

Heterogeneity in the Spending Response to Stimulus: Evidence from the Pulse Survey*

Woongchan Jeon^{†1} and Kieran James Walsh^{‡2}

¹UC Santa Barbara

²ETH Zürich

May 2023

Abstract

The US Census Bureau asked households how they spent stimulus payments over 2020 and 2021. Controlling for many demographic variables, we find that while for the 2020 payments the fraction of mostly-spending households was declining in pre-COVID income, in 2021 this stimulus spending distribution was U-shaped. The theory of Miranda-Pinto et al. (2022, 2023) offers an explanation for these results: in crisis times, such as 2020, liquidity constraints are binding for poorer households, rendering them anxious to consume, whereas in a normalized economy (2021) many poorer households are anxious to save or service debt due to expenditure shocks.

Keywords: Marginal Propensity to Consume, Stimulus Payments

JEL Codes: D12, H31, G51

*We thank Vivian Nguyen for excellent research assistance, Daniel Murphy, Jorge Miranda-Pinto, and seminar participants at UCSB (May 2021) for helpful comments, and UCOP Diversity Initiative for financial support.

[†]Email: wjeon@ucsb.edu

[‡]Email: kwalsh@ethz.edu

1 Introduction

In response to economic disruptions caused by COVID-19, the US government distributed three Economic Impact Payments (EIPs) to American households over 2020 and 2021. Concurrently, the US Census administered its new Household Pulse Survey (HPS), which contains detailed information on how the pandemic affected the socioeconomic status of households across the nation. The HPS asked households how they used their EIP stimulus checks, and in this paper, we study the heterogeneity in the responses. In 2020 the rate of mostly spending (vs. saving or paying off debt) from an EIP was declining in household income. But in 2021, after the economy had largely recovered, the relationship between spending and income became much flatter and U-shaped, with spending increasing in income over much of the distribution. We explain how the recent expenditure shock theory of Miranda-Pinto et al. (2022) offers a simple explanation for both this macro state-dependence of the spending propensity distribution and the exotic shape of the distribution in 2021.

An intuitive idea from consumer theory is the consumption function that is concave and increasing in wealth or income, concavity resulting from borrowing constraints or precautionary saving. As current resources fall, the consumer becomes either less able or less willing to consume more presently. The marginal propensity to consume (MPC) is the derivative of the consumption function, and when consumption is concave, it follows that the MPC is decreasing in wealth or income. Poorer households are more likely to be borrowing-constrained or more concerned about future income fluctuations, leading them to disproportionately reduce current consumption. Hence, poorer households are relatively desperate to consume immediately out of new income. In this world, when faced with a random splash of money (e.g., a government stimulus check), richer households mostly save (in line with the Permanent Income Hypothesis (PIH)). Poorer households were already constrained in their ability to consume and thus consume a higher fraction of the same splash of income. See the top row of Figure 1 for a numerical example of this intuitive phenomenon (the full details of this two period model are presented in Section 3). With constant relative risk aversion (CRRA) utility and random income in the second period, first period consumption is concave in present income, and the MPC is decreasing in present income (solid lines). With a tight borrowing constraint (dotted lines), consumption concavity and the declining MPC become more pronounced.

Interestingly, a number of recent papers (referenced below) have uncovered evidence seemingly at odds with this standard theoretical MPC distribution in which MPCs decline monotonically with wealth or income. These papers observe MPCs *increasing* in income

and, combined with existing evidence on higher MPCs amongst the poor, suggest the potential for a non-monotonic relationship in which higher MPCs come from richer households and the very poor. Looking across papers, it seems possible that “moderately-low-income” households, using the parlance of Miranda-Pinto et al. (2023), have the MPCs closest to zero.

