

Boutros, Michael; Mijakovic, Andrej

Working Paper

The macroeconomic implications of coholding

Bank of Canada Staff Working Paper, No. 2024-16

Provided in Cooperation with:

Bank of Canada, Ottawa

Suggested Citation: Boutros, Michael; Mijakovic, Andrej (2024) : The macroeconomic implications of coholding, Bank of Canada Staff Working Paper, No. 2024-16, Bank of Canada, Ottawa, <https://doi.org/10.34989/swp-2024-16>

This Version is available at:

<https://hdl.handle.net/10419/299351>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

The Macroeconomic Implications of Coholding

by Michael Boutros¹ and Andrej Mijakovic²



¹ Financial Stability Department
Bank of Canada
Michael.Boutros@bankofcanada.ca

² European University Institute

Acknowledgements

The presented views are those of the authors, not necessarily of the Bank of Canada. This paper subsumes “Household Finances and Fiscal Stimulus in 2008” (Boutros, 2019) and benefited from discussions with Manuel Adelino, Lukas Altermatt, David Berger, Zahi Ben-David, Francesco Bianchi, Edouard Challe, Russell Cooper, Mory Elsaify, Daniel Greenwald, Andrea Lanteri, David Min, James Pinnington, Maarten van Rooij, Tao Wang, and Taylor Webley. We would also like to thank seminar and conference participants at the European University Institute, Bank of Canada, 5th Biennial Consumer Finance and Macroeconomics Conference, ESEM 2023, Australian National University, 2023 OzMac Workshop, University of Sydney, Reserve Bank of Australia, and Toronto Metropolitan University for helpful comments and suggestions.

Abstract

In the United States, 30% of households are coholders who simultaneously borrow on credit cards and hold liquid assets. This generates a rich distribution of gross wealth positions that underpins the distribution of net wealth often used to calibrate macroeconomic models. We show that, beyond their role in constructing net wealth, gross positions in liquid assets and liquid debt are important in determining how households consume, save, and repay debt in response to positive income shocks. We build a quantitative model that generates the coholding observed in the data and matches observed marginal propensities to consume, save, and repay debt. The model highlights that fiscal transfers are more effective in stimulating demand if targeted at households with low gross positions instead of low net liquid wealth, while debt relief is less effective overall in the short run but achieves large consumption gains in the long run.

Topics: Central bank research; Fiscal policy; Monetary policy; Economic models

JEL codes: E21, E44, E62, G51

Résumé

Aux États-Unis, 30 % des ménages sont considérés comme emprunteurs-épargnants, c'est-à-dire qu'ils empruntent sur leurs cartes de crédit alors qu'ils détiennent des actifs liquides. Une riche distribution des positions du patrimoine brut en résulte, laquelle sous-tend la distribution du patrimoine net souvent utilisée pour calibrer les modèles macroéconomiques. Nous montrons l'importance des positions du patrimoine brut – au-delà de leur utilité dans le calcul du patrimoine net – pour déterminer comment les ménages consomment, épargnent et remboursent leurs dettes à la suite de l'augmentation de leur revenu. Nous utilisons un modèle quantitatif qui illustre les situations d'emprunt-épargne sur la base de données et dont les résultats correspondent aux propensions marginales observées empiriquement à consommer, à épargner et à rembourser les dettes. Le modèle fait ressortir que les transferts budgétaires sont plus efficaces pour stimuler la demande lorsqu'ils ciblent les ménages dont le patrimoine brut est faible plutôt que ceux dont les liquidités nettes sont faibles. Les mesures d'allégement de dette sont, quant à elles, moins efficaces dans l'ensemble à court terme, mais donnent lieu à une importante croissance de la consommation à long terme.

Sujets : Recherches menées par les banques centrales, Politique budgétaire, Politique monétaire, Modèles économiques

Codes JEL : E21, E44, E62, G51

1 Introduction

Modern heterogeneous agent models attribute a fundamental role to the distribution of wealth in the economy. At the individual level, wealth affects the allocation of resources across consumption and savings. Aggregated up, the distribution of wealth affects the aggregate response of consumption to a change in resources and, as such, the efficacy of fiscal and monetary policy.

The relevant definition of wealth is typically considered to be *net total* wealth, while more recent contributions highlight the role of *net liquid* wealth (Kaplan and Violante, 2014). At the same time, *net* liquid wealth positions can differ vastly from their *gross* positions. In the United States, around one-third of households report regularly borrowing on their credit card by revolving balances month-to-month. Of those that borrow on their credit cards, at an average interest rate of 14%, the vast majority report simultaneously holding liquid assets such as cash. These facts, and others, have been well documented in the household finance literature since Gross and Souleles (2002).

In this paper, we integrate the concept of coholding from the household finance literature into a wider macroeconomic context and study the relevance of the joint distribution of liquid assets and debt for macroeconomic policy. We make two key contributions to the literature. First, we show empirically that households' consumption, saving, and borrowing responses to income changes are a function of both their liquid asset *and* debt position. Grouping households across the distribution of net liquid wealth, instead, masks substantial heterogeneity. Second, we develop a quantitative model with fully rational coholding of liquid assets and debt, which matches the empirically observed household behavior. We apply our model to study the implications of coholding for fiscal and monetary policy, highlighting the role that the joint distribution of assets and debt plays in the aggregate marginal propensities to consume, save, and repay debt.

In the first part of the paper, we document the tight empirical connection between household behavior and gross wealth positions. Extending beyond the typical analysis of net liquid wealth and the marginal propensity to consume, we study the relation between both liquid assets and liquid debt with consuming, saving, and repaying debt. We show that gross wealth positions matter beyond net positions in the determination of marginal propensities. Our key finding is that liquid debt dampens the marginal propensity to consume. All else equal, a decrease in wealth due to an additional dollar of liquid debt decreases the marginal propensity to consume and increases the marginal propensity to repay debt. On the other hand, a decrease in wealth due to one fewer dollar of liquid assets leaves the marginal propensity to consume roughly unchanged and increases the

marginal propensity to save.

The contribution of our empirical analysis is to demonstrate that using low liquid wealth to proxy for hand-to-mouth households may confound two distinct groups of households: true hand-to-mouth households that have low net positions and low gross positions, and coholding households that have low net positions because of large and offsetting gross positions. While true hand-to-mouth households have large consumption responses, coholders have small consumption responses.

The second part of the paper contributes a model that endogenously features both true hand-to-mouth and coholding households. The model combines two novel features. First, households make explicit saving and borrowing decisions, represented by two distinct controls and corresponding state variables. Second, we add a liquidity-in-advance constraint, in the style of Lucas (1982), which induces households to target a specific amount of liquid assets. Combined with the standard mechanism of holding a buffer-stock level of net wealth to smooth consumption over time, households have distinct demand for both liquid assets and liquid debt. As *ex ante* identical households are exposed to heterogeneous income shocks, a subset become true hand-to-mouth households while another subset become coholders.

We use the model to understand the implications of coholding for fiscal and monetary policy. Our analysis centers on the role that the joint distribution of liquid assets and liquid debt plays in determining aggregate consumption, saving, and repaying debt. Direct fiscal transfer programs have become a common policy tool to stimulate consumption. Untargeted programs distribute cash to all households in the economy, coholders and hand-to-mouth alike. As the intensity of coholding has increased over time, the efficacy of untargeted programs in stimulating consumption has decreased. Policymakers designing targeted fiscal transfers face the daunting challenge of correctly identifying true hand-to-mouth households, and our analysis shows that poorly targeted programs can easily be less effective than untargeted programs.

A key contribution of our model is the ability to consider the dynamic effects of different fiscal policies. We study not only the immediate consumption response of untargeted and targeted programs, but also how saving and repaying debt using fiscal transfers can lead to elevated consumption dynamics in the long-run. For example, compared to untargeted or other targeted programs, direct cash transfers to coholders have the smallest short-run and largest long-run consumption responses.

We also use the model to study fiscal debt relief programs that have recently become part of the fiscal policy toolkit. We consider the immediate and dynamic impacts of debt relief programs that implicitly target coholding households, and compare them to fiscal

cash transfer programs. We again highlight that the short- and long-run impacts of these programs can differ significantly. An important implication of our analysis is that any discussion of the efficacy of fiscal policies must clearly state the horizon over which policies should be evaluated.

Finally, we discuss the implications of coholding for monetary policy in a large class of general equilibrium models. We construct the model's intertemporal Keynesian cross, as defined by Auclert et al. (2023), for not only consumption, but also saving and repaying debt. We demonstrate that the structure of financial markets, which will determine the pass-through of monetary policy to saving and borrowing rates, is crucial for fully understanding how coholding affects the macroeconomy. Future work in this area will enrich existing models with heterogeneous banks and endogenous credit spreads to include coholding households.

Related Literature This paper adds to the large literature on coholding of liquid assets and debt. Several theoretical explanations of the coholding puzzle have been put forward (Bertaut et al., 2009; Telyukova and Wright, 2008; Telyukova, 2013; Fulford, 2015; Druedahl and Jørgensen, 2018; Gorbachev and Luengo-Prado, 2019). We build on the idea proposed in Telyukova (2013) that households cohold debt and assets due to liquidity demand. Our contribution is to integrate this mechanism into a standard consumption-savings model to generate endogenous coholding and study its implications for fiscal policy. Our model generates coholding of liquid assets and debt, while the model built in Kaplan and Violante (2014) generates coholding of liquid and illiquid wealth. Kosar et al. (2022) and Lee and Maxted (2023) also focus on understanding the relevance of debt for stimulative fiscal transfers. Kosar et al. (2022) introduce a debt price schedule into a standard incomplete markets model and show that by using transfers to reduce debt, households increase their individual welfare by reducing the interest rate paid on debt. Lee and Maxted (2023) show that in an economy with present bias, credit card borrowers do not need to be close to their borrowing constraint to have an elevated marginal propensity to consume. Relative to these papers, our focus is on studying the marginal propensities to consume and repay debt along the joint distribution of liquid assets and debt.

We also contribute to the empirical literature on marginal propensities to consume by focusing on the role of gross wealth instead of net wealth. In general, the literature has documented a negative relation between the marginal propensity to consume and wealth, and standard economic models generate this behavior (Kaplan and Violante, 2022). Recently, however, several studies have found a flat MPC across the distribution of liquid wealth (Bunn et al., 2018; Christelis et al., 2019; Fuster et al., 2021). At the same time, other

studies have more closely examined the role of debt, documenting the decreasing MPC across the distribution of debt (Jappelli and Pistaferri, 2014; Sala and Trivin, 2021; Kosar et al., 2022) and how households adjust their debt positions in response to income changes (Agarwal et al., 2007; Sahm et al., 2015; Boutros, 2019; Coibion et al., 2020; Fagereng et al., 2021). We reconcile these findings by showing that consumption, saving, and debt repayment behavior are functions of the joint distribution of assets and debt. Studying assets, debt, or net wealth in isolation is not sufficient to characterize household behavior. Moreover, we provide evidence that the relevant debt statistic is credit card debt.

Finally, our fiscal policy analysis adds to the literature on debt-dependent fiscal multipliers. Previous studies primarily find multipliers that are increasing in the level of debt, often based on micro-level evidence (Dynan et al., 2013; Mian et al., 2013; Klein, 2017; Baker, 2018; Bernardini and Peersman, 2018; Demyanyk et al., 2019; Bernardini et al., 2020). Our results do not necessarily contradict this literature as we focus on credit card debt only, while most of the existing literature focuses on either aggregate debt or other debt components.

2 Credit Card Debt and the Coholding Puzzle

The coholding of low-return cash and high-cost credit card debt has been thoroughly documented in the household finance literature. In this section, we reproduce a baseline set of facts of coholding over time in the United States, and document coholding in several other jurisdictions in the appendix. Our goal is to then take coholding as given and study the implications for the macroeconomy in a structural model. As such, we survey the theoretical literature that proposes explanations for this “puzzle,” and explain our motivation for modeling the demand for coholding as arising from the rational demand for liquidity due to a liquidity-in-advance constraint.


2.1 The Prevalence of Credit Card Borrowing

To establish a baseline set of facts regarding coholding in the United States, we use the Survey of Consumer Finances (SCF), a nationally representative sample of US households fielded roughly every three years.¹ Figure 1a plots the distribution of credit card holders in the SCF. According to the survey, roughly 70% of households have credit cards, but 40% of households are convenience users that report paying their entire balance in full and therefore never borrow on their credit cards. Almost one-third of all households

¹We restrict our sample to households aged 25–65 with annual income above 1,000 USD.

report having at least one credit card and paying less than the full statement balance each month.

Figure 1: Extensive Margin of Credit Card Holding and Borrowing in the United States



Notes: Source: 2016 SCF.

Figure 1b shows that credit card holders and borrowers are found across the entire distribution of income. In fact, higher income households are both more likely to have credit cards and, except for the last quartile, to revolve debt on their credit cards.


2.2 Coholding of Cash and Credit Card Debt

In this section, we focus on the subset of credit card users who are borrowers and study the composition of their balance sheets on the intensive margin. Figure 2a plots the average level of liquid assets and debt for coholders in each quartile of income.² In the lowest quartile of income, households hold roughly \$3,000 in liquid debt and \$2,000 in liquid assets, yielding a negative net liquid wealth. Liquid assets and debt both increase by \$1,000 in the second quartile, yielding again a negative net liquid wealth. For those in the third quartile of income, net liquid wealth is only slightly higher by a few hundred dollars, but liquid assets and debt both increase to around \$8,000 and \$7,000, respectively. In the top quartile of income, liquid assets increase to almost \$15,000, while liquid debt increases to only \$10,000, yielding a positive net liquid wealth of just over \$5,000.

Figure 2b provides additional detail for the degree of coholding for each income quartile. In the lowest income quartile, only 10% of households report no coholding, and just under 40% report holding enough liquid assets to completely pay off their credit card

²Liquid assets are defined as funds held in checking and savings accounts. Liquid debt is defined as credit card debt that is measured by the balance due on the credit card after the last statement was paid.

Figure 2: Intensive Margin of Coholding for Credit Card Borrowers



Notes: Data from the 2016 SCF. In the left panel, we additionally restrict the sample to households between the 1st and 95th percentile of the liquid asset, debt, and wealth distribution.

debt. As income increases, the fraction of coholding increases; for the top income quartile, over 60% of households report enough liquid assets to completely pay off their credit card debt, while under 2% report no coholding.

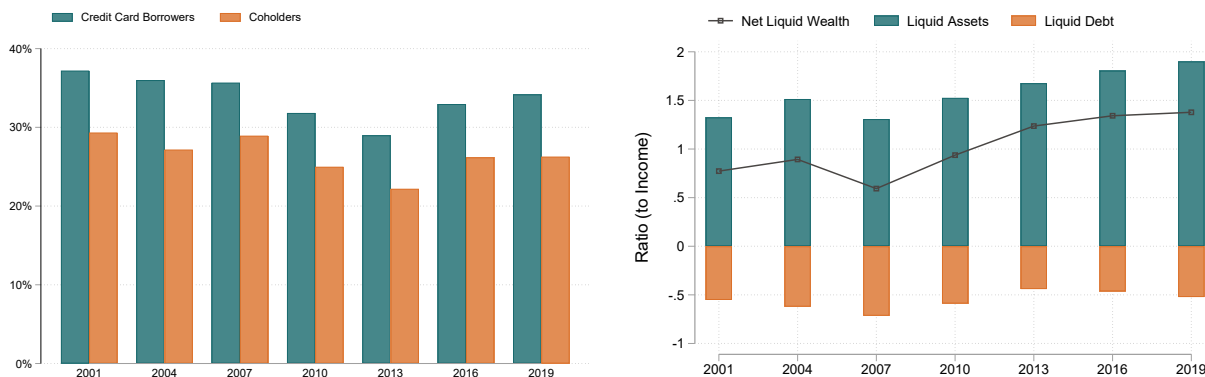
2.3 Coholding Over Time

So far, we have inspected coholding behavior at one particular point in time. However, the coholding of liquid assets and debt has been prevalent since at least the beginning of the 2000s. Figure 3a reports the percentage of credit card borrowers and coholders across several waves of the Survey of Consumer Finances from 2001–2019. We define coholders as credit card borrowers who hold gross liquid assets and debt equal to at least 10% of their monthly income.³ The share of credit card borrowers has been fluctuating between 30% and 40% of the population, except for a brief period of deleveraging after the global financial crisis. Similarly, the share of coholders has largely been tracking the share of credit card borrowers and has fluctuated between 20% and 30% for the majority of the sample period.

