), 294–309.
- Graham, James and Robert A McDowall**, “Mental Accounts and Consumption Sensitivity Across the Distribution of Liquid Assets,” Technical Report, Working paper 2024.
- Guttman-Kenney, Benedict and Andrés Shahidinejad**, “Unraveling information sharing in consumer credit markets,” *Available at SSRN 4629496*, 2024.
- Hisnanick, John J and Andreas Kern**, “The 2008 tax rebate and US household debt,” *Applied Economics Letters*, 2018, *25* (9), 592–596.
- Holm, Martin Blomhoff, Pascal Paul, and Andreas Tischbirek**, “The transmission of monetary policy under the microscope,” *Journal of Political Economy*, 2021, *129* (10), 2861–2904.
- Ilut, Cosmin and Rosen Valchev**, “Economic agents as imperfect problem solvers,” *The Quarterly Journal of Economics*, 2023, *138* (1), 313–362.
- Kaplan, Greg and Giovanni L Violante**, “A model of the consumption response to fiscal

- stimulus payments,” *Econometrica*, 2014, *82* (4), 1199–1239.
- **and** —, “The marginal propensity to consume in heterogeneous agent models,” *Annual Review of Economics*, 2022, *14* (1), 747–775.
- , **Benjamin Moll**, **and Giovanni L Violante**, “Monetary policy according to HANK,” *American Economic Review*, 2018, *108* (3), 697–743.
- Kotsogiannis, Christos and Plutarchos Sakellaris**, “MPCs Estimated from Tax Lottery, Survey and Administrative Data,” 2025.
- Kreiner, Claus Thustrup, David Dreyer Lassen, and Søren Leth-Petersen**, “Liquidity constraint tightness and consumer responses to fiscal stimulus policy,” *American Economic Journal: Economic Policy*, 2019, *11* (1), 351–379.
- Krueger, Dirk, Kurt Mitman, and Fabrizio Perri**, “Macroeconomics and household heterogeneity,” in “Handbook of macroeconomics,” Vol. 2, Elsevier, 2016, pp. 843–921.
- Kueng, Lorenz**, “Excess sensitivity of high-income consumers,” *The Quarterly Journal of Economics*, 2018, *133* (4), 1693–1751.
- Lewis, Daniel J, Davide Melcangi, and Laura Pilossoph**, “Latent heterogeneity in the marginal propensity to consume,” *FRB of New York Staff Report*, 2021, (902).
- Manski, Charles F**, “Measuring expectations,” *Econometrica*, 2004, *72* (5), 1329–1376.
- Miranda-Pinto, Jorge, Daniel Murphy, Kieran Walsh, and Eric Young**, “A Model of Expenditure Shocks,” Working Papers 20-04, Federal Reserve Bank of Cleveland February 2020.
- Parker, Jonathan A. and Nicholas S. Souleles**, “Reported Effects vs. Revealed-Preference Estimates: Evidence from the propensity to spend tax rebates,” Working Paper 23920, National Bureau of Economic Research October 2017.
- Parker, Jonathan A and Nicholas S Souleles**, “Reported effects versus revealed-preference estimates: Evidence from the propensity to spend tax rebates,” *American Economic Review: Insights*, 2019, *1* (3), 273–290.
- , **Jake Schild, Laura Erhard, and David Johnson**, “Household Spending Responses to the Economic Impact Payments of 2020: Evidence from the Consumer Expenditure Survey,” Working Paper 29648, National Bureau of Economic Research January 2022.
- 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, 2010, pp. 69–110.

Shapiro, Matthew D. and Joel Slemrod, “Consumer Response to Tax Rebates,” *The American Economic Review*, 2003, *93* (1), 381–396.

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

Stantcheva, Stefanie, “How to run surveys: A guide to creating your own identifying variation and revealing the invisible,” *Annual Review of Economics*, 2022, *15*.

Zeldes, Stephen P, “Consumption and liquidity constraints: an empirical investigation,” *Journal of Political economy*, 1989, *97* (2), 305–346.