We show that this U-shaped cross-household spending propensity distribution arises in the recent HPS. The HPS has surveyed adults across all 50 states and the District of Columbia in weekly or biweekly cross-sectional samples drawn from the Census Bureau’s Master Address File to track socioeconomic developments, including consumption/saving/debt behavior, over the course of the COVID pandemic. Central to our paper, many of the survey weeks have contained questions about the use of stimulus checks. However, the U-shape appears only later in the sample, in 2021, after the US economy had substantially recovered. In contrast, during the throes of the COVID economy in 2020, spending propensities are declining in income.¹

More specifically, for both June-July 2020 and January-July 2021, we first plot the fraction of US households reporting “mostly spend” by the 2019 income bin. Spending rates in 2020 are high and decline in income. Spending rates in 2021, in contrast, are much lower, and the distribution is much flatter with a slight U-shape. To control for confounding factors that may explain these patterns, we use the HPS microdata to estimate a linear probability model for the outcome “mostly spend.” With numerous demographic controls and state-by-week fixed effects, the same patterns emerge as in the US aggregate plot without controls. For both 2019 and 2020, we reject the hypothesis that the spending rate does not vary by income, and the spending rate is clearly downward sloping in income in 2020. Moreover, in 2021 both the lowest income coefficient and highest income coefficients are statistically significantly higher compared to the reference group of \$25,000 to \$34,999. Coefficients related to the effect of job loss on spending also flip signs between 2020 and 2021. In short, the effects of demographics on spending out of stimulus did not just change in magnitude going from 2020 to 2021. They also changed qualitatively.

The odd U-shaped spending distribution is a key feature of the consumption theories established in Miranda-Pinto et al. (2022) and Miranda-Pinto et al. (2023).² In the models

¹See Boutros (2020), Garner et al. (2020), and Garner and Schild (2021) for general analyses of the HPS. Boutros (2020) and Garner et al. (2020) look at 2020 and highlight stimulus spending rates declining in income. In contemporaneous work, Garner and Schild (2021) examine 2021. Their Table 6 logit regressions reveal the U-shape, but they do not emphasize it, instead stating that “households with higher income are more likely [to] ‘mostly spend’ the stimulus payment.”

²While we use the model of Miranda-Pinto et al. (2022) and Miranda-Pinto et al. (2023) to interpret the HPS, the present analysis is distinct from those papers. The former is a quantitative, infinite horizon model used to match the joint dynamics of consumption and income in the Panel Study of Income Dynamics, and the latter is about the cross-country relationship between inequality and the interest rate response to fiscal stimulus.

of these two papers, households face consumption thresholds, which, if violated, yield a utility cost. The idea is that households face a variety of expenditure shocks – education expenses, medical bills, home or car repairs, for example – that impose substantial costs if not serviced. These consumption thresholds, modeled as a kink in the utility function, create a flat portion of the consumption function where the MPC is zero. Households for which consumption thresholds are binding are “saving-constrained” and use additional income to save or delever, as they had previously borrowed or dissaved simply to meet the threshold. At the same time, very poor households are borrowing constrained, meaning they have high MPCs, while rich households have moderate MPCs, consistent with the PIH. Thus, the consumption threshold model generates a U-shaped MPC distribution. The bottom row of Figure 1 illustrates these cases when credit is loose (solid lines). But in recessions with tight credit, the COVID crisis for example, more households are in the borrowing-constrained portion of the consumption function and tighter borrowing constraints make consumption thresholds harder to meet (bottom row of Figure 1 with dotted lines). So in deep recessions the MPC distribution becomes more like the monotonic one from standard consumption theory. Section 3 presents a simple two period version of Miranda-Pinto et al. (2022) that illustrates these mechanisms.³

Our paper relates to the existing literature by presenting new evidence on the spending propensity distribution. Our results indicate a U-shape with respect to income, consistent with the theory of Miranda-Pinto et al. (2022). The potential for this U-shape is discussed in some previous papers. Pistaferri and Saporta Eksten (2012) write, in their MPC literature review, “. . . at least some of the papers document a U-shaped response of consumption to transitory changes both in income and assets. Hence, both the very poor (in terms of income and assets) and the very rich seem to have large consumption responses.” Citing the 2001/2008 tax rebate studies of Shapiro and Slemrod (2003) and Sahm et al. (2010), Campbell and Hercowitz (2019) emphasize that the mostly-spend-rebate fraction of households is U-shaped in stock wealth. And the U-shape is mentioned in footnote 11 of Carroll et al. (2017). But all of these authors note the large standard errors in this literature and none takes a firm stance on the U-shape. In contrast, in our 2021 sample with over 205,000 households and a rich set of control variables, our linear probability model regression reveals a statistically significant U-shape in the “mostly spend” fraction by income (in the sense that we reject no difference in spending rates, and the lowest and highest income groups have a significantly larger spend fraction relative to the reference group). Since we are using new

³Relatedly, Miranda-Pinto et al. (2023) argue that this tension between saving constraints (from their expenditure shock model) and borrowing constraints has important general equilibrium implications at the macro level: in normal times, saving-constrained agents dominate the distribution of households and fiscal stimulus has a muted effect on interest rates. In crises, borrowing-constrained households dominate and fiscal stimulus pushes interest rates up. See also Auerbach et al. (2022) for recent work on the state-dependence of the effects of fiscal policy.