While the extensive margin of coholding has been fairly stable over time, the intensive margin has experienced larger changes. Figure 3b reports the evolution of the aggregate stock of liquid assets and debt relative to income over time. After a contraction around the time of the global financial crisis, both net and gross wealth positions have expanded. Between 2007 and 2013, net liquid wealth increased due to both an increase in liquid as-

³In Appendix A.1, we show that the discussion in this section is robust to alternative measures of credit card borrowing and coholding used in the literature.

Figure 3: Coholding and Composition of Liquid Wealth Over Time



(a) Share of Credit Card Borrowers and Coholders

(b) Composition of Liquid Wealth

Notes: Data from triennial SCF waves between 2001 and 2019. The left panel plots the fraction of all households that are credit card borrowers and coholders. Credit card borrowers are households that report having a credit card and not paying off their credit card balance fully. Coholders are households that hold more than 10% of monthly income in liquid assets and debt and report revolving credit card debt. The right panel reports average liquid wealth, asset, and debt holdings relative to income. Liquid assets and debt are winsorized at the 99% level.

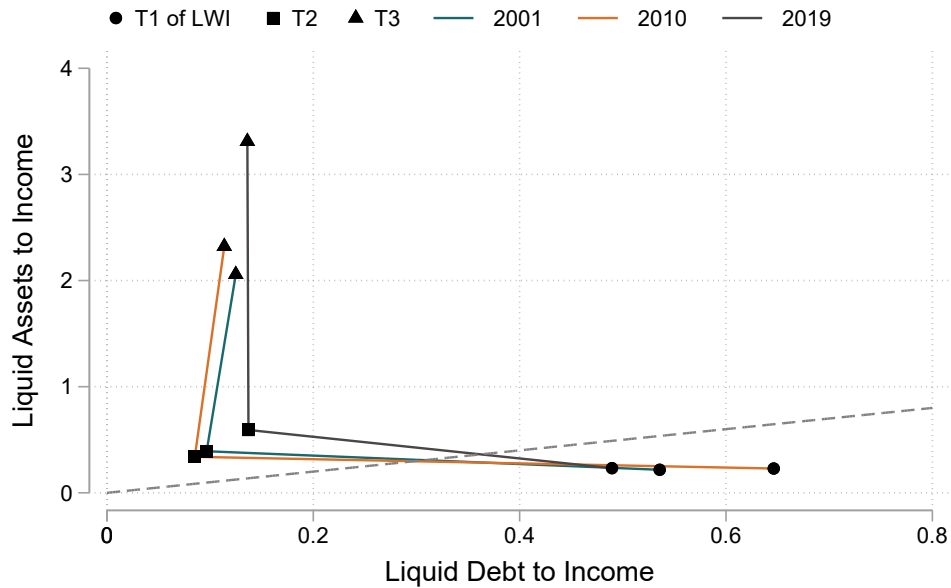
sets and decrease in liquid debt. More recently, the stock of liquid debt has also begun to increase again, reflecting partly a larger share of credit card borrowers. From 2013 to 2016, net wealth increased by less than half of the increase in prior years, fueled by a smaller increase in liquid assets but also an increase in liquid debt. From 2016 to 2019, net wealth was almost flat, because despite an increase in liquid assets, liquid debt increased by almost as much. Focusing solely on changes in net positions masks this compositional shift of liquid wealth which, as we will demonstrate, has important implications for changes in aggregate household behavior.

2.4 The Joint Distribution of Liquid Assets and Debt

An important implication of coholding is that focusing solely on the distribution of liquid wealth without examining the gross components masks heterogeneity in household balance sheets. After documenting the heterogeneity in this section, we turn to analyzing the importance of this heterogeneity and how the net level of wealth is an insufficient statistic for fully characterizing household behavior.

Figure 4 plots one approach to characterizing the joint distribution of liquid assets and debt as it relates to the distribution of liquid wealth. Households are sorted into terciles of liquid wealth and average liquid assets and liquid debt are calculated within tercile. The

Figure 4: Joint Distribution of Liquid Assets and Debt



Notes: Data from SCF waves 2001, 2010, and 2019. Liquid assets and debt are winsorized at the 99% level.


dashed 45-degree line represents a zero net wealth position: above this line, households have positive net positions; below it, they have negative net positions.

In the presence of no coholding, all households would lie along either the horizontal or vertical axis. This is clearly not the case as all points are on the interior of the graph. In each year, the joint distribution has a distinct L-shape, which we confirm remains if we increase the number of percentiles from 3 to 5 or 10. This figure shows that coholding is prevalent throughout the entire distribution of liquid wealth. In the first tercile, households have low but positive cash balances and large amounts of credit card debt. These households all hold negative net positions. In the third tercile, households have large cash balances and moderate credit card debt, yielding positive net positions.

Households in the second tercile have a similar level of liquid assets as the lowest tercile but less liquid assets than the top tercile. These households also have net positive positions despite holding the same level of liquid assets as the poorest households. Focusing solely on net wealth hides this heterogeneity, which we will show is an important determinant of household behavior.

This figure also demonstrates the increasing prevalence of coholding over time as the L-distribution moves north-east in each year. If the distribution were to shift out exactly along the 45-degree line, then the gross positions would increase proportionally and the

Figure 5: Coholding and Hand-to-Mouth Households Over Time



Notes: Population shares of hand-to-mouth and coholder households in SCF waves from 2001 to 2022. See Section 2.5 for more details and definitions.

net levels of wealth would remain constant. Although this is not exactly the case, this figure shows that large changes in gross positions underpin the relatively small changes in net positions over the years.

2.5 Coholding and Hand-to-Mouth Households

The fraction of “spender” or “hand-to-mouth” households that have large and immediate consumption responses to positive income shocks is an important calibration target for many modern macroeconomic models. Since hand-to-mouth status is a latent variable, it is often proxied through the same household balance sheet variables that determine coholding status, generating an overlap in households that are both coholders and hand-to-mouth.

In seminal work, Kaplan and Violante (2014) identify two different groups of hand-to-mouth households: poor hand-to-mouth based on their total net wealth, as had been prevalent in the literature, and a new measure of wealthy hand-to-mouth that may have large total net worth but low liquid net wealth. Hand-to-mouth status is measured using net liquid wealth, but as the previous section demonstrates, unpacking the joint distribution of liquid assets and debts allows for a decomposition into coholding status. As such, we can further divide poor and wealthy hand-to-mouth households into coholders and non-coholders.

Panel (a) of Figure 5 plots the population shares for these four groups of households. Following the literature, we label households as hand-to-mouth if their net worth is less than two weeks of earnings; for poor hand-to-mouth, we measure net worth as net total

wealth; and for wealthy hand-to-mouth, we use net liquid wealth. In this panel, we use the narrow definitions of assets and debt employed in this paper, while Panel (b) uses the more common broader definition of liquid wealth (Kaplan et al., 2014).

As our analysis will show, coholders do not have large consumption responses out of positive income shocks, which is a key attribute of hand-to-mouth households. As such, failing to account for coholders will significantly overcount the share of “true” hand-to-mouth non-coholder households. In 2016, for example, 51.3% of households are hand-to-mouth, but excluding coholders reduces this to 34.3%. The majority of this reduction comes from the 13.5% of wealthy hand-to-mouth households that are also coholders. By construction, wealthy hand-to-mouth households have low net liquid wealth, and almost one-third are also coholders. Similarly, in the broad definition of liquid wealth, the share of hand-to-mouth households decreases from 41.2% to 26.3% when excluding coholders, with 11.4% of the reduction from wealthy hand-to-mouth households.

2.6 Proposed Explanations for Coholding

Given the empirical prevalence of coholding (Gross and Souleles, 2002; Vihriälä, 2020; Gathergood and Weber, 2014; Gathergood and Olafsson, 2022; Greene and Stavins, 2022; Pulina, 2024), a large body of work has sought to understand why households engage in this “puzzling” behavior.

One strand of the literature resolves the puzzle by arguing that households are not fully rational and coholding occurs due to behavioral biases such as mental accounting (Choi and Laschever, 2018; Batista et al., 2023; Medina and Pagel, 2023). Others have put forth models of rational coholding driven by credit access risk (Fulford, 2015; Druedahl and Jørgensen, 2018; Gorbachev and Luengo-Prado, 2019), bankruptcy (Lehnert and Maki, 2002; Mankart, 2014; Lopes, 2008), household composition (Bertaut et al., 2009), or liquidity (Telyukova and Wright, 2008; Telyukova, 2013; Zinman, 2007).

In this paper, our goal is to take coholding as given and study the implications for the macroeconomy. We remain agnostic as to the exact source of coholding and acknowledge that all of the mechanisms described above are likely active to varying degrees, both across households and even within a single household. Although we eventually use a liquidity-in-advance constraint to generate rational coholding, the common denominator among all of these mechanisms is persistent coholding. Households cohold because of their preferences and structure of the economy, and not simply because of transient circumstances such as a sequence of negative income shocks. As we discuss later, insofar as all of these mechanisms generate persistent coholding behavior, the implications for the

macroeconomy will be similar.

3 Coholding and Marginal Propensities in the Data

Having established the prevalence of coholding, we begin our analysis of its implications for the macroeconomy by studying the implications of coholding for individual household behavior. After establishing how coholding affects each household, we integrate over the distribution of households to understand aggregate behavior.

The starting point for this analysis is the large literature studying the marginal propensity to consume out-of-income shocks and how it relates to liquid wealth. As demonstrated in the previous section, the distribution of liquid wealth masks heterogeneity with respect to the underlying distributions of liquid assets and debt. We study the marginal propensity to consume as it relates to these underlying joint distributions. Further, instead of focusing only on the marginal propensity to consume, we separately study the propensities to save and repay debt, which are typically joined together as the propensity to increase wealth and ignored since they are the residual of increasing consumption.

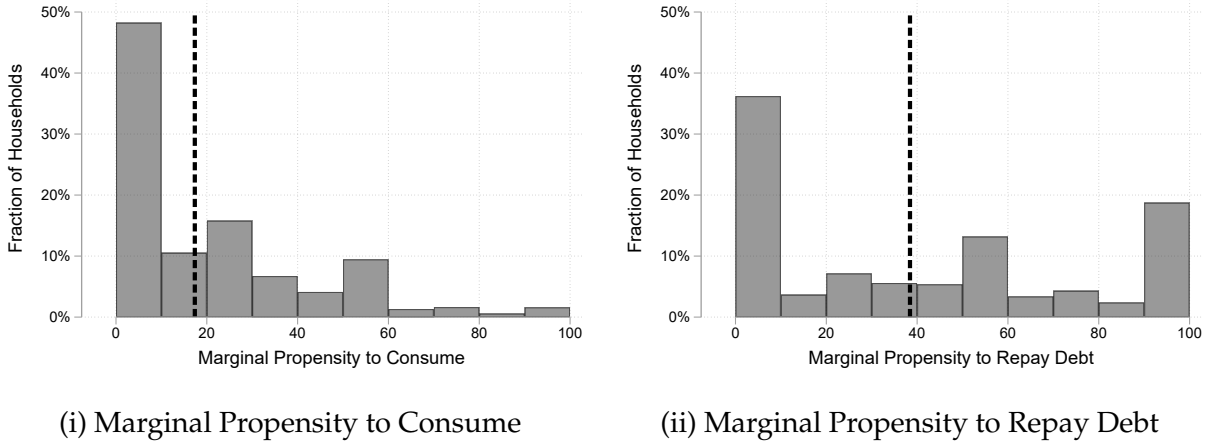
We extend the literature in two directions: documenting facts about three distinct actions (consume, save, repay debt) across the joint distribution of liquid assets and debt. We begin with laying out simple summary statistics on the marginal propensities to consume, save, and repay debt, and demonstrate how these actions vary across the distribution of net liquid wealth. We then turn to liquid assets and debt, showing how household behavior varies across each marginal distribution and across the joint distribution. To establish a baseline set of facts regarding coholding and marginal propensities in the United States, we use the Survey of Consumer Expectations, a nationally representative survey of US households fielded by the Federal Reserve Bank of New York. The data, which has been used extensively in numerous papers such as Fuster et al. (2021) and Mijakovic (2023), is described in Appendix A.2.

3.1 Distributions of Marginal Propensities

Figure 6 plots histograms of the marginal propensities to consume and repay debt.⁴ The average propensity to consume is 16.8%, in line with the literature, especially recent estimates that revisit previous stimulus programs using new techniques for event study

⁴As in the SCF, we restrict the analysis to households aged 25–65 with annual income above 1,000 USD. We drop households for which we do not observe income, marginal propensities, or liquid assets and debt. Whenever we report marginal propensities across balance sheet items, we also trim the top 1% of the liquid asset and debt distribution, and the top 1% and bottom 1% of the liquid wealth distribution.

Figure 6: Marginal Propensities to Consume and Repay Debt



Notes: Marginal propensities from the SCE 2015–2018, unweighted distributions.

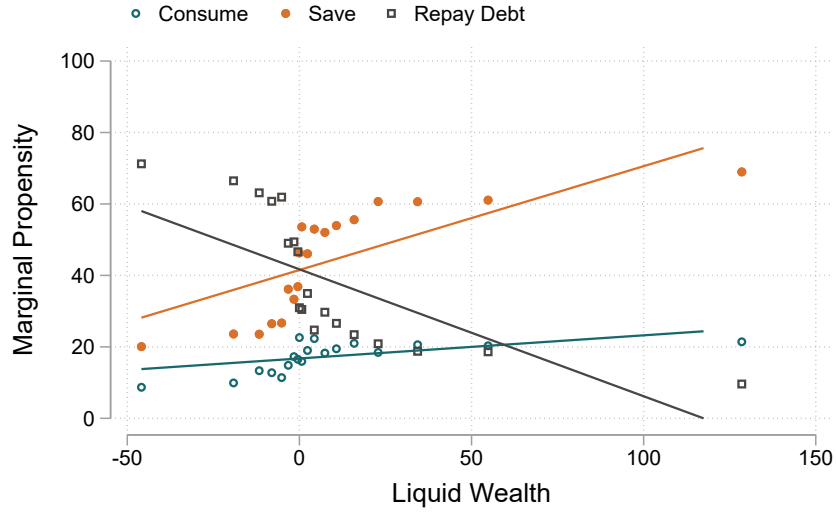
designs (Borusyak et al., 2024). The average propensity to repay debt is 39.2% and the propensity to save is 44.0%. These averages include a nontrivial number of households that report a zero response in one or multiple categories, consistent with findings in other settings (Fuster et al., 2021; Misra and Surico, 2014). The distribution of propensities to consume features few households with MPCs larger than 50%, while the distribution of propensities to repay debt is bimodal, with large fractions at both 0% and 100%.

3.2 Household Behavior and Net Liquid Wealth

Figure 7 plots the marginal propensities to consume, save, and repay debt over the distribution of liquid wealth. Households are divided into 20 equal groups based on liquid wealth. Each marker plots the average propensity within that group, and the lines of best fit are estimated over all households.