Figure 1
See Footnote 12 for parameter values.

data from a new time period and generating precise estimates, our findings lend credence to the hypothetical U-shape only hinted at in previous papers.

While only a few papers directly discuss the U-shape, arguably it also emerges from a general reading of the literature. On one hand, an abundance of papers finds that lower income or lower wealth households have higher spending propensities. See, for example, Johnson et al. (2006) and Blundell et al. (2008). On the other hand, Kueng (2018) and Lewis et al. (2019) show MPCs increasing with income, and Misra and Surico (2014) find both the highest and lowest MPCs concentrated amongst richer households (so moderate MPCs come from the poor). Other papers, Shapiro and Slemrod (2003) and Souleles (2002) for example, show higher point estimates for higher income households but fail to reject that spending propensities do not differ by income. A theory with a macro state-dependent spending propensity distribution that sometimes has a U-shape – in which moderately-low-income households have low MPCs, very poor households have high MPCs, and wealthy households have moderate MPCs on average – provides a simple explanation for the seemingly contradictory evidence in the literature: the *empirical* shape of the MPC distribution is sensitive to (1) the state of the economy (boom vs. bust) in the particular sample, (2) the part of the income/wealth distribution sampled, and (3) as Misra and Surico (2014) observe, the exogenous cutoffs used by the researcher to group households.

The expenditure shock model of Miranda-Pinto et al. (2022) gives a clear explanation for both what we observe in the HPS as well as the conflicting findings in the previous literature. Additionally, the quantitative analysis of Miranda-Pinto et al. (2022) shows that their model performs better than a standard Bewley model (with or without measurement error) in matching the joint dynamics of consumption and income in the post-1999 Panel Study of Income Dynamics (PSID). But there are at least two other possible theoretical explanations for the U-shape. In the model of Campbell and Hercowitz (2019), households face a borrowing constraint and need to save for large periodic expenses on a second good. In their quantitative assessment of the model, the MPC distribution is U-shaped in wealth (they do not show MPCs by income): the poor and rich have the highest MPCs, the former because of the borrowing constraint and the latter because their wealth is tied-up in anticipation of the major expense. Similarly, the model Kaplan and Violante (2014) exhibits “wealthy hand-to-mouth” agents, whose wealth is tied-up in an illiquid asset. This setting could in principle generate the U-shape, although its presence is not clear in the calibration of Kaplan and Violante (2022) (and they do not show MPCs by income).

Section 2 describes the HPS and presents our empirical results. Section 3 draws on Miranda-Pinto et al. (2022, 2023) to offer a simple theoretical explanation for our empirical observations. Section 4 links the theory and empirics and concludes.

2 Data and Results

The recession caused by COVID-19 led the US government to distribute the three EIPs to American households over 2020 and 2021. Under the Coronavirus, Aid, Relief, and Economic Security (CARES) Act of 2020, tax filers with adjusted gross income (AGI)⁴ up to \$75,000 for individuals and up to \$150,000 for married couples typically received \$1,200 and \$2,400, respectively. The payment amount decreased by \$5 for every \$100 above this threshold. Therefore, without qualifying child dependents, the check amount was completely phased out above \$99,000 for single filers and \$198,000 for joint filers. They could additionally receive \$500 for each child dependent; thus, the total phaseout amount increased by \$10,000 for each qualifying child dependent.⁵ The Internal Revenue Service used the 2018 or 2019 tax returns information to issue payments. The CARES Act was followed by the Consolidated Appropriations Act of 2021 and the American Rescue Plan Act of 2021.⁶

⁴AGI is the total income that is subject to income tax and is defined as gross income minus specific deductions that each filer is eligible to take. Some examples of deductions include alimony payments, 401k contributions, health savings accounts, and education expenses. In general, AGI is lower than or equal to total gross income.