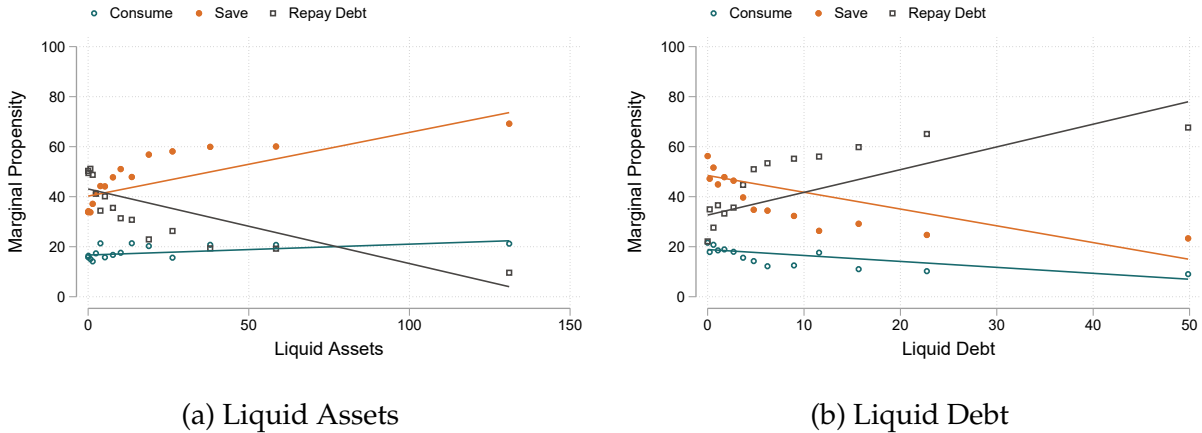
The consumption gradient with respect to liquid wealth is slightly positive, which is consistent with a number of recent papers but inconsistent with a large literature documenting the negative relation between the propensity to consume and liquidity (see the discussion in Fuster et al. (2021)). The flat consumption response implies that the residual of consumption, that is, the propensity to increase wealth, is also flat. However, decomposing the propensity to increase wealth into the propensities to save and repay debt reveals vast heterogeneity across the distribution of liquid wealth. The average propensity to repay debt for the poorest households is almost 75%, but decreases by more than 60 pp (percentage points) to just above 10% for the richest households. Almost all of this change comes from a switch into repaying debt, not increasing consumption; the average

Figure 7: Marginal Propensities Across the Distribution of Liquid Wealth



Notes: Marginal propensities from the SCE 2015–2018. Liquid wealth is defined as liquid assets minus liquid debt.

Figure 8: Marginal Propensities Across the Marginal Distributions of Gross Liquid Wealth



Notes: Marginal propensities from the SCE 2015–2018. Liquid assets are defined as the sum of checking and savings accounts plus idle money in brokerage accounts. Liquid debt is defined as credit card debt, measured by the balance due after the last statement was paid.

propensity to save is 20% for the poorest households and increases by 50 pp to 70% for the richest households.

3.3 Household Behavior and Gross Liquid Wealth

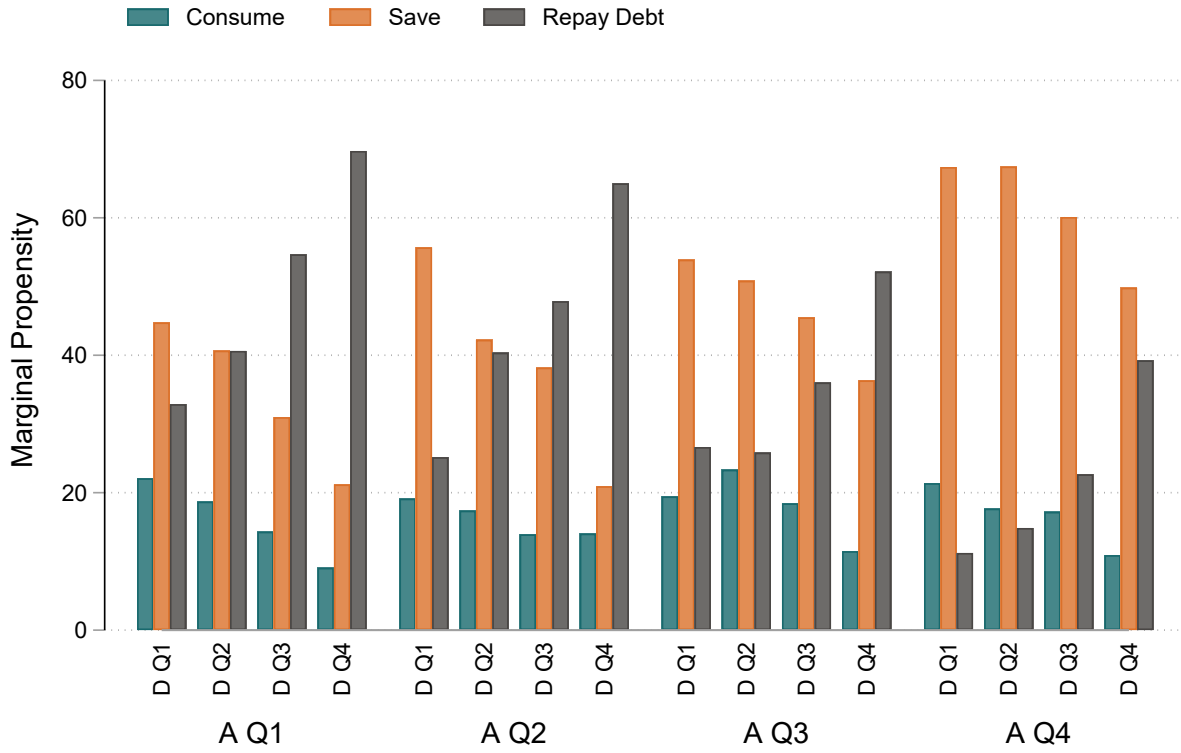
Figure 8 plots the relationship between the marginal propensities to consume, save, and repay debt across the marginal distribution of assets and debt. Intuitively, the relationship between household behavior and liquid assets is almost identical to that of liquid wealth. As liquid assets increase, the propensity to consume is essentially flat, but the propensity to repay debt decreases significantly and the change is made up completely in an increasing propensity to save.

The relationship for liquid debt tells the opposite story. As debt increases, the propensity to repay debt increases, and both the propensities to save and consume decrease commensurately. It is intuitive that households with more debt would choose to pay down their debt instead of save because credit card debt has a much higher interest rate than cash. The tradeoff between consumption versus debt repayment is less clear but is elucidated with our model. In the data and in our model, households with high debt that are not at the borrowing constraint will choose to repay debt instead of consume, because doing so will increase lifetime wealth and therefore permanent consumption. Thus, households with low liquid wealth due to high credit card balances will have smaller propensities to consume than might be expected based on their net position alone. This insight may help resolve the inconsistency noted above with respect to the relationship between liquid wealth and the consumption response, since different papers in the literature define “liquidity” as either liquid assets or net liquid wealth.

To study the joint distribution of liquid assets and debt, we divide households into quartiles along each dimension. Figure 9 plots the average propensities to consume, save, and repay debt for each of the 16 groups. The takeaway from this analysis is that *within* each quartile of liquid assets, the patterns discussed above across the marginal distribution of liquid debt still hold. In the first quartile of liquid assets, the propensities to consume and save are decreasing in liquid debt, while the propensity to repay debt is increasing. The same pattern remains true even in the fourth quartile of liquid assets.

Given that households in the fourth quartile of liquid assets typically have very large positive net wealth, this implies that even the wealthiest or most liquid households are responsive to their individual debt positions. Similarly, even the poorest or least-liquid households vary their behavior depending on their debt positions. This highlights the importance of studying and modeling the joint distribution of liquid wealth for modeling household behavior.

Figure 9: Marginal Propensities Across the Joint Distributions of Gross Liquid Wealth



Notes: Marginal propensities from the SCE 2015–2018. The figure reports average propensities by quartile of liquid assets and liquid debt. Liquid assets are defined as the sum of checking and savings accounts plus idle money in brokerage accounts. Liquid debt is defined as credit card debt, measured by the balance due after the last statement was paid.

3.4 Regression Analysis

In this section, we use linear regressions to confirm the visual findings above. In the first column of Table 1, we verify that the propensity to consume is flat in liquid assets. The estimate is both statistically and economically insignificant; increasing liquid assets by \$1 increases the MPC by only 0.01 pp. On the other hand, increasing liquid debt by \$1 decreases the MPC by 0.20 pp. Together, these estimates imply that a \$1 increase in liquid wealth due to increasing liquid assets has a negligible impact on the propensity to consume, while the same increase from a reduction in debt will meaningfully increase the propensity to consume.

Liquid assets are more relevant for the propensities to save and repay debt. The marginal propensity to save increases by 0.24 pp for a \$1 increase in liquid assets, but decreases by 0.67 pp for a \$1 increase in liquid debt. Again, this implies that a change in net liquid wealth from an increase in assets or decrease in debt will have opposite pre-

Table 1: Regressions of Marginal Propensities on Household Liquid Balance Sheet

	(1) Spend	(2) Spend	(3) Save	(4) Save	(5) Repay Debt	(6) Repay Debt
Liquid Wealth	0.043** (0.014)		0.321*** (0.025)		-0.364*** (0.025)	
Liquid Assets		0.007 (0.016)		0.244*** (0.027)		-0.251*** (0.024)
Liquid Debt		-0.204*** (0.038)		-0.667*** (0.059)		0.872*** (0.078)
N	2,742	2,742	2,742	2,742	2,742	2,742
R^2	0.042	0.051	0.143	0.158	0.155	0.185

Notes: Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Controls include age, gender, race, marital status, education, geography, and survey date.

dictions for the propensity to save. The same also holds true for the propensity to repay debt, which is the most sensitive to both changes in liquid assets and debt. Increasing liquid assets by \$1 decreases the propensity to repay debt by 0.25 pp, while increasing liquid debt by \$1 increases the propensity to repay debt by 0.87 pp.

In Appendices A.3 and A.4, we illustrate that our results remain virtually unchanged if we estimate the regression equation using standardized measures of balance sheet items instead of USD values, or if we include a wider set of controls such as income, illiquid assets and debt, home equity lines of credit, housing status, and measures of financial literacy. This corroborates the evidence put forward in Kaplan et al. (2014) that the liquid component of households' balance sheets matters more for marginal propensities than the illiquid component.

In Appendix A.5, we provide suggestive evidence for the relevance of the joint distribution of liquid assets and debt for marginal propensities that goes beyond the households surveyed in the SCE. In particular, we revisit the empirical evidence in Jappelli and Pistaferri (2014) and Christelis et al. (2019), who study the relation between marginal propensities to consume and liquid assets using Italian and Dutch data, respectively. We extend their analysis by explicitly focusing jointly on assets and debt and find results that are consistent with the ones in the SCE. Debt matters for the determination of marginal propensities beyond its effect on the level of net wealth.

4 Model

To explore the prevalence of credit card debt and the implications of coholding on the response to income shocks, we build a model of consumption and savings in which households optimally and rationally cohold both liquid assets and debt. Our approach to arriving at the macroeconomic implications of coholding is to aggregate microeconomic coholding at the individual household level across the entire distribution of households in the economy.

Building on the empirical analysis above, our modeling approach can be divided into two interacting blocks. First, the model must generate the correct marginal propensities to consume, save, and repay debt for each household along the distributions of liquid assets and debt. Second, the model's steady state distributions of liquid assets and debt must resemble the empirical distributions. Then, by integrating individual household behavior over the distribution of households, we arrive at the aggregate consumption, saving, and borrowing functions that are built on the proper microfoundations.

4.1 Environment and Financial Markets

The model environment is populated by a continuum of identical households that are *ex ante* identical and live infinitely. Time is discrete and the model period is one month. In each period, households receive stochastic endowment income, y_t , which will be calibrated in the next section. A representative financial institution serves households by inelastically providing two financial instruments. Households can save in a one-period liquid asset, a_{t+1} , and simultaneously borrow in one-period liquid debt, d_{t+1} . The rate of return on saving is $R_a = 1 + r$, and there is a positive wedge between borrowing and saving, $\delta > 0$, such that $R_d = R_a + \delta$.

4.2 Households

In each period, households take as given their stochastic income process and the interest rates set by the financial institution. Each household then makes consumption, saving, and borrowing decisions subject to its budget constraint, borrowing constraint, and liquidity-in-advance constraint.

4.2.1 Budget and Borrowing Constraints

In each period, the household's budget constraint equates income to expenditure:

$$y_t + a_t - d_t = c_t + \frac{a_{t+1}}{R_a} - \frac{d_{t+1}}{R_a + \delta}.$$

Income consists of endowment income, y_t , and liquid assets, a_t , net of liquid debt, d_t . On the expenditure side, the household chooses consumption, c_t , saves into the liquid asset at price R_a^{-1} , and borrows with liquid debt at price $(R_a + \delta)^{-1}$. The household is subject to an exogenous borrowing constraint, ϕ , such that it cannot borrow more than this amount:

$$d_{t+1} \leq \phi.$$

In addition, both financial instruments must be weakly positive. Net wealth is defined as $w_t \equiv a_t - d_t$. When $\delta = 0$, the budget constraint collapses to that of the standard model in which only the net level of wealth is relevant.

4.2.2 Liquidity-in-Advance Constraint

The key ingredient in our model is the addition of a liquidity-in-advance constraint for consumption:

$$c_t \leq \omega(y_t, a_t).$$

The household's total consumption, c_t , is limited by its access to cash-on-hand in the form of income and liquid assets and, specifically, liquid assets chosen in the previous period. This yields an intertemporal connection between saving and consumption in the future.

This constraint captures the theory developed and tested in Telyukova and Wright (2008) and Telyukova (2013). The core idea is that a fraction of household consumption can only be paid for using the liquid asset. Telyukova uses the Survey of Consumer Expenditures and partitions goods by their ability to be purchased using cash. She finds that households that spend more on these goods are those most likely to hold large amounts of cash and hold credit card debt. While we use a liquidity-in-advance constraint to remain as close as possible to a standard model of consumption and savings, Telyukova builds a model in which every period is divided into two stages and the money demand decision is made explicitly prior to consumption.

The intuition behind this mechanism for the demand for liquidity is very similar to the notion of demand for money from the macroeconomic literature. Indeed, the liquidity-in-advance constraint we use is inspired directly by Lucas (1982). In addition, Telyukova

(2013) notes that “it is useful to think about consumer debt through the lens of modern monetary economics.” In the money demand literature, money is held due to its transactional value, resolving the puzzle of coholding zero-return cash instead of positive-return government bonds. Similarly, households demand liquid assets in order to transact for a fraction of their consumption goods. The household simultaneously wishes to borrow from its future self, generating a demand for liquid debt.

4.2.3 Preferences and Optimization Problem

Households value consumption and have standard time preferences governed by a discount factor, β . Suppressing notation indicating their state-dependence, the household chooses a sequence of consumption, saving, and borrowing to maximize lifetime utility:

$$\max_{\{c_t, a_{t+1}, d_{t+1}\}} E_0 \sum_{t=0}^{\infty} \beta^t u(c_t),$$

subject to the borrowing, budget, and liquidity-in-advance constraints above.

The first order conditions for saving and borrowing yield two distinct Euler equations. Focusing on the interior solution in which no constraints are binding, these can be written as:

$$u'(c_t) = \beta R E_t \left[u'(c_{t+1}) + \mu_{LIA,t+1} \left(\frac{\partial \omega(y_{t+1}, a_{t+1})}{\partial a_{t+1}} - 1 \right) \right], \quad (1)$$

$$u'(c_t) = \beta(R + \delta) E_t [u'(c_{t+1}) - \mu_{LIA,t+1}], \quad (2)$$

where μ_{LIA} is the Lagrange multiplier on the liquidity-in-advance constraint, that is, the marginal benefit of relaxing the constraint. Note that even when the constraint is not binding, the positive likelihood that it will bind in the future influences behavior in the current period. The first Euler relates consumption and saving, and the specific form of the liquidity-in-advance constraint plays key role. The second Euler relates consumption and borrowing, as evidenced by the presence of the interest rate wedge between saving and borrowing, δ . It features the liquidity-in-advance constraint indirectly through the Lagrange multiplier.