⁵For example, for joint filers with 3 qualifying child dependents, the phaseout amount was \$228,000.

⁶The second check was \$600 for individuals and \$1,200 for married couples with \$600 for each qualifying child dependent. The third check was \$1,400 for single filers and \$2,800 for joint filers with \$1,400 for each qualifying child dependent.

The HPS microdata contains many socioeconomic variables (employment, housing security, household spending, food sufficiency, etc.) that reveal how the pandemic affected the US households. The HPS is designed to construct estimates at three geographical levels—the 15 largest Metropolitan Statistical Areas, the 50 states plus the District of Columbia, and the nation—using the Bureau’s Master Address File as the source of sampled housing units. Approximately one million housing units were selected for each wave out of 145 million addresses. About 68,000-108,000 respondents answered questionnaires in the waves used in our study. To generate estimates representative of the US households at the state level, the HPS also provides sampling weights that account for nonresponse and sampling stratification. In interviews conducted from June 11, 2020 to July 21, 2020 (wave 7 to 12) and from January 6, 2021 to July 5, 2021 (wave 22 to 33), the Bureau asked respondents to report how the EIPs changed their consumption and borrowing behavior. We utilize these two sets of pooled cross-sectional waves to examine the evolution of consumer spending responses over time.⁷ The main goal of our empirical exercise is to test heterogeneity in consumer spending behavior across various income groups after receiving stimulus payments. Thus, we restrict our sample to respondents who said someone in their household received or expected to receive a payment. Summary statistics for our sample are in Table 1.


The main outcome of interest is reported changes in recipients’ consumption behavior. Specifically, the Bureau asked whether a household that received or expected to receive a payment mostly used the stimulus (1) to pay for expenses, (2) to pay off debt or (3) to add to savings. To see how the HPS responses align with consumption theory, we look at the fraction of households in different pre-2020 income groups reporting mostly paying for expenses (“mostly spend”), which we call the average spending rate.

Before turning to our main results, we first plot raw spending rates by income, calculated using the HPS household weights (Figure 2). In the summer 2020 interviews (wave 7 to 12), the fraction of households reporting “mostly spend” is decreasing in income. Over 80% of the < \$25,000 income group reported “mostly spend,” whereas the fraction was less than 60% for the > \$100,000 groups. The relationship is monotonic. The 2021 interviews (wave 22 to 33) reveal a completely different distribution of spending by income. In 2021, the spending rates are much lower (all groups less than 40%), the relationship is much flatter, and there is a slight U-shape, with the minimum at \$25,000 – \$34,999.⁸

But it is difficult to interpret Figure 2 due to omitted variables. Pre-2020 income could easily be correlated with many factors affecting subsequent MPCs, such as other demographic variables, exposure to different regional economic developments and government policies,

⁷A small fraction of respondents repeated surveys for one or two additional weeks from the 7th to 12th waves. Therefore, we only include each respondent’s first survey.

⁸Online Appendix Figure 4 shows the same graph for the 50 states and the District of Columbia individually. Many states, Texas and Illinois for example, have figures almost identical to the national one, and nearly all places exhibit the general pattern described.



Note: b1 = (less than \$25,000), b2 = (\$25,000 - \$34,999), b3 = (\$35,000 - \$49,999), b4 = (\$50,000 - \$74,999), b5 = (\$75,000 - \$99,999), b6 = (\$100,000 - \$149,999), b7 = (\$150,000 - \$199,999), b8 = (\$200,000 and above). This figure plots the proportion of the respondents reporting “mostly spend” by income among households who received or were expected to receive a payment.

Figure 2
Average Spending Rate (Fraction Reporting Mostly Spend) Conditional on 2019 Household Income before Tax

and job loss. This motivates us to account for time-varying state characteristics and other demographics to isolate the association between income and consumption responses.


Using the ordinary least squares (OLS) method,⁹ we estimate a linear probability model where the dependent variable is one for households that mostly used the stimulus to pay for expenses and zero otherwise:

$$\mathbb{I}\{\text{Mostly Spend}\}_{it} = \alpha + \sum_g \beta_g \mathbb{I}\{\text{Income Bin} = g\}_{it} + \boldsymbol{\gamma}' \mathbf{X}_{it} + \delta_{s(i)t} + u_{it}, \quad (1)$$

where i indexes survey respondents, t refers to the week, and u_{it} is the error term. Our main interest is in using the HPS microdata to uncover how households’ consumption responses to fiscal stimulus vary with pre-2020 income (β_g coefficients). Other demographics (X_{it}) include indicators for labor force status, education, housing tenure, marital status, Hispanic origin, race, gender, age, household size, and the number of children in the household.¹⁰ Our specification includes state-by-week fixed effects $\delta_{s(i)t}$ as controls to account for time-varying state-level macroeconomic conditions, COVID-related policies, and other unobservable characteristics. The regression coefficient associated with each covariate measures its average

⁹Logit regressions yield very similar results.

¹⁰Summary statistics are in Table 1.



Note: b1 = (less than \$25,000), b2 = (\$25,000 - \$34,999), b3 = (\$35,000 - \$49,999), b4 = (\$50,000 - \$74,999), b5 = (\$75,000 - \$99,999), b6 = (\$100,000 - \$149,999), b7 = (\$150,000 - \$199,999), b8 = (\$200,000 and above). This figure plots the average partial effect of an income bin b_i on the probability of reporting “mostly spend.” The baseline income category is b_2 . Each point corresponds to a coefficient from the regression specification in equation (1). Vertical bars represent 95% confidence intervals based on the standard errors clustered at the state level. The regression includes demographics and state-by-week fixed effects.

Figure 3

Average Partial Effect on Spending Rate by 2019 Household Income before Tax

partial effect. All regressions use standard errors clustered at the state level to account for within-state interdependence and are weighted by household weights to be representative at the household level. Table 2 presents the OLS estimates.

Figure 3 shows the average partial effect of income, along with the 95% confidence interval. We use the second income bin, \$25,000 – \$34,999, as a reference. Wave 7 to 12 shows a declining spending propensity with respect to income. The F-test rejects that all income coefficients are equal (p-value = 0.000), and the average partial effects are statistically significant.

In contrast, wave 22 - 33 exhibits a U-shaped pattern. Again, the F-test rejects that all income coefficients are equal (p-value = 0.000), and the average partial effects are statistically significant (except for the income bin just above the reference). Compared to the households with income between \$25,000 and \$ 34,999, the households with income less than \$25,000 have a spending rate 2 percentage points higher, and the two highest income groups have a spending rate over 4 percentage points higher.

The pattern with respect to household labor force status is also striking. The survey asks respondents whether they had job losses in their household in the past 4 weeks and whether they expect to have job losses in their household in the next 4 weeks. In 2020,

among the households that did not expect any job losses in the future, those who indeed had job losses in the past are about 14 percentage points more likely to spend. Moreover, among the households that did not have job losses in the past, those who expect job losses in the future are about 13 percentage points more likely to spend. The average spending rate increases by about 21 percentage points when households both had job losses in the past and expect losses in the future. In 2021, however, the spending rate decreases when households had job losses, while the expectation of job losses in the future and the interaction term are statistically not discernible from zero. This seemingly puzzling evolution in reported consumption behavior also holds for other potential indicators of financial hardship such as education and housing tenure. While the theories of Miranda-Pinto et al. (2022, 2023) do not explicitly address labor markets or housing, these patterns may also be explained by the tension between borrowing constraints and saving constraints: in crisis times when borrowing constraints are binding, financial hardship leads to higher spending out of stimulus, while in normal times, many struggling households have borrowed to meet consumption thresholds and are anxious to save or repay debt.

3 Theory

The evidence in Section 2 seems puzzling from the perspective of standard consumption theory. But it turns out a minor and intuitive extension to the standard model generates precisely the patterns we observe in the Pulse Survey. This extension is fully studied in the saving constraint model of Miranda-Pinto et al. (2022, 2023). Here we briefly describe and solve a simple numerical example based on this framework.