Combining the Eulers yields a single expression that relates the expected marginal utility of consumption to the expected marginal benefit of relaxing the LIA constraint:

$$E_t u'(c_{t+1}) = E_t \mu_{LIA,t+1} + \frac{R}{\delta} E_t [\mu_{LIA,t+1} \cdot \omega_a(y_{t+1}, a_{t+1})]$$

If we further assume that the LIA constraint is linear in liquid assets (as it will be in the baseline calibration) and rewrite the derivative as a constant, $\omega_a \equiv \omega_a(y_{t+1}, a_{t+1})$, then we can further simplify this expression:

$$\frac{E_t u'(c_{t+1})}{E_t \mu_{LIA,t+1}} = 1 + \omega_a \frac{R}{\delta}$$

In the limit, when $\omega_a \rightarrow 0$, households choose saving and borrowing to equate the expected marginal benefits of consumption and relaxing the constraint. When $\omega_a > 0$, there exists an optimal wedge between the two expected marginal benefits, which also depends on the saving and borrowing interest rates. As ω_a increases, that is, the strength of the constraint increases, the wedge increases, implying that households demand less consumption (i.e., more marginal utility of consumption) to be made equally well off. This trade-off generates simultaneous demand for both saving and borrowing, which we describe as rational coholding.

4.3 Rational Coholding

In this section, we outline the model mechanisms that generate coholding of liquid assets and liquid debt. As discussed in Section 2.6, the household finance literature proposes and validates many mechanisms that generate coholding, many of which rely on bounded rationality or households making suboptimal decisions. In contrast, the liquidity-in-advance constraint generates rational coholding in the sense that despite the large wedge between saving and borrowing rates, fully rational and optimizing households opt to simultaneously hold liquid assets and debt.

Rational coholding arises from the composition of gross positions that balances desired net wealth, w^* , and desired liquid assets, a^* . The desired level of net wealth comes from the standard buffer-stock mechanisms common to all models in this class: the household wishes to hold precautionary saving to insure against negative income shocks and wishes to borrow against positive income shocks. This yields a level of desired wealth, w^* . In standard models, this is the only mechanism. In our model, the liquidity-in-advance constraint generates a desired level of liquid assets, a^* , driven by the (expected) level of consumption in the future and the exact specification of the constraint. Together, these two targets yield a desired level of liquid debt, $d^* = w^* + a^*$, which is satisfied subject to the borrowing constraint. Thus the household will rationally cohold due to the joint behavior stemming from buffer-stock behavior and liquidity-in-advance constraint.

4.4 Recursive Formulation and Computational Details

Stated recursively, the household’s problem is to find the policy functions for consumption, saving, and borrowing, given by $c(a, d, y)$, $a'(a, d, y)$, and $d'(a, d, y)$, respectively, to maximize its value function,

$$V(a, d, y) = \max_{c(\cdot), a'(\cdot), d'(\cdot)} u(c(a, d, y)) + \beta E[V(a'(a, d, y), d'(a, d, y), y')|y],$$

subject to the exogenous income process, detailed below, and the budget and liquidity-in-advance constraints. We solve the model using 100 gridpoints for assets and debt.

The marginal propensities to consume, save, and repay debt are the difference between the policy functions with the shock minus the policy functions without the shock, divided by the size of the shock. While technically the “marginal” propensity is defined in the limit as the size of the shock approaches zero, we follow the terminology in the literature of marginal propensities for shocks that are small relative to income.

Our empirical analysis guides the choice to not include a life-cycle or illiquid asset in the model. We do not find strong evidence for age patterns affecting marginal propensities or coholding status, and Telyukova (2013) provides further evidence for the latter. Furthermore, our empirical analysis suggests that it is primarily liquid wealth that matters for marginal propensities and not illiquid wealth. Our calibration strategy is consistent with this modeling choice.

5 Calibration and Model Performance

This section presents the model’s baseline calibration. We calibrate external parameters to standard values in the literature and target the distribution of wealth in the economy using the discount rate and liquidity-in-advance constraint. In the next section, we show that despite targeting wealth, the model does remarkably well in reproducing the empirical patterns for the marginal propensities to consume, save, and repay debt.

5.1 External Calibration: Income, Preferences, and Financial Markets

The model is calibrated to a monthly periodicity. Table 2 presents the calibrated parameters. We select a standard value of risk aversion, $\gamma = 2$. We take the interest rate on savings and the interest rate spread on credit card debt from Telyukova (2013) and set it to $r = 0.0033$ and $\delta = 0.0074$, respectively. This corresponds to an annual interest rate on saving of 4% and an annual interest rate on credit card debt of 14%.

Borrowing is allowed up to approximately two months of average monthly income, $\phi = 2.2$, in line with the analysis by Kaplan and Violante (2014) who find a limit of 74% of quarterly income. We take the income process from Gelman (2021), who estimates an AR-1 income process using financial accounts data. This yields a persistence parameter $\rho = 0.096$ and variance of the innovation $\sigma_y^2 = 0.039$. This process is meant to capture transitory variations in income and abstract from permanent differences across households.

Table 2: Baseline External Calibration

Parameter	Description	Value	Source
γ	Risk aversion	2	Standard
r	Interest rate	0.0033	4.00% APR
δ	Credit card spread	0.0074	9.63% APR
ϕ	Borrowing limit	2.2	74% of quarterly income
ρ_y	Persistence of y_t	0.096	Gelman (2021)
σ_y^2	Variance of innovation in y_t	0.039	Gelman (2021)

5.2 Internal Calibration: Liquidity-in-Advance Constraint and Discount Rate

We calibrate the liquidity-in-advance constraint and the discount rate to match the distribution of liquid assets and debt observed in the data. As above, liquid assets in the data are defined as the sum of checking and savings accounts plus idle money in brokerage accounts. Liquid debt is defined as credit card debt, measured by the balance due after the last statement was paid. We normalize by monthly income to bring the scaling in line with our model.

5.2.1 Data and Model Targets

We target distributional moments from the SCF instead of the SCE. This is because the SCF provides a more accurate picture of *revolving* credit card debt, whereas the SCE only asks about the current stock of total credit card debt.⁵ Note that key moments of our analysis are predominantly untargeted: the share of coholders; levels of marginal propensities to consume, save, and repay debt; and the slopes of marginal propensities across the

⁵Appendix A.6 compares household characteristics across SCE and SCF and shows that coholding shares are similar. Targeting moments from the SCE instead of the SCF does not materially affect our conclusions.

distribution of liquid wealth. We compute marginal propensities by simulating quarterly responses to an unanticipated change in average monthly income of roughly 10%.⁶

We target the median level of liquid assets to monthly income, which is 0.61. A cursory glance at the data, shown in Figure 1a, reveals an important fact not captured by our model: 30% of households have no credit card, and another 38% report having a credit card but using it only for convenience. Thus, the median level of credit card debt in the data is zero. In the model, it would be impossible to generate such low levels of median credit card debt holding without imposing additional costs on the extensive margin of opening a credit card or additional benefits from being a convenience user (such as reward points or multi-account incentives). Instead, we target the 75th percentile of liquid debt, 0.41, and report the median level of debt in the data and model as an untargeted moment.

5.2.2 Parameterization

Discount Rate The (annual) discount rate is set to 0.90943. We note that our model does not require an extremely low discount rate to induce impatience and therefore borrowing. This is because in our case, households borrow not only to smooth consumption across time, but because within a given period, they wish to both satisfy the future liquidity-in-advance constraint and finance current consumption. This generates demand for debt without the need for impatience.

Liquidity-in-Advance Constraint We model the liquidity-in-advance constraint to dictate that the household must pay for a fraction, $\theta \in (0, 1]$, of consumption using its stock of liquid assets in the current period, a_t :

$$\theta c_t \leq a_t.$$

In the baseline calibration, we set $\theta = 0.50$; that is, half of consumption must be paid for using liquid assets, which is in the lower range of empirical estimates of the percentage of consumption paid for using non-credit products. In the 2016 Survey of Consumer Payments, households reported that they pay for 58% of transactions using cash or debit, another 10% using direct transfers from their bank accounts, and only 18% using credit cards. This is consistent with earlier work by Telyukova (2013) that finds cash payments accounted for 65% of the total value of all consumer transactions in the 2002 Survey of Consumer Expenditure. Over time, there has been a marked shift from cash to digital

⁶Note that we compute MPCs out of lump-sum changes while the SCE asks about proportional changes. Recomputing MPCs out of proportional changes yields similar results.

payments, but the majority of that transition has been to digital debit payments. Greene and Stavins (2022) document that over 80% of households’ regular consumption goods and services, such as shelter and utilities, are paid for using liquid assets directly from bank accounts, and that many households hold exactly for this purpose.

This specification yields the following relationship in the first-order conditions between the expected marginal utility of future consumption and expected marginal benefit of relaxing the liquidity-in-advance constraint:

$$\frac{E_t u'(c_{t+1})}{E_t \mu_{LIA,t+1}} = 1 + \frac{R}{\theta \delta}$$

Since $\theta \in (0, 1]$, the constraint effectively decreases the difference between the saving and borrowing rates, δ . In other words, the larger is the required level of liquidity for financing consumption, the less is the effective premium to borrow liquid debt while holding liquid assets.

We note that this specification does not include income, y_t , in the measure of liquidity. This is meant to tighten the interpretation of the constraint on the household’s endogenous choice for liquid assets. In this setup, the household must “convert” income in the current period into liquidity in the next period before it can be used to loosen the constraint. This might reflect the increasingly uncommon case in which households receive their income via check and need to convert it into cash with some short delay. Because this technical feature of the model may distort the contemporaneous consumption response to a positive income shock for households bound by the liquidity-in-advance constraint, we compute marginal propensities out of shocks to assets instead, as is standard in one-asset models. For the sake of exposition, we will nevertheless refer to them as marginal propensities out of income shocks.⁷ Appendix A.7 presents an alternative calibration of the model in which current income is allowed to enter the liquidity-in-advance constraint. The inclusion of income worsens the model fit on the asset side, but preserves the fit with regards to the coholder share and marginal propensities.

5.3 Analysis of Model Performance

As our primary calibration targets, the model matches the 75th percentile of liquid debt holdings and median liquid asset holdings exactly. Both in the data and the model, a


⁷Irrespective of that distinction, our focus on the quarterly response corrects for this slight distortion since by then households will have been able to convert income into the desired levels of liquidity and consumption. To illustrate this point, we explicitly analyze in Appendix A.8 the impact of income shocks by introducing a fourth state variable, Δ , that represents a one-period income shock and enters directly into the budget constraint. We show that the MPC out of income is similar to the MPC out of assets.

household at the 75th percentile of the distribution holds debt equal to 41% of monthly income. Along the asset distribution, the median household holds liquid assets worth 61% of monthly income. Despite its simplicity, the model also matches untargeted distributional moments fairly well. Consistent with the data, the median household holds a negligible level of liquid debt at 7% of monthly income. The model slightly overstates the median holdings of liquid wealth with 48% of monthly income, compared to 31% in the data.

The model also matches the joint distribution of liquid assets and debt reasonably well. Along the external margin of coholding, roughly 46% of households hold simultaneously liquid assets and debt, compared to 27% in the data. This difference is partly a result of applying a stricter definition of coholding in the data. While in both model and data, coholders need to hold more than 10% of monthly income in liquid assets and debt, in the data they need to additionally report to revolve debt habitually. This restriction can reduce the coholder share by as much as 9 pp, as illustrated in Figure A.1. Along the internal margin, the model performs better and predicts a liquid debt-to-asset ratio for a household at the 75th percentile of the distribution of 0.81 compared to 0.80 in the data. We again focus on the 75th percentile because the median household does not hold liquid debt. As in Figure 4, Figure 10 plots the joint distribution of liquid assets and debt over the distribution of liquid wealth. The model generates a similarly L-shaped distribution as in the data. Households in the lowest tercile of liquid wealth have low cash holdings and high debt holdings, while households in the highest tercile of liquid wealth have high cash holdings and no debt. Households in the middle tercile hold less liquid debt than the lowest tercile and less liquid assets than households in the top tercile.

This figure highlights a shortcoming of the model relative to the data: while we observe coholding for even very high levels of liquid wealth, there is no coholding in the model for these households. Specifically, in Figure 4, liquid debt is positive even for the highest tercile households, while in the model counterpart, liquid debt is zero for the highest tercile. This is because the rational coholding generated by the liquidity-in-advance constraint is not necessary for high-wealth households that, by virtue of positive income shocks, hold significant savings in liquid savings and therefore have no need to borrow. Generating coholding for these households would require modeling one of the other (perhaps non-rational) mechanisms discussed in Section 2.6. As we demonstrate in the next section, however, there are relatively few of these high-wealth coholders, and our main analysis is unaffected by missing them in the model.

Figure 10: Joint Distribution of Liquid Assets and Debt



Notes: Joint distribution of liquid assets and debt in baseline model specification.

6 Coholding and Marginal Propensities in the Model

With the calibrated model at hand, we turn to studying household behavior of coholders and across the full distributions of liquid assets and debt. We compare simulated data from the model against the survey data analyzed in Section 2. To provide a direct comparison, we simulate a one-time positive income shock for 30,000 households and regenerate the empirical figures from above on the simulated data instead of the SCF/SCE data. Along several dimensions, the model generates household behavior consistent with the observed empirical behavior.

6.1 Distributions of Marginal Propensities

Panel (c) of Table 3 shows the model's performance in matching several untargeted moments. The model performs well in matching the average propensities to consume, save, and repay debt. The average household in the model consumes around 16% of an unexpected income windfall, uses 43% to repay debt, and saves the remainder. The model

Table 3: Internal Calibration and Model Moments

Panel A: Internally Calibrated Parameters

Parameter	Description	Value
β	Annual discount factor	0.90943
θ	Share of liquid consumption	0.50

Panel B: Targeted Moments

	Data	Model
Liquid Assets (Median)	0.61	0.61
Liquid Debt (75 th Pct.)	0.41	0.41

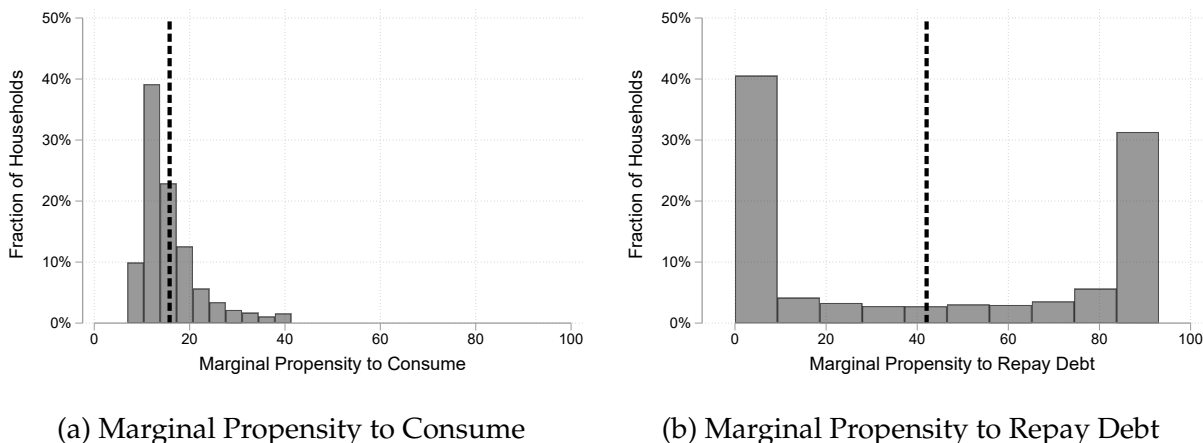
Panel C: Untargeted Moments

	Data	Model
<i>Distribution of assets, debt and net wealth</i>		
Liquid Assets (Mean)	2.04	0.73
Liquid Debt (Mean)	0.50	0.26
Liquid Debt (Median)	0.00	0.07
Liquid Wealth (Mean)	1.54	0.47
Liquid Wealth (Median)	0.31	0.48
<i>Joint distribution of assets and debt</i>		
Share of co-holders (%)	26.6	46.3
Liquid debt-to-asset ratio (75 th Pct.)	0.80	0.81
<i>Marginal propensities</i>		
Avg. MPC	16.8	15.9
Avg. MPRD	39.2	42.5
Avg. MPS	44.0	42.9

Notes: Liquid assets and debt are expressed relative to monthly income. Liquid assets are defined as the sum of checking and savings accounts plus idle money in brokerage accounts in the data. Liquid debt is defined as credit card debt, measured by the balance due after the last statement is paid. Households are coholders if they hold more than 10% of monthly income in liquid assets and debt and additionally report revolving credit card debt habitually. Balance sheet data are taken from the SCF 2016, while marginal propensities are taken from the SCE 2015–2018.

matches the average MPRD particularly well, which is 39% in the data. The average MPC generated by the model is slightly lower than in the data, but close to the one generated by a conventional two-asset model with a liquid and illiquid asset (Kaplan and Violante, 2014). This is particularly noteworthy as our model does not require the presence of wealthy but illiquid households—commonly referred to as wealthy hand-to-mouth—to generate high MPCs.