Suppose households solve the following two-period consumption/saving problem:

$$\max_{c_1, c_{2s}, a'} \left\{ \frac{c_1^{1-\gamma}}{1-\gamma} - \lambda \max(\underline{c} - c_1, 0) + \sum_{s=1}^S \pi_s \frac{c_{2s}^{1-\gamma}}{1-\gamma} \right\} \text{ subject to} \quad (2)$$

$$c_1 + a' = y_1 \quad (3)$$

$$c_{2s} = y_{2s} + a' \quad (4)$$

$$a' \geq \underline{a}, \quad (5)$$

where c_1 and y_1 are first period consumption and income, c_{2s} and y_{2s} are (state-dependent) second period consumption and income, and a' is saving (negative values of a' mean borrowing). There are S possible states in the second period, which correspond to different realizations of income, and π_s is the probability of state $s \in S$. The household chooses initial consumption, state-dependent second period consumption, and saving to maximize expected utility over consumption, subject to the budget constraints (Equations 3 and 4) and a bor-

rowing constraint (5), where $\underline{a} \leq 0$ is the lower bound on a' . Second period flow utility has the CRRA form with risk aversion γ . First period flow utility is the sum of two terms, CRRA utility with risk aversion γ and a proportional cost of consuming below a threshold $\underline{c} \geq 0$: $c_1^{1-\gamma}/(1-\gamma) - \lambda \max(\underline{c} - c_1, 0)$, where $\lambda \geq 0$ governs the strength of the cost.

Miranda-Pinto et al. (2022, 2023) argue we can think of the consumption threshold as a reduced-form for expenditures on necessities. Car repairs, education expenses, medical bills, home repairs, and family emergencies require the consumption of goods and services and put a lower bound on household spending. Of course, households are not forced to meet these expenses, but ignoring them entails substantial costs (e.g., poor health, inconvenience, car/house depreciation). Moderately-low-income households facing the threshold borrow or dissave just enough to meet the expenditure necessity. These households, in the parlance of Miranda-Pinto et al. (2022, 2023), are saving-constrained in the sense that (1) without the threshold, they would have consumed less presently and (2) on the margin they entirely save or delever out of additional income. On the other hand, the poorest households need to pay the cost from violating the threshold and thus have very high MPCs. Richer households have lower but non-zero MPCs in line with standard theory.¹¹

Figure 1 shows the relationship between current income (y_1) and either consumption (c_1) or the MPC ($\Delta c_1/\Delta y_1$) implied by household optimization.¹² Note that in this partial equilibrium framework, there is a one-to-one mapping between current income and liquid wealth, so varying y_1 can be interpreted as varying liquid wealth.

As described in the introduction, the top row of Figure 1 shows the solution to Problem 2 without consumption thresholds ($\underline{c} = 0$). The solid lines correspond to the case with a loose borrowing constraint ($\underline{a} = -1$), while for the dotted lines we have a tight constraint ($\underline{a} = 0$). Even without a binding borrowing constraint, we see concave consumption and a declining MPC with respect to current income, stemming from precautionary saving. With the tight borrowing constraint, consumption concavity is more pronounced and the MPC rises to one below the income level where borrowing would have begun (≈ 0.5). The bottom row of Figure 1 uses the same parameters as the top row, except with \underline{c} pushed up from 0 to 0.35. With the loose borrowing constraint (solid lines), the MPC by income follows a U-shape, and there is a flat portion of the consumption corresponding to moderately-low-income. But now, when the borrowing constraint tightens, the shapes of consumption and the MPC drastically

¹¹Here, we are using the saving constraint model to understand the MPC distribution. While the model is taken from Miranda-Pinto et al. (2022, 2023), the applications in those papers are quite different, and the present theory is much simpler. Miranda-Pinto et al. (2022) calibrate and solve an infinite horizon version of the model with stochastic thresholds in order to explain the joint dynamics of income and consumption in micro panel data. Miranda-Pinto et al. (2023) solve a general equilibrium version of the model to explain cross-country heterogeneity in the interest rate response to fiscal stimulus.

¹²For our numerical example, we set $\gamma = 2$, $\lambda = 100$, $\underline{c} \in \{0, 0.35\}$, and $\underline{a} \in \{-1, 0\}$. y_{2s} is equally likely to be 0.3, 1, or 1.3, and we consider 50 different values for y_1 , equally-spaced between 0.01 and 1. Therefore, in calculating the MPC, $\Delta y_1 \approx 0.02$.

change. The saving-constrained households become borrowing-constrained (they are unable to borrow enough to get above the consumption threshold), consumption becomes concave again, and the MPC is declining in income.

4 Discussion

Our model, the bottom row of Figure 1 in particular, sheds light on two striking features of the Pulse Survey. The first is the U-shaped spending propensity distribution post-2020. The second is the contrast of this pattern with the spending propensity distribution in the middle of 2020. There was a deep recession in 2020, with a dramatic spike in unemployment and a fall in real GDP in the first half of the year. And as Online Appendix Figure 5 shows, there was an extremely large decline in consumer credit over March-May 2020. While it is difficult to precisely decompose the collapse of credit into supply and demand effects, economists within the Federal Reserve System have argued there was a tightening of consumer lending standards.¹³ And in the Federal Reserve Board’s July 2020 Senior Loan Officer Opinion Survey on Bank Lending Practices, banks on net reported tightening lending standards for all categories of household credit in the second quarter of 2020.¹⁴ The general decline in incomes pushes more households left along their consumption functions (although recall that in the Pulse Survey household income is pre-pandemic), moving households into the credit-constrained region with high MPCs. The lower supply of credit renders more households credit-constrained vs. saving-constrained, as borrowing to meet the consumption threshold becomes more difficult. The result is a cross-household pattern of spending propensities declining with income, as in the standard model. But by 2021, the economy had substantially normalized, with a rapid fall in unemployment and a quick recovery of real GDP. The non-crisis economy with recovered incomes moves households away from credit constraints, allowing them to once again meet consumption thresholds. The result is the normal-times U-shaped MPC distribution.

In summary, our analysis, which is based on the large and nationally representative HPS, reveals that heterogeneity in the response to the stimulus was strikingly different in 2021 vs. 2020. While the 2020 response was roughly consistent with standard theory, with spending clearly declining in income, the distribution of spending propensities in 2021 was much flatter and slightly U-shaped. We argue that the U-shape is explained by the recent theories of Miranda-Pinto et al. (2022, 2023), in which many lower income households are eager to repair their balance sheets that are damaged by spending on necessities. And we suggest there was no U-shape in 2020 because the deep recession made it difficult to satisfy necessities, rendering households anxious to spend.

¹³See, for example, Santucci (2020), Horvath et al. (2020), and Lu and van der Klaauw (2021).

¹⁴<https://www.federalreserve.gov/data/sloos/sloos-202007.htm>

The MPC distribution is central to the propagation of fiscal and monetary shocks, and it informs policymakers about the needs and challenges of households. Our results lend support to a growing literature suggesting that spending out of stimulus is not always monotonically decreasing in income, and we show that the distribution of household spending behavior can vary greatly with the state of the macroeconomy. We have emphasized one particular new theory, but our broader point is that theorists and policymakers should take seriously the possibility of both the low-MPC poor and a strongly state-dependent MPC distribution.

References

- Auerbach, A., Gorodnichenko, Y., McCrory, P. B., & Murphy, D. (2022). Fiscal multipliers in the covid19 recession. *Journal of International Money and Finance*, 126, 102669.
- Blundell, R., Pistaferri, L., & Preston, I. (2008). Consumption inequality and partial insurance. *American Economic Review*, 98(5), 1887–1921.
- Boutros, M. (2020). Evaluating the impact of economic impact payments. *Working Paper*.
- Campbell, J. R., & Hercowitz, Z. (2019). Liquidity constraints of the middle class. *American Economic Journal: Economic Policy*, 11(3), 130–55.
- Carroll, C., Slacalek, J., Tokuoka, K., & White, M. N. (2017). The distribution of wealth and the marginal propensity to consume. *Quantitative Economics*, 8(3), 977–1020.
- Garner, T., Safir, A., & Schild, J. (2020). Receipt and use of stimulus payments in the time of the Covid-19 pandemic. *Beyond the Numbers: Prices and Spending*, 9(10).
- Garner, T., & Schild, J. (2021). Consumer response to economic impact payments during the COVID-19 pandemic and the role of subjective assessments of well-being: A view from the US using a rapid response survey. *Working Paper*.
- Horvath, A., Kay, B., & Wix, C. (2020). The COVID-19 shock and consumer credit. *VOXEU Column*, <https://cepr.org/voxeu/columns/covid-19-shock-and-consumer-credit>.
- Johnson, D. S., Parker, J. A., & Souleles, N. S. (2006). Household expenditure and the income tax rebates of 2001. *American Economic Review*, 96(5), 1589–1610.
- Kaplan, G., & Violante, G. L. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4), 1199–1239.
- Kaplan, G., & Violante, G. L. (2022). The marginal propensity to consume in heterogeneous agent models. *Annual Review of Economics*, 14, 747–775.
- Kueng, L. (2018). Excess sensitivity of high-income consumers. *The Quarterly Journal of Economics*, 133(4), 1693–1751.
- Lewis, D. J., Melcangi, D., & Pilossoph, L. (2019). Latent heterogeneity in the marginal propensity to consume. *FRB of New York Staff Report*, (902).
- Lu, J., & van der Klaauw, W. (2021). Consumer credit demand, supply, and unmet need during the pandemic. *Liberty Street Economics, Federal Reserve Bank of New York*.
- Miranda-Pinto, J., Murphy, D., Walsh, K. J., & Young, E. R. (2022). A model of expenditure shocks. *Mimeo University of Virginia*.
- Miranda-Pinto, J., Murphy, D., Walsh, K. J., & Young, E. R. (2023). Saving constraints, inequality, and the credit market response to fiscal stimulus. *European Economic Review*, 151, 104355.
- Misra, K., & Surico, P. (2014). Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs. *American Economic Journal: Macroeconomics*, 6(4), 84–106.