Figure 11: Marginal Propensities to Consume and Repay Debt



Notes: Marginal propensities from the model. Top and bottom percent of marginal propensities are win-sorized to avoid outliers caused by numerical error.


Figure 11 plots the distributions of propensities to consume and repay debt in the model. Relative to the data, the model’s propensities to consume have the correct right-skewness but no mass at zero, which can be generated in models with fixed costs to responding to income shocks (Fuster et al., 2021; Boutros, 2023). The model correctly generates the empirically observed bi-modality (Panel (b) of Figure 6). Most households either do not repay any debt at all in response to an income windfall, potentially because they do not hold any debt, or they use almost all of it to repay debt, with few households in-between.

6.2 Household Behavior and Net Liquid Wealth

Figure 12 plots the marginal propensities to consume, save, and repay debt along the distribution of liquid wealth in the model. As in the data, the propensity to consume is essentially flat over the distribution of liquid wealth. The propensity to save is increasing in liquid wealth, with the entire change coming from a commensurate decrease in the propensity to repay debt.

This stands in stark contrast to the predictions of conventional models, which assign high MPCs to low-wealth or liquidity-constrained households. In our model, a large share of low-wealth households holds substantial amounts of debt, reducing their consumption response in favor of higher debt repayments. Conventional models, instead, classify households as either savers or debtors, thereby abstracting from coholding and masking substantial differences across allegedly hand-to-mouth households.

Figure 12: Marginal Propensities Across the Distributions of Balance Sheet Items




Notes: Marginal propensities from the model. Top and bottom percent of marginal propensities are winsorized to avoid outliers caused by numerical error.

6.3 Household Behavior and Gross Liquid Wealth

Turning to the distributions of liquid assets and debt, Figure 13 plots the marginal propensities to consume, save, and repay debt over each marginal distribution. These figures closely resemble their empirical counterparts. Again, behavior across the distribution of liquid assets is nearly identical to behavior across the distribution of liquid wealth. The propensities to save and repay debt are increasing and decreasing, respectively, in liquid assets, while the consumption response is essentially flat. For liquid debt, the opposite story holds, with the propensity to repay debt sharply increasing in debt.

Figure 13 plots average propensities by asset and debt quartile. The empirical version of this figure is crucial in illustrating the importance of considering the joint distribution of assets and debt. Again, the model produces behavior strongly consistent with the empirical patterns. Abstracting for a moment from the debt dimension, both data and model show that MPCs are relatively stable across the distribution of assets, while the MPS is increasing and the MPRD is decreasing. Focusing on the distribution of debt for a given quartile of assets, we see that across most levels of assets, the MPC is decreasing in debt, while the MPRD is increasing. This stands in contrast to what conventional models would predict. More debt, holding assets constant, implies lower wealth, and thus a higher MPC. In our model, however, households have a strong desire to repay debt due to high interest rates. As noted above, a limitation of the model is that it does not match

Figure 13: Marginal Propensities Across the Distributions of Balance Sheet Items



Notes: Marginal propensities from the model. Top and bottom percent of marginal propensities are win-sorized to avoid outliers caused by numerical error.


the right tail of coholding, that is, households that are in the highest quartile of both liquid assets and liquid debt. Through the lens of the model, this degree of coholding cannot be justified even with the desire for liquidity-in-advance, which results in the “missing” bars on the right-hand side of Figure 13. This figure highlights that the model misses some of these households, but they account for only a small fraction of the total distribution.

6.4 Regression Analysis

Finally, we again confirm the visual findings in the last several sections using linear regressions. Table 4 reports the estimated coefficients and, compared to the data, eight of the nine slopes are the same sign (and all are statistically significant). The exception is the estimated relation between liquid assets and the propensity to consume, which is -0.351 in the model and slightly positive but statistically insignificant in the data. The model preserves the traditional role of liquidity as insurance against temporary income changes, and higher liquid assets imply more consumption smoothing.

Otherwise, the model captures well the relations between marginal propensities and balance sheet items. With regards to marginal propensities to save and repay debt, the model agrees with the data on the sign of the relation with liquid assets and debt, but not on the magnitude. This is explained by the fact that the model does not match well the joint distribution of assets and debt in *levels*, in particular the respective right tails of the distributions. However, note that the model captures the idea that liquid assets and debt explain relatively more of the variation in the MPS and MPRD than in the MPC, as indicated by the higher R-squared.

Figure 14: Marginal Propensities Across the Joint Distribution of Assets and Debt



Notes: Marginal propensities from the model. Top and bottom percent of marginal propensities are winsorized to avoid outliers caused by numerical error.

Table 4: Regressions of Marginal Propensities on Household Liquid Balance Sheet: Model

	(1) Spend	(2) Spend	(3) Save	(4) Save	(5) Repay Debt	(6) Repay Debt
Liquid Wealth	0.005* (0.002)		1.993*** (0.009)		-2.025*** (0.010)	
Liquid Assets		-0.351*** (0.004)		2.186*** (0.023)		-1.857*** (0.022)
Liquid Debt		-0.315*** (0.005)		-1.826*** (0.019)		2.171*** (0.022)
N	30,000	30,000	30,000	30,000	30,000	30,000
R ²	0.000	0.212	0.584	0.585	0.566	0.567

Notes: Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Balance sheet variables in the model are rescaled to match the average debt level in data. Top and bottom percent of marginal propensities in the model are winsorized to avoid outliers caused by numerical error.

7 The Macroeconomic Implications of Coholding

With the calibrated model in hand, we explore the macroeconomic implications of coholding through the lens of the aggregate household consumption, saving, and borrowing functions, which are composed of the household-level functions integrated over the distribution of households. We begin by studying stimulative fiscal transfers such as in 2001, 2008, and 2020. We show that for these completely untargeted transfers, the documented changes in the joint distributions of liquid assets and debt have stark implications for the aggregate impact of these programs. We then demonstrate that to generate the largest immediate consumption impact, fiscal stimulus should be targeted to true hand-to-mouth households with little credit card debt since they have the largest propensities to consume.

We then turn to the implications of coholding for general equilibrium macroeconomic models. Instead of taking a stance on a particular model, we characterize the model's intertemporal marginal propensities to consume, save, and repay debt, which are near-sufficient statistics for the household block in modern Heterogeneous Agent New Keynesian (HANK) models (Auclert et al., 2023).

7.1 Aggregate Consumption, Saving, and Borrowing Functions

The aggregate consumption function at a given point in time integrates each household's consumption behavior, $c_{i,t}$, over the entire distribution of households, $\lambda_t(i)$. In reality, the distribution of households varies along an infinite number of characteristics, but classic models in macroeconomics summarize behavior based on net wealth. In our model, the aggregate consumption function depends on assets, debt, and income:

$$C_t = \int_{(a,d,y)} c_t(a, d, y) \lambda_t(a, d, y) d(a, d, y).$$

The analysis in previous sections demonstrates that the model's marginal propensities to consume, save, and repay debt reflect behavior across the distribution of households, and our model calibration targets moments from the distributions of assets and debt. Altogether, then, we are able to combine these two components to generate an accurate aggregate consumption function.

The aggregate saving and borrowing functions, S_t and B_t , are similarly defined by integrating the household-level functions over the distributions of gross wealth. These will play an important role in our analysis, especially when we study aggregate implications of coholding in general equilibrium.

7.2 Stimulative Fiscal Transfers

Stimulative fiscal transfers are an important policy tool employed to increase aggregate demand in times of economic downturn. These policies have been popular in the United States in the last 25 years in different forms, from rebates to stimulate the specific sector (e.g., automobile rebates) to direct cash transfers in 2001, 2008, and 2020–2021. In this section, we show that the distribution of coholding over time affects the impact of stimulative fiscal transfers in two ways. First, for untargeted fiscal transfers, the unconditional distribution of wealth in the economy matters because as coholding increases in the general population, the aggregate marginal propensity to consume decreases. Second, we show that targeting based on “gross liquid wealth,” that is, the sum of gross positions, is more effective for stimulating consumption since this measure more accurately identifies true hand-to-mouth households than targeting based on income or net wealth.


7.2.1 Untargeted Fiscal Transfers

The starting point for our analysis is an untargeted fiscal stimulus program in which all households receive an identical and unexpected one-time income shock. The aggregate propensities to consume, save, and repay debt are exactly the averages reported in Table 3 when evaluating the model’s calibration. These represent the integration of the household-level consumption function over the joint distributions of liquid assets and debt calibrated to the 2016 Survey of Consumer Finances.

As we noted in Section 2, the distribution of coholding has changed significantly over time. In this section, we use the model to study how the aggregate consumption response to an untargeted transfer program has changed over time due solely to the composition of wealth in the economy. We integrate the model’s household-level consumption function over the distribution of liquid assets and debt from the 2001–2016 vintages of the Survey of Consumer Finance. For each year, we divide the data into quintiles of liquid wealth and calculate the average level of liquid assets and debt in each group. We then calculate the propensity to consume for each quintile using the model’s consumption policy function, and combine each quintile’s consumption response to arrive at the aggregate marginal propensity to consume.

Figure 15 plots this aggregate consumption response, relative to the baseline calibration in 2016, across time. The aggregate propensity to consume is larger in every year prior to 2016. This striking result arises from the fact that coholding has increased over time and, based on our previous analysis, coholding households typically have smaller consumption responses and larger debt repayment responses. For example, in 2010, the

Figure 15: Relative Model-Based Aggregate Marginal Propensity to Consume (2001–2016)



Notes: This figure plots the aggregate consumption response to an untargeted fiscal transfer. The consumption response in 2016 is normalized to one, and each bar represents the relative response using the distribution of households from each vintage of the Survey of Consumer Finances between 2001 and 2013.

aggregate consumption response to an untargeted fiscal transfer program is 16.5% larger than in 2016, while in 2013, the aggregate response is 12.2% larger than in 2016. This downward trend over time in the aggregate MPC out of an untargeted fiscal stimulus program is driven by the corresponding upward trend in coholding on both the intensive and extensive margins (see Section 2). The model captures that households with more liquid debt have smaller consumption responses, and thus as the number of indebted households in the SCF grows, the aggregate consumption response decreases.

7.2.2 Targeted Fiscal Transfers

The previous section focuses on the aggregate consumption response using the distribution of households over time. We next study the aggregate consumption response if the fiscal authority is able to target specific segments of the distribution of households. In particular, we use the baseline model, which is calibrated to the distribution of households in 2016, to study targeted fiscal transfers to three groups of households sorted on three characteristics: the bottom 10%, bottom 30%, and bottom 50% of households along the distributions of income, net wealth, and gross wealth (defined as the sum of liquid assets and liquid debt). To facilitate comparison across programs, we hold the aggregate transfer size fixed across all programs and vary the size of each household's transfer. Ta-

ble 5 reports the percentage change in aggregate consumption and debt as a fraction of the aggregate fiscal transfer size for each characteristic and group.

Table 5: Consumption and Debt Response to Fiscal Transfers

	Net Wealth		Gross Wealth		Income	
	Cons.	Debt	Cons.	Debt	Cons.	Debt
Bottom 10%	11.7	-88.3	19.2	-29.1	14.8	-54.8
Bottom 30%	13.3	-81.7	19.7	-37.2	15.0	-54.3
Bottom 50%	16.9	-70.8	19.9	-39.4	15.3	-51.9

Notes: This table reports the response of aggregate consumption and debt to targeted fiscal transfers as a percentage of the aggregate transfer size. The aggregate transfer size is held fixed across scenarios. Transfers are lump-sum and amount to 10% of average monthly income for the scenario in which the bottom 30% of the distribution are targeted. In the benchmark untargeted transfer, the changes are 17.2% for consumption and -43.4% for debt.

Net Wealth From the first row and first two columns of Table 5, targeting the bottom 10% of the net wealth distribution increases consumption by 11.7% of the total fiscal transfer and decreases debt by 88.3% of the total transfer. Increasing the group of households that receives the transfer to the bottom 50% of the distribution significantly increases the consumption response to 16.9% and decreases the debt response (in absolute terms) to 70.8%. Contrary to traditional wisdom, targeting households with *more* net wealth leads to a *larger* consumption response. This occurs because the poorest households have the most credit card debt and therefore a smaller consumption response than households with more net wealth due to less credit card debt. Targeting wealthier households (that are still in the lower half of the distribution) increases the consumption response because these households are true hand-to-mouth with low liquid wealth due to low assets and low debt.

Gross Wealth Building on this insight, we construct a new measure, “gross wealth,” which sums liquid assets and liquid debt. Households with low gross wealth have low gross positions and by construction have low net positions, whereas households with low net positions may have large gross positions. As such, only true hand-to-mouth households can have low gross wealth. Indeed, targeting the bottom 10% of households based on gross wealth, the aggregate consumption response is 19.2% of the total transfer and the aggregate debt response is -29.1%. Increasing the scope of the program to the bottom 50% of households only marginally increases the aggregate consumption response

to 19.9%, indicating that the consumption response is similar for all households along the lower part of the gross wealth distribution.

Further, the debt response increases, but only to 39.4%, much less than under any of the net wealth targeting programs. This indicates that the gross wealth targeting program is also distributing transfers to households with moderate gross wealth positions due to moderate liquid assets and no liquid debt. These households have positive net wealth positions. Although they may not strictly be hand-to-mouth, they still have larger consumption responses than the coholding households with small consumption responses that would be mislabeled as hand-to-mouth due to their large negative positions from holding credit card debt.

Income Finally, we consider income-based targeting in the last two columns of Table 5. Targeting the bottom 10% of the income distribution increases consumption by 14.8% of the total fiscal transfer and decreases debt by 54.8%. As the transfer targets more households in the bottom 50% of the population, the aggregate consumption response marginally increases to 15.3% and the aggregate debt response marginally decreases to 51.9%. Targeting higher-income households increases the aggregate consumption response because higher-income households typically hold less credit card debt and therefore have larger propensities to consume (and save).

Compared to targeting net wealth, targeting the bottom 10% of income generates a larger consumption response because the bottom 10% of net wealth are heavy credit card users. In contrast, if targeting the bottom 30%, sorting based on net wealth generates a larger consumption response because it captures a larger share of true hand-to-mouth households. However, compared to targeting gross wealth, the average consumption response when targeting income is uniformly smaller because lower-income households typically have more credit card debt and therefore a larger average debt response. These results show that while income-targeting may yield similar or better consumption responses relative to net wealth targeting, in our experiments targeting gross wealth is ultimately superior to both other measures in generating the largest aggregate consumption response.

Comparison to Untargeted Transfer Connecting to the analysis in the previous section, Figure 16 plots the change in aggregate consumption when targeting the bottom 30% of each group relative to an equally-sized program that distributes a transfer to every household in the economy. In the untargeted plan, the change in consumption is 17.2% of the total program size and the change in debt is -43.4%. Relative to the untargeted program, the consumption response in the program targeting the bottom 30% of net wealth is 22.7% smaller, and the consumption response in the similar income-based program

Figure 16: Consumption Response to Targeted Relative to Untargeted Fiscal Transfers



Notes: This figure reports the response of aggregate consumption to fiscal transfers targeting the bottom 30% of the distribution relative to untargeted transfers holding the aggregate transfer size fixed. Transfers are lump-sum and amount to 10% of average monthly income for the targeted scenario.

is 12.8% smaller. These programs generate less aggregate changes than from the untargeted program because credit card borrowing is pervasive across the lower distributions of income and net wealth, driving down the consumption response out of fiscal stimulus payments. In contrast, the gross wealth targeting program generates an aggregate consumption response 14.5% larger than the untargeted program. This is accomplished by directing stimulus towards true hand-to-mouth households and away from coholders.

7.3 The Dynamics of Stimulative Fiscal Transfers

All of the analysis in the previous section focuses on the immediate impact of fiscal transfers. The model we develop also allows us to study the dynamics of consumption over time. In this class of models, since households ultimately only value consumption, the cumulative consumption response to any income shock eventually reaches 100%, and can surpass this if the non-consumed portion of the shock is invested into a financial instrument with positive returns.

Untargeted Fiscal Transfers Panel (a) of Figure 17 plots the cumulative consumption and debt repayment responses of the same untargeted fiscal transfer studied in previous sections. As above, the immediate consumption response is 17.2% of the total transfer, while 43.3% of the transfer is used for debt repayment. After one year, the cumulative

consumption response is only 50%, implying that in quarters two to four, a significant amount of the fiscal transfer remains in the form of wealth. Finally, after almost four years, the cumulative consumption response reaches 100% of the initial cash transfer.

Over that time, the non-consumed wealth portion of the transfer accumulates returns, both directly through saving in the liquid asset and indirectly through the savings from smaller debt service payments. Cumulative debt repayment and saving eventually become zero, implying that average household balance sheets return to their pre-transfer levels. However, because of the large front-loading of wealth increases, the cumulative consumption response grows beyond 100%, ultimately reaching 115.2% of the initial cash transfer. We note that this number must be interpreted with caution because our analysis is undertaken in partial equilibrium, and the effects of general equilibrium will be especially important over time.


Targeted Fiscal Transfers We perform a similar dynamic analysis for each of the targeted fiscal transfers studied in the previous section. Panel (b) of Figure 17 plots the cumulative response for each program relative to the untargeted program in Panel (a). Consistent with the analysis above, the immediate consumption response in the income and net wealth programs is well below the untargeted program, while only the gross wealth targeting program has a larger immediate consumption response than the untargeted program.

Over time, however, both the income and gross wealth targeting programs converge to roughly the same long-run cumulative response as the untargeted program. With gross wealth targeting, the primary beneficiaries are true hand-to-mouth households, which have the largest spending responses and smallest saving responses. As such, the response is completely front-loaded relative to the untargeted program. Since more of the transfers are immediately spent rather than saved, there are less cumulative returns to the non-consumed wealth, and thus the cumulative response in the long run is slightly less than in the untargeted case.

On the other hand, the income targeting program directs a significant sum of fiscal transfers to coholders, who pay down debt and have a smaller initial consumption response. Eventually, the savings from having less debt become increased consumption, and most of the response is back-loaded until the cumulative responses eventually equalize after four years. In the long run, the cumulative response from the income targeting program is slightly greater than the untargeted program.

Since the net wealth program almost directly targets coholders with the largest levels of credit card debt, this program has an even smaller immediate consumption response than the income targeting program in every period. As with the other programs, the net

Figure 17: Dynamic Aggregate Responses to Fiscal Transfers



Notes: The left panel reports the dynamic response of aggregate consumption and debt repayment to untargeted fiscal transfers as a percentage of the transfer size. The right panel reports the dynamic response of aggregate consumption to targeted fiscal transfers relative to the untargeted scenario. In each scenario, the bottom 30% of the respective distribution (income, net wealth or gross wealth) are targeted. The aggregate transfer size is held fixed across scenarios. Transfers are lump-sum and amount to ten percent of average monthly income for the targeted scenario.

wealth program equalizes with the untargeted program after roughly four years. However, instead of remaining roughly in line with the untargeted program, the cumulative response in the net wealth program continues increasing, ultimately growing roughly 5% more than any other program. This occurs because coholders immediately pay down their credit card debt instead of increasing consumption, and the cumulative wealth returns from avoiding high-cost debt payments yields increased consumption in the long run.

Discussion and Implications Our analysis highlights that each fiscal transfer program generates unique short- and long-run dynamics for aggregate household consumption. Depending on the policymaker’s objectives, each program may be optimal and thus comparison across programs necessitates specifying both the desired outcome and horizon. If the goal is to immediately stimulate demand, our analysis suggests targeting away from coholders. If the goal is to increase consumption over the medium- or long-term, then targeting coholders may be optimal since these households immediately increase wealth through decreased debt, which eventually translates into higher consumption. Again, although the analysis in this section is in partial equilibrium, we believe that the main insights still apply.

7.4 Debt Relief as Fiscal Policy


In addition to stimulative cash transfers, fiscal authorities have also experimented with various debt-relief policies during economic downturns. In the United States, the federal government has implemented debt relief for both mortgages in 2008 and student debt in 2020 (Ganong and Noel, 2020; Dinerstein et al., Forthcoming), which researchers have argued may be too broad or expensive relative to other alternatives (Catherine and Yanelis, 2023; Boutros et al., 2023). In Canada, the federal government has implemented debt relief for credit cards in the form of payment pauses (Allen et al., 2022). These recent advances into consumer credit markets demonstrate a willingness for governments to implicitly target borrowers instead of the broad population.

Our model allows us to study and compare debt relief policies relative to direct cash transfers. Panel (a) of Figure 18 plots the aggregate consumption and debt responses to a debt relief program that forgives outstanding credit card debt up to 10% of monthly income. Households with no outstanding credit card debt are completely unaffected by the program, while households with debt equal to less than 10% of monthly income see their entire balance forgiven. The debt relief program increases aggregate consumption by 13.2% of the total amount of debt forgiven. After almost 4.5 years, the aggregate consumption response is 100%, and eventually reaches 118.7% in the long run. On the other hand, aggregate debt decreases by 69.4% of the total debt relief package. This is less than 100% because in response to debt forgiveness, many households optimally choose to re-borrow a fraction of that debt in the next period. As with the direct fiscal transfer program, the cumulative aggregate debt response is zero and long-run balance sheets return to their pre-program levels.

Comparison to Fiscal Transfers Despite a smaller initial response, the long-run aggregate consumption response to debt relief is 3.5 pp greater than the aggregate response from the direct transfer program. This difference arises for two reasons that our model can help us understand. First, the MPC out of income is different than the MPC out of debt relief; and second, debt relief implicitly targets only households that hold debt, which are a subset of the total population.

Panel (b) of Figure 18 plots two alternative fiscal programs that shed light on how each of these reasons contributes to the difference between direct cash transfers and debt relief. First, we consider an alternative “Targeted Fiscal Transfer” program that targets each household that receives debt relief and, instead of forgiving debt up to 10% of income, gives them the same amount as cash. This allows us to isolate the difference between the type of positive shock and the corresponding consumption response. On impact, the

Figure 18: Dynamic Aggregate Responses to Debt Relief



Notes: The left panel reports the dynamic response of aggregate consumption to an untargeted debt relief of 10% of average monthly income as a percentage of the total debt relief. The right panel reports the response relative to a uniform transfer (untargeted transfers to all households) and a transfer replicating debt relief (targeted cash transfer equal in magnitude to debt relief).

aggregate consumption response of the cash program is 25% larger than the debt-relief program, but after four years, the cumulative consumption response for both programs is roughly equal. When debt is outright forgiven, households primarily respond by further decreasing debt, not by increasing consumption, but this ultimately leads to the same long-run increase in consumption.

Since debt relief only reaches households that hold debt, the total fiscal cost of the program studied in Panel (a) is far less than the untargeted transfer studied in the previous section. We therefore consider an alternative “Untargeted Fiscal Transfer” program in which all households receive a cash transfer, but the aggregate size of this program is equal to the size of the debt-relief program. Again, the immediate consumption response of this program is greater than the debt-relief program, but the long-run cumulative consumption response of the untargeted program is roughly 3.5% less than the debt-relief program. In the debt-relief program, heavily leveraged households see larger increases in wealth that they convert into larger increases in long-run consumption. This effect is washed out when debt forgiveness is replaced by uniform transfers for all households.

7.5 Fiscal and Monetary Policy in General Equilibrium

Previous sections analyze fiscal and monetary policy in isolation, holding constant other relevant variables determined in general equilibrium. We now turn to analyzing coholding in general equilibrium New Keynesian models. Relative to seminal New Keynesian

models that feature a single representative household, our model features a distribution of households with heterogeneous liquid asset and debt households. However, what ultimately matters is the aggregate consumption, saving, and borrowing response of all households, which are exactly the functions characterized above.

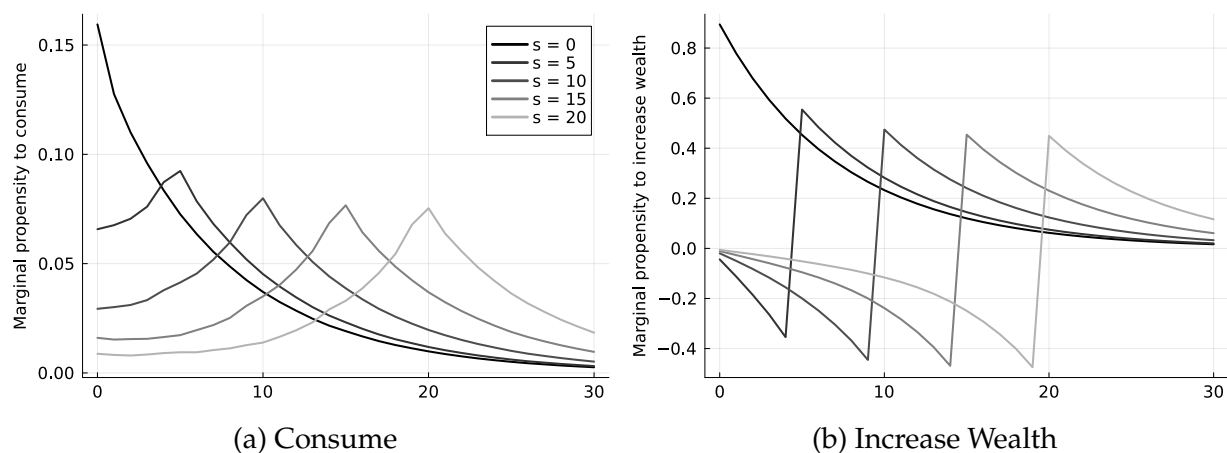
In particular, to study the impact of aggregate shocks in a first-order linear approximation around the steady state, Auclert et al. (2023) demonstrate that the matrix of intertemporal marginal propensities to consume (iMPCs) is nearly sufficient to characterize the impact of household heterogeneity on macroeconomic aggregates.⁸ The elements of the iMPC matrix are the derivatives of the aggregate consumption response function in each period in response to a shock that realizes in times $t = 0, 1, 2, \dots$. In particular, element (i, j) is the change in aggregate consumption at time i from an infinitesimal change in income at period j . The partial equilibrium analysis above, and the classic Keynesian cross, focus on the first column of this matrix, which studies the change in aggregate consumption in periods $t = 0, 1, 2, \dots$ from a shock that occurs in period $t = 0$.

As Auclert et al. (2023) show, the remaining columns of this matrix are necessary to study general equilibrium in an economy with dynamic households, firms, and policy-makers. For example, the second column is the response of aggregate consumption to a shock in period 2. In the first row is the change in consumption from news of the shock, and from the second period onward is the impact of the realized shock. In the third column, the first two rows are the impact of news of a shock in period 3, and so on.

Panel (a) of Figure 19 plots several columns of the iMPC matrix for the baseline model calibrated above. The solid black line is the first column of the matrix for a shock $s = 0$ periods in the future, that is, a contemporaneous income shock. This line simply plots the aggregate marginal propensity to consume over time. As is standard, the propensity to consume peaks on impact of the shock and decreases over time. Each subsequent line is marginally lighter and plots the responses over time to shocks in future periods, $s \in \{5, 10, 15, 20\}$. Each of these lines takes on a “tent” shape: the consumption response in anticipation of the shock is increasing, peaks in the period of the shock, and then decreases towards zero. In Panel (b), we plot the residual matrix of the intertemporal marginal propensity to increase wealth, iMPIW. As implied by the shape of the iMPC matrix, households decrease wealth to increase consumption in anticipation of the shock,

⁸This matrix is necessary and sufficient to characterize household behavior in a class of models detailed by Auclert et al. (2023). This class of models features constant interest rates and is typically used to study fiscal policy. In standard models that study monetary policy, the iMPC matrix is necessary but not sufficient to characterize household behavior. However, Auclert et al. (2023) show that the iMPC matrix is still nearly sufficient, in the sense that while other aspects of household behavior are also important (such as the propensity to consume out of asset returns), the primary determinant of aggregate behavior is the iMPC matrix.

Figure 19: Intertemporal Marginal Propensities to Consume and Increase Wealth




Notes: The left panel plots the aggregate marginal propensity to consume out of a positive income shock. Each line represents the response to a fully anticipated shock at period s relative to the initial period, $t = 0$. The darkest line, $s = 0$, is the response over time to a shock at period 0. The lightest line, $s = 25$, is the response over time to a shock at period 25. Each line represents the corresponding column of the iMPC matrix detailed in Section 7.5. The right panel plots the residual function, the propensity to increase wealth out of a positive income shock.

and then decrease wealth in the period of the shock and beyond.

Overall, the dynamics embedded in the iMPC matrix from the baseline model with coholding are remarkably similar to the iMPC matrices in Auclert et al. (2023), including the matrix from a standard two-asset model calibrated to match the average MPC observed in the data. The main difference is that the average level of the consumption response is lower in the model with coholding, which directly follows the earlier analysis that the average MPC is lower when accounting for coholders who otherwise appear as hand-to-mouth households. This implies that dynamic aggregate household behavior in response to a positive income shock will be similar with coholding and without.

By extension, it follows that the iMPIW matrix in the model of coholding resembles that in the standard two-asset model. However, the model of coholding allows a further decomposition into the intertemporal marginal propensities to save (iMPS) and repay debt (RD). Figure 19 plots columns of the iMPS and iMPRD matrices, which sum to the iMPIW matrix plotted in Panel (b) of Figure 19. In the model of coholding, the aggregate propensity to increase wealth out of a fully unanticipated income shock (i.e., when $s = 0$) is evenly divided between an increase in saving and an increase in debt repayment. In anticipation of future shocks, however, there is a stronger negative savings response than negative debt repayment response. Intuitively, in anticipation of a future positive income shock, households deplete savings to increase consumption more so than they increase

Figure 20: Intertemporal Marginal Propensities to Save and Repay Debt



Notes: Panels (a) and (b) plot the aggregate marginal propensities to save and repay debt, respectively, out of a positive income shock in period $t + s$ relative to the current period $t = 0$. See Figure 19 and Section 7.5 for more details.

borrowing.

The implications of this decomposition of household behavior into saving and borrowing depend crucially on the supply of liquid assets and liquid debt. In standard one- or two-asset models, the interest rate adjusts in equilibrium such that a single financial intermediary supplies the market-clearing level of net assets. In the model of coholding with separate gross instruments for borrowing and saving, the exact specification of financial intermediaries will determine the extent to which household coholding changes aggregate dynamics. This specification includes not only how interest rates are determined in equilibrium, but also on which households own each financial intermediary and how dividends are distributed. Thus a full understanding of the implications of coholding in general equilibrium requires incorporating the household block developed in this paper with recent advances on the financial block, such as heterogeneous banks (Belifemine et al., 2023) and endogenous credit spreads (Faccini et al., 2024), which we leave for future work.

8 Conclusion

We build a quantitative model of household consumption, saving, and borrowing built on insights from the household finance literature on the coholding of credit card debt and liquid assets. We use the model to understand the empirical evidence on the marginal propensities to consume, save, and repay debt, which is infeasible in standard mod-

els that only consider net wealth positions instead of the joint distribution of liquid assets and debt. We adapt the standard model by adding a parsimonious liquidity-in-advance constraint and, without explicitly targeting them, generate relationships between the marginal propensities to consume, save, and repay debt and the joint distribution of liquid assets and debt that largely resemble the data.

The model's key insight is that there are two groups of households with low liquid wealth: the true hand-to-mouth, who have low net wealth and low liquid assets, and the coholders, who have low net wealth and high liquid assets. These households appear identical if considering only net wealth, but behave very differently in response to transitory income shocks; the former have a large MPC and the second have a low MPC. Coholders have a low MPC because it is optimal for them to repay debt instead of increase consumption. This has important implications for fiscal stimulus policy, especially as household indebtedness grows. Future work will continue to study the implications of coholding for optimal fiscal and monetary policy, especially in a general equilibrium environment with aggregate shocks and endogenous supplies of assets and debt.

References

- Agarwal, Sumit, Chunlin Liu, and Nicholas S Souleles**, “The reaction of consumer spending and debt to tax rebates—evidence from consumer credit data,” *Journal of Political Economy*, 2007, 115 (6), 986–1019.
- Allen, Jason, Robert Clark, Shaoteng Li, and Nicolas Vincent**, “Debt-relief programs and money left on the table: Evidence from Canada’s response to COVID-19,” *Canadian Journal of Economics/Revue canadienne d’économique*, Feb. 2022, 55 (S1), 9–53.
- Auclert, Adrien, Matthew Rognlie, and Ludwig Straub**, “The Intertemporal Keynesian Cross,” Technical Report, NBER Working Paper #25020 2023.
- Baker, Scott R**, “Debt and the response to household income shocks: Validation and application of linked financial account data,” *Journal of Political Economy*, 2018, 126 (4), 1504–1557.
- Batista, Rafael M, Ella Mao, and Abigail B Sussman**, “A Psychological Account of Co-Holding: Why Some Choose to Simultaneously Save and Borrow,” *Available at SSRN 4558490*, 2023.
- Bellifemine, Marco, Rustam Jamilov, and Tommaso Monacelli**, “HBANK: Monetary Policy with Heterogeneous Banks,” Technical Report, Working Paper Dec. 2023.
- Bernardini, Marco and Gert Peersman**, “Private debt overhang and the government spending multiplier: Evidence for the United States,” *Journal of Applied Econometrics*, 2018, 33 (4), 485–508.
- , **Selien De Schryder, and Gert Peersman**, “Heterogeneous government spending multipliers in the era surrounding the great recession,” *Review of Economics and Statistics*, 2020, 102 (2), 304–322.
- Bertaut, Carol C, Michael Haliassos, and Michael Reiter**, “Credit card debt puzzles and debt revolvers for self control,” *Review of Finance*, 2009, 13 (4), 657–692.
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess**, “Revisiting Event-Study Designs: Robust and Efficient Estimation,” *Review of Economic Studies*, Feb. 2024.
- Boutros, Michael**, “Household Finances and Fiscal Stimulus in 2008,” *Available at SSRN 3716571*, 2019.

- , “Windfall Income Shocks with Finite Planning Horizons,” Technical Report, Bank of Canada Staff Working Paper 2022-40 2023.
- , **Nuno Clara, and Francisco Gomes**, “Borrow Now, Pay Even Later: A Quantitative Analysis of Student Debt Payment Plans,” Technical Report, Working Paper Oct. 2023.
- Bunn, Philip, Jeanne Le Roux, Kate Reinold, and Paolo Surico**, “The consumption response to positive and negative income shocks,” *Journal of Monetary Economics*, 2018, 96, 1–15.
- Catherine, Sylvain and Constantine Yannelis**, “The distributional effects of student loan forgiveness,” *Journal of Financial Economics*, Feb. 2023, 147 (2), 297–316.
- Choi, Hwan-sik and Ron A Laschever**, “The credit card debt puzzle and noncognitive ability,” *Review of Finance*, 2018, 22 (6), 2109–2137.
- 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?,” Technical Report, National Bureau of Economic Research 2020.
- Demyanyk, Yuliya, Elena Loutskina, and Daniel Murphy**, “Fiscal stimulus and consumer debt,” *Review of Economics and Statistics*, 2019, 101 (4), 728–741.
- Dinerstein, Michael, Constantine Yannelis, and Ching-Tse Chen**, “Debt Moratoria: Evidence from Student Loan Forbearance,” *American Economic Review: Insights*, Forthcoming.
- Druedahl, Jeppe and Casper Nordal Jørgensen**, “Precautionary borrowing and the credit card debt puzzle,” *Quantitative Economics*, 2018, 9 (2), 785–823.
- Dynan, Karen, Wendy Edelberg et al.**, “The relationship between leverage and household spending behavior: evidence from the 2007-2009 survey of consumer finances,” *Federal Reserve Bank of St. Louis Review*, 2013, 95 (5), 425–448.
- Faccini, Renato, Seungcheol Lee, Ralph Luetticke, Morten O. Ravn, and Tobias Renkin**, “Financial Frictions: Micro vs Macro Volatility,” Technical Report, Working Paper Jan. 2024.

- Fagereng, Andreas, Martin B Holm, and Gisle J Natvik**, “MPC heterogeneity and household balance sheets,” *American Economic Journal: Macroeconomics*, 2021, 13 (4), 1–54.
- Fulford, Scott L**, “How important is variability in consumer credit limits?,” *Journal of Monetary Economics*, 2015, 72, 42–63.
- 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*, 2021, 88 (4), 1760–1795.
- Ganong, Peter and Pascal Noel**, “Liquidity versus Wealth in Household Debt Obligations: Evidence from Housing Policy in the Great Recession,” *American Economic Review*, Oct. 2020, 110 (10), 3100–3138.
- Gathergood, John and Arna Olafsson**, “The co-holding puzzle: New evidence from transaction-level data,” *Available at SSRN 3607560*, 2022.
- **and Jörg Weber**, “Self-control, financial literacy & the co-holding puzzle,” *Journal of Economic Behavior & Organization*, 2014, 107, 455–469.
- Gelman, Michael**, “What drives heterogeneity in the marginal propensity to consume? Temporary shocks vs persistent characteristics,” *Journal of Monetary Economics*, 2021, 117, 521–542.
- 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.
- Greene, Claire and Joanna Stavins**, “Credit card debt puzzle: Liquid assets to pay bills,” Technical Report, FRB of Boston Working Paper 2022.
- Gross, David B and Nicholas S Souleles**, “Do liquidity constraints and interest rates matter for consumer behavior? Evidence from credit card data,” *The Quarterly journal of economics*, 2002, 117 (1), 149–185.
- Jappelli, Tullio and Luigi Pistaferri**, “Fiscal policy and MPC heterogeneity,” *American Economic Journal: Macroeconomics*, 2014, 6 (4), 107–136.
- Kaplan, Greg and Giovanni L Violante**, “A model of the consumption response to fiscal stimulus payments,” *Econometrica*, 2014, 82 (4), 1199–1239.

- **and Giovanni L. Violante**, “The Marginal Propensity to Consume in Heterogeneous Agent Models,” *Annual Review of Economics*, Aug. 2022, 14 (1), 747—775.
- **, Giovanni L Violante, and Justin Weidner**, “The wealthy hand-to-mouth,” Technical Report, National Bureau of Economic Research 2014.
- Klein, Mathias**, “Austerity and private debt,” *Journal of Money, Credit and Banking*, 2017, 49 (7), 1555–1585.
- Kosar, Gizem, Davide Melcangi, Laura Pilosso, and David Wiczer**, “Stimulus and Insurance: The Marginal Propensity to Discharge Debt,” Working Paper 2022.
- Lee, Sean Chanwook and Peter Maxted**, “Credit Card Borrowing in Heterogeneous-Agent Models: Reconciling Theory and Data,” Technical Report, Mimeo 2023.
- Lehnert, Andreas and Douglas M Maki**, “Consumption, debt, and portfolio choice: testing the effect of bankruptcy law,” Finance and Economics Discussion Series 2002-14, Board of Governors of the Federal Reserve System 2002.
- Lopes, Paula**, “Credit card debt and default over the life cycle,” *Journal of Money, Credit and Banking*, 2008, 40 (4), 769–790.
- Lucas, Robert**, “Interest rates and currency prices in a two-country world,” *Journal of Monetary Economics*, 1982, 10 (3), 335–359.
- Mankart, Jochen**, “The (Un-) importance of Chapter 7 wealth exemption levels,” *Journal of Economic Dynamics and Control*, 2014, 38, 1–16.
- Medina, Paolina C and Michaela Pagel**, “Does Saving Cause Borrowing? Implications for the Co-Holding Puzzle,” Technical Report 2023.
- Mian, Atif, Kamalesh Rao, and Amir Sufi**, “Household balance sheets, consumption, and the economic slump,” *The Quarterly Journal of Economics*, 2013, 128 (4), 1687–1726.
- Mijakovic, Andrej**, “Marginal Propensities to Consume with Behavioural Agents,” Technical Report, Working Paper Oct. 2023.
- Misra, Kanishka and Paolo Surico**, “Consumption, Income Changes, and Heterogeneity: Evidence from Two Fiscal Stimulus Programs,” *American Economic Journal: Macroeconomics*, Oct. 2014, 6 (4), 84—106.
- Pulina, Giuseppe**, “Credit card debt puzzle: Evidence from the euro area,” *Economics Letters*, 2024, p. 111586.

Sahm, Claudia R, Matthew D Shapiro, and Joel Slemrod, “Balance-sheet households and fiscal stimulus: Lessons from the payroll tax cut and its expiration,” Technical Report, National Bureau of Economic Research 2015.

Sala, Hector and Pedro Trivin, “Reported MPC in the Presence of Debt,” *Economics Letters*, 2021, 205, 109929.

Telyukova, Irina A, “Household need for liquidity and the credit card debt puzzle,” *Review of Economic Studies*, 2013, 80 (3), 1148–1177.

– **and Randall Wright**, “A model of money and credit, with application to the credit card debt puzzle,” *The Review of Economic Studies*, 2008, 75 (2), 629–647.


Vihriälä, Erkki, “Untangling the credit card debt puzzle,” Technical Report, Working paper 2020.

Zinman, Jonathan, “Household borrowing high and lending low under no-arbitrage,” *Journal of Banking & Finance*, 2007, 31 (5), 1311–1334.

A Appendix

A.1 Alternative Definitions of Credit Card Borrowers and Coholders

Figure A.1: Credit Card Borrowing and Coholding Over Time



Notes: Data from SCF waves 2001–2019. Liquid assets and debt are winsorized at the 99% level.

The literature uses various measures to define credit card borrowers and coholders in the data. Figure A.1 plots borrowers and coholders over time using a range of definitions to take these into account. The upper bound of credit card borrowers is given by households that report not paying off their credit card balance fully. The lower bound is given by households that report revolving debt habitually, conditional on owning a credit card. The upper bound of coholders is given by households that hold more than 10% of monthly income in liquid assets and debt. The lower bound is given by households that additionally report revolving debt habitually. Overall, regardless of the specific definitions used, the patterns remain identical: on the extensive margin, credit card borrowing and coholding has experienced a slight decrease after the global financial crisis and has been increasing ever since.

A.2 Data Description

The New York FED Survey of Consumer Expectations (SCE) is a monthly online survey of a rotating panel of around 1,300 households. It collects information on household expectations and decisions on a variety of topics and provides detailed accounts of household income, balance sheets, and demographics.

We combine the monthly SCE core survey with two additional modules at lower frequency, the Spending Survey and the Household Finance Survey, which contain information on marginal propensities and household balance sheets. Our merged dataset covers the period 2015–2018. We restrict the analysis to households aged 25–65 with income of 1000 USD or more. We drop households for which we do not observe income, marginal propensities, or liquid assets and debt. For the regression analysis and binscatters, we also trim the top 1% of the liquid asset and debt distribution, and the top 1% and bottom 1% of the liquid wealth distribution.

Next, we describe how selected variables are measured.

Marginal propensities to consume, save, and repay debt (SCE). *Suppose next year you were to find your household with 10% more income than you currently expect. What would you do with the extra income?*

Participants are asked to give both a qualitative and a quantitative response in which they specify what percentage of additional income they would spend, save, or use to pay down debt.

Liquid assets (SCE). *Approximately what is the total current value of your [and your spouse's/partner's] savings and investments (such as checking and savings accounts, CDs, stocks, bonds, mutual funds, Treasury bonds), excluding those in retirement accounts? What proportion of the money in your [and your spouse's/partner's] saving and investment accounts (excluding funds in retirement accounts) is invested in: 1: Checking/saving accounts, 2:...*

We define liquid assets as the value of savings held in category 1, checking and savings accounts.

Liquid debt (SCE). *Next consider all outstanding debt you [and your spouse/partner] have, including balances on credit cards (including retail cards), auto loans, student loans, other personal loans, as well as medical or legal bills, but excluding all housing related debt (such as mortgages, home equity lines of credit, home equity loans). Approximately, what is the total amount of your [and your spouse's/partner's] current outstanding debt? What proportion of your [and your spouse's/partner's] current total outstanding debt (excluding all housing related debt) is due to: 1: Credit cards, 2:...*

We define liquid debt as the value of debt in category 1, credit cards. For comparability, liquid debt in the SCF is measured as follows:

Liquid debt (SCF). *After the last payment(s) (was/were) made, what was the total balance still owed on (this account/all these accounts)?*

Accounts refer to credit/company/store-branded/charge cards.

A.3 Additional Evidence on Marginal Propensities and Household Characteristics

In this section, we consider extensions to Table 1. We add controls for illiquid assets and debt, income, housing status, and financial literacy. We also investigate if introducing balance sheet variables and income in standardized terms instead of USD values affects our results. Table A.1 reports the results of our robustness exercises. Overall, our main results are preserved qualitatively, irrespective of the specification.

Table A.1: Regressions of Marginal Propensities on Household Characteristics

	in '000 USD						Standardized					
	(1) Spend	(2) Save	(3) Repay Debt	(4) Spend	(5) Save	(6) Repay Debt	(7) Spend	(8) Save	(9) Repay Debt	(10) Spend	(11) Save	(12) Repay Debt
Liquid Assets	0.007 (0.016)	0.244*** (0.027)	-0.251*** (0.024)	0.008 (0.016)	0.228*** (0.028)	-0.235*** (0.025)	0.235 (0.523)	7.874*** (0.876)	-8.086*** (0.777)	0.257 (0.529)	7.349*** (0.891)	-7.582*** (0.798)
Liquid Debt	-0.204*** (0.038)	-0.667*** (0.059)	0.872*** (0.078)	-0.218*** (0.039)	-0.702*** (0.060)	0.921*** (0.076)	-2.557*** (0.471)	-8.366*** (0.744)	10.935*** (0.973)	-2.730*** (0.491)	-8.810*** (0.747)	11.551*** (0.953)
Illiquid Assets				-0.000 (0.000)	0.000** (0.000)	-0.000 (0.000)				-0.306 (0.256)	0.984** (0.328)	-0.670 (0.482)
Illiquid Debt				0.004** (0.001)	0.002 (0.002)	-0.007** (0.002)				0.902** (0.300)	0.551 (0.493)	-1.458** (0.501)
Income				-0.001 (0.000)	0.005*** (0.001)	-0.004*** (0.001)				-0.267 (0.143)	1.534*** (0.207)	-1.270*** (0.206)
Mortgager				-5.988*** (1.458)	1.911 (2.285)	4.061 (2.420)				-5.988*** (1.458)	1.911 (2.285)	4.061 (2.420)
Homeowner				-3.095 (1.764)	7.490** (2.564)	-4.473 (2.620)				-3.095 (1.764)	7.490** (2.564)	-4.473 (2.620)
Moderate financial literacy				0.418 (1.938)	1.087 (3.082)	-1.239 (3.288)				0.418 (1.938)	1.087 (3.082)	-1.239 (3.288)
High financial literacy				2.850 (2.167)	-0.456 (3.422)	-2.106 (3.687)				2.850 (2.167)	-0.456 (3.422)	-2.106 (3.687)
N	2742	2742	2742	2615	2615	2615	2742	2742	2742	2615	2615	2615
R ²	0.051	0.158	0.185	0.065	0.174	0.200	0.051	0.158	0.185	0.065	0.174	0.200

Notes: Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Additional controls include age, gender, race, marital status, education, geography, and survey date. Income and balance sheet variables enter the regression either in '000 USD or standardized. Mortgager refers to households with a mortgage; Homeowner refers to households that own a home without having a mortgage. Households are classified as moderately (highly) financially literate if more than half (all) of the questions on financial literacy were answered correctly.

A.4 Home Equity Lines of Credit

This section analyzes the role of home equity lines of credit (HELOCs) in the determination of marginal propensities. We retrieve information on HELOCs from the housing module of the SCE, which is fielded once a year for a subset of the core sample. Our analysis comes with a few caveats. First, the availability of information on HELOCs reduces our sample to around one-fifth of our baseline sample. Second, HELOCs are not necessarily measured at the same time of the year as other balance sheet items. Third, we cannot distinguish between home equity loans and HELOCs. For simplicity, we refer to this composite as HELOC.

Table A.2 reports our baseline regression for the subsample of households for which we observe HELOCs, with and without HELOCs included in the regression. Our results remain largely unaffected by the inclusion of HELOCs in the regression as the coefficients are virtually identical across the specification with and without HELOCs. Similarly to credit card debt, a higher amount of HELOC debt is associated with a lower MPC. In terms of magnitude, however, an additional 1,000\$ of HELOC debt only reduces the MPC by 0.09 pp compared to 0.26 pp for credit card debt. This is not surprising given that interest rates on HELOCs tend to be substantially lower than interest rates on credit card debt.

Table A.2: Regressions of Marginal Propensities on HELOCs

	(1) Spend	(2) Save	(3) Repay Debt	(4) Spend	(5) Save	(6) Repay Debt
Liquid Assets	-0.042 (0.026)	0.203*** (0.058)	-0.162** (0.059)	-0.048 (0.026)	0.201*** (0.058)	-0.154** (0.059)
Liquid Debt	-0.266*** (0.056)	-0.606*** (0.109)	0.872*** (0.126)	-0.255*** (0.055)	-0.602*** (0.109)	0.857*** (0.125)
HELOC				-0.088*** (0.025)	-0.032 (0.047)	0.120* (0.061)
N	537	537	537	537	537	537
R^2	0.201	0.261	0.306	0.210	0.261	0.310

Notes: Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Additional controls include age, gender, race, marital status, education, geography, and survey date. Balance sheet variables enter the regression in '000 USD.

A.5 Empirical Evidence from Other Datasets

This section explores to what extent our empirical results on the relation between marginal propensities and household liquid balance sheets are generalizable to other settings. For this purpose, we repeat our analysis using two alternative datasets, one for Italy (Jappelli and Pistaferri, 2014) and one for the Netherlands (Christelis et al., 2019).

A.5.1 Italian Survey of Household Income and Wealth

We first revisit the empirical evidence in Jappelli and Pistaferri (2014) (JP14). The authors find a negative gradient between the MPC and cash-on-hand, defined as financial assets plus income. They also provide suggestive evidence for indebted households having lower MPCs. We will explore this aspect in more detail.

The authors use the following question from the 2010 wave of the Italian Survey of Household Income and Wealth (SHIW) to elicit MPCs:

Imagine you unexpectedly receive a reimbursement equal to the amount your household earns in a month. How much of it would you save and how much would you spend? Please give the percentage you would save and the percentage you would spend.

The question is broadly comparable to the one posed in the SCE, but does not distinguish between saving and repaying debt. For this reason, we only analyze the MPC. Empirically, we closely follow the strategy in JP14. We estimate a cross-sectional Tobit regression of the MPC on quantiles of liquid wealth, liquid debt, and a set of covariates that includes age, gender, marital status, education, location, and family size. Different from JP14, we focus on liquid wealth and financial debt instead of cash-on-hand to be closer to the specification estimated in Table 1. However, our results are similar when we use cash-on-hand instead of liquid wealth.

Liquid, or financial wealth, is composed of financial assets and debt. Financial assets include deposits, government securities, trade credit, and other securities. Financial debt includes liabilities to banks and financial companies, trade debt, and liabilities to other households. Most liquid debt consists of bank liabilities, as credit cards are much less common in Italy than in the United States.

Table A.3 reports the results of this exercise. Column 1 shows that the MPC is decreasing in liquid wealth, in line with the findings in JP14. The higher the quintile of the liquid wealth distribution, the lower the MPC compared to the first quintile that serves as the comparison group. Once we additionally control for the composition of liquid wealth by including terciles of the liquid debt distribution, we observe that for a given quantile of wealth, the MPC is again decreasing in the amount of debt that is held. This effect

Table A.3: Regressions of MPC on Household Liquid Balance Sheet in SHIW

	(1) MPC	(2) MPC	(3) MPC
II net financial wealth quintile	-0.025 (0.022)	-0.093*** (0.023)	
III net financial wealth quintile	-0.066*** (0.018)	-0.144*** (0.019)	
IV net financial wealth quintile	-0.131*** (0.018)	-0.208*** (0.020)	
V net financial wealth quintile	-0.175*** (0.019)	-0.262*** (0.021)	
I financial debt tercile given pos. debt		-0.153*** (0.022)	
II financial debt tercile given pos. debt		-0.138*** (0.024)	
III financial debt tercile given pos. debt		-0.224*** (0.026)	
II gross financial wealth quintile			-0.093*** (0.019)
III gross financial wealth quintile			-0.153*** (0.020)
IV gross financial wealth quintile			-0.228*** (0.020)
V gross financial wealth quintile			-0.243*** (0.022)
N	7950	7950	7950
R^2	0.071	0.078	0.075

Notes: This table reports results from a Tobit regression of the MPC on household balance sheets and a set of covariates that includes age, gender, marital status, education, location, and family size. For net and gross financial wealth quintiles, the first (lowest) quintile serves as the comparison group. For financial debt terciles, households without financial debt form the comparison group. Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

becomes stronger for higher quantiles of debt. Note that here, the comparison group consists of households without liquid debt and terciles are constructed conditional on holding positive amounts of liquid debt. In Column 3, we control for gross instead of net financial wealth, defined as the sum of financial assets and debt. We find that, similarly to the results in the SCE, the MPC decreases in the amount of gross wealth held.

A.5.2 CentER Internet Panel

We next revisit the evidence in Christelis et al. (2019) that uses the Dutch CentER Internet panel maintained by CentERdata at Tilburg University. The authors find a negative gradient between the MPC and cash-on-hand, defined as financial assets plus income. MPCs are elicited through the following question:

Imagine you unexpectedly receive a one-time bonus from the government equal to the amount of net income your household earns in three-months. In the next 12 months, how would you use this unexpected income transfer?

The survey allows households to choose between savings, repaying debt, durable consumption, and non-durable consumption. We focus on non-durable consumption as our measure of the MPC. Compared to the SCE, the income change is larger and the question explicitly mentions the horizon over which the money would be spent.

For the regression analysis, we closely follow the empirical strategy in Christelis et al. (2019). In particular, we estimate a cross-sectional OLS regression of the MPC on quantiles of financial wealth, financial debt, and a set of covariates that includes age, gender, and household size. Different from Christelis et al. (2019), we focus on financial wealth and financial debt instead of cash-on-hand to be closer to the specification estimated in Table 1. However, our results also hold when we use cash-on-hand instead of liquid wealth. Note that financial debt excludes mortgage debt.

Table A.4 reports the results of this exercise. Column 1 shows that the MPC is decreasing in financial wealth, in line with the findings in Christelis et al. (2019). Households in higher quartiles of the financial wealth distribution have higher MPCs on average compared to the lowest quartile, which serves as the comparison group. Once we additionally control for the composition of financial wealth by including measures of financial debt, we observe that for a given quartile of wealth, the MPC is again decreasing in the amount of debt. Due to the relatively small sample, we split households into three groups: no financial debt (around 80%), low financial debt holdings (below median conditional on positive debt), and high financial debt holdings (above median conditional on positive debt). The comparison group consists of households without financial debt. In Column 3, we control for gross instead of net financial wealth, defined as the sum of financial assets and debt. We find that, similar to the results in the SCE, the MPC is decreasing in the amount of gross wealth held. With respect to the MPS and MPRD, we observe similar patterns as in our baseline regression. The MPS is increasing in net wealth, but decreasing in debt. The MPRD instead is decreasing in net wealth but increasing in debt.

Table A.4: Regressions of MPC on Household Liquid Balance Sheet in Center

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	MPC	MPC	MPC	MPS	MPS	MPS	MPRD	MPRD	MPRD
II net financial wealth quartile	-0.010 (0.013)	-0.024 (0.014)		0.088*** (0.023)	0.055* (0.024)		-0.106*** (0.022)	-0.052* (0.021)	
III net financial wealth quartile	-0.038** (0.013)	-0.054*** (0.014)		0.097*** (0.023)	0.058* (0.025)		-0.102*** (0.022)	-0.038 (0.023)	
IV net financial wealth quartile	-0.014 (0.014)	-0.030 (0.015)		0.143*** (0.025)	0.105*** (0.026)		-0.170*** (0.022)	-0.109*** (0.022)	
Low financial debt given pos. debt		-0.039** (0.015)			-0.080** (0.029)			0.142*** (0.029)	
High financial debt given pos. debt		-0.044** (0.016)			-0.108*** (0.029)			0.171*** (0.031)	
II gross financial wealth quartile			-0.039** (0.013)			0.009 (0.023)			-0.004 (0.021)
III gross financial wealth quartile			-0.041** (0.013)			0.023 (0.024)			0.014 (0.022)
IV gross financial wealth quartile			-0.026 (0.014)			0.080** (0.025)			-0.082*** (0.019)
N	1332	1332	1332	1326	1326	1326	1332	1332	1332
R ²	0.025	0.034	0.028	0.028	0.040	0.010	0.071	0.118	0.040

Notes: This table reports results from an OLS regression of the MPC, MPS, and MPRD on household balance sheets and a set of covariates that includes age, gender, and household size. For net and gross financial wealth quartiles, the first (lowest) quartile serves as the comparison group. For financial debt, households without financial debt form the comparison group. Low financial debt includes households with debt below median conditional on positive debt; high financial debt includes households with debt above median conditional on positive debt. Heteroskedasticity-robust standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

A.6 Comparison of SCE and SCF

This section compares households' balance sheets in the SCE and the SCF. Table A.5 provides an overview. On average, households in the SCE hold fewer assets but more liquid debt than in the SCF. They also hold fewer total assets but more total debt. Inspecting median wealth levels, households hold comparable levels of liquid and total wealth across the two surveys. In terms of income, households are very similar across the entire distribution. Coholding is somewhat more prevalent in the SCF. Figure A.2 provides a visual comparison of the entire distribution of liquid assets and debt.

A.7 Income in the Liquidity-in-Advance Constraint

In this section, we determine if our results are sensitive to the specification of the liquidity-in-advance constraint. Specifically, we solve a version of the model in which current income enters the constraint; that is:


$$\theta(c_t - \underline{c}) \leq a_t + y_t,$$

Table A.5: Comparison of SCE and SCF

	SCE				SCF			
	Mean	p25	p50	p75	Mean	p25	p50	p75
Liquid assets	17.59	0.00	1.50	12.00	22.19	0.50	3.20	12.50
Liquid debt	9.38	0.00	1.30	7.50	2.88	0.00	0.00	2.20
Liquid wealth	8.21	-4.50	0.00	10.00	19.31	0.00	1.20	10.16
Total assets	442.76	15.00	185.00	456.00	718.63	26.10	176.70	457.00
Total debt	142.39	5.00	35.00	142.00	114.79	3.00	44.95	153.60
Total wealth	298.27	-1.84	90.00	308.50	603.83	8.25	81.44	307.08
Income	104.19	32.00	63.00	105.00	111.05	31.39	60.76	105.31
Co-holder share	0.31				0.34			
Observations	2871				4580			

Notes: This table compares the distribution of household balance sheets and income across the SCE and SCF. Survey weights are applied. Liquid assets are defined as funds in checking and savings accounts. Liquid debt is defined as credit card debt. Households are considered coholders if they hold more than 10% of monthly income in liquid assets and debt. SCE data refer to the period 2015–2018, SCF data to 2016.

Figure A.2: Distribution of Liquid Assets and Debt in SCE and SCF



Notes: SCE data refer to the period 2015–2018, SCF data to 2016. Unweighted distributions are reported. Liquid assets and debt are reported in thousand USD and truncated at 100 and 30, respectively.

We recalibrate the model to again match the 75th percentile of the debt-to-income distribution, but target the share of coholders instead of the median asset-to-income ratio.

Table A.6 reports the resulting moment. The model with income in the LIA generates a coholding share that is very close to the one in the data. At the same time, it also matches the 75th percentile of the liquid debt-to-income distribution. However, the addition of income worsens the model fit on the asset side. Because households can now finance part of their consumption with current income, they hold fewer assets, which yields coun-

terfactually low asset holdings. Nonetheless, the model still matches average marginal propensities quite well. In unreported analysis, we also observe that marginal propensities across household balance sheets are similar to the baseline model. This suggests that the precise specification of the LIA-constraint does not materially affect our qualitative conclusions.

Table A.6: Comparison of Model and Data Moments with Income in LIA

Panel A: Internally Calibrated Parameters

Parameter	Description	Model with income in LIA
β	Discount factor	0.991995
θ	Share of liquid consumption	0.8615

Panel B: Targeted Moments


	Data	Model with income in LIA
Liquid Assets (Median)	0.61	0.16
Liquid Debt (75 th Pct.)	0.41	0.41
Share of co-holders (%)	0.34	0.34

Panel C: Untargeted Moments

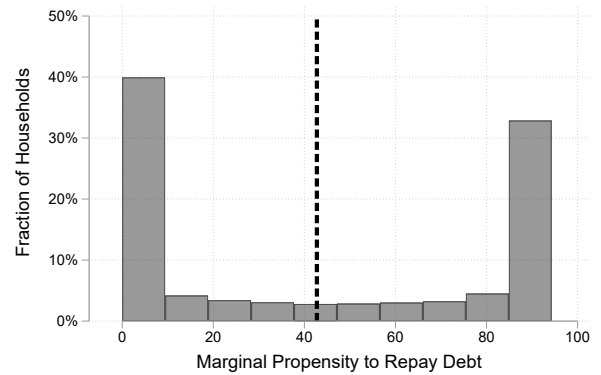
	Data	Model with income in LIA
<i>Distribution of assets and debt</i>		
Liquid Assets (Mean)	2.04	0.32
Liquid Debt (Mean)	0.50	0.09
Liquid Debt (Median)	0.00	0.09
Liquid Wealth (Mean)	1.54	0.06
Liquid Wealth (Median)	0.31	0.06
Liquid debt-to-asset ratio (75 th Pct.)	0.80	3.78
<i>Marginal propensities</i>		
Avg. MPC	16.8	16.4
Avg. MPRD	39.2	42.3
Avg. MPS	44.0	42.6

Notes: Liquid assets and debt are expressed relative to monthly income. Liquid assets are defined as the sum of checking and savings accounts plus idle money in brokerage accounts in the data. Liquid debt is defined as credit card debt, measured by the balance due after the last statement was paid. Households are coholders if they hold more than 10% of monthly income in liquid assets and debt. Balance sheet data are taken from the SCF 2016, while marginal propensities are taken from the SCE 2015–2018.

Figure A.3: Marginal Propensities to Consume and Repay Debt out of Income



(i) Marginal Propensity to Consume



(ii) Marginal Propensity to Repay Debt

Notes: Top and bottom percent of marginal propensities in the model are winsorized to avoid outliers caused by numerical error.

A.8 Marginal Propensities out of Income Versus Assets

This section reports marginal propensities out of income shocks and illustrates that they are not too different from those out of asset shocks. Figure A.3 shows the unconditional distribution of MPCs and MPRDs out of income. The correlation between the MPC out of income and asset shocks is around 0.75.