- Pistaferri, L., & Saporta Eksten, I. (2012). Changes in the income distribution and aggregate consumption. *Government of the Italian Republic (Italy), Ministry of Economy and Finance, Department of the Treasury Working Paper*, (11).
- Sahm, C. R., Shapiro, M. D., & Slemrod, J. (2010). Household response to the 2008 tax rebate: Survey evidence and aggregate implications. *Tax Policy and the Economy*, *24*(1), 69–110.
- Santucci, L. (2020). How has the COVID-19 pandemic affected the supply of consumer credit? *Consumer Finance Institute Special Report, Federal Reserve Bank of Philadelphia*.
- Shapiro, M. D., & Slemrod, J. (2003). Consumer response to tax rebates. *American Economic Review*, *93*(1), 381–396.
- Souleles, N. S. (2002). Consumer response to the reagan tax cuts. *Journal of Public Economics*, *85*(1), 99–120.

Table 1
Sample summary statistics

	Wave 7 - 12 (06/11/20-07/21/20)	Wave 22 - 33 (01/06/21-07/05/21)
	Fraction of Households / Mean	Fraction of Households / Mean
Use of economic impact payments		
Mostly to pay for expenses	0.7168	0.2432
Mostly to pay off debt	0.1558	0.2511
Mostly to add savings	0.1273	0.5058
Household income before taxes in 2019		
Less than \$25,000	0.1906	0.1784
\$25,000 - \$34,999	0.1379	0.1392
\$35,000 - \$49,999	0.1474	0.1528
\$50,000 - \$74,999	0.1958	0.2038
\$75,000 - \$99,999	0.1329	0.1314
\$100,000 - \$149,999	0.1298	0.1354
\$150,000 - \$199,999	0.0468	0.0417
\$200,000 and above	0.0188	0.0173
Job loss status/expectation in households		
Had job losses in the last 4 weeks	0.4841	0.410
Expect job losses in the next 4 weeks	0.3343	0.2065
Education		
Less than high school	0.4072	0.3962
Some college/Associate's	0.3150	0.3067
Bachelor's	0.1640	0.1770
Graduate	0.1138	0.1202
Housing tenure		
Owner/occupied without rent	0.2019	0.2385
Mortgage	0.4304	0.4258
Rent	0.3678	0.3357
Currently married	0.4942	0.5128
Hispanic origin	0.1551	0.1656
Race		
White	0.7423	0.7544
Black	0.1507	0.1363
Asian	0.0463	0.0534
Other	0.0607	0.0559
Female	0.5293	0.5395
Age	47.4882	50.2475
Household size	2.9748	2.8702
Number of children in households	0.7426	0.6872
Observations	316,343	205,369

Note: This table reports the summary statistics for the variables used in the regression analysis. Observations are at the individual respondent level. Age, household size, and the number of children in households are continuous variables, but the rest are discrete variables.

Table 2
Pooled regression of economic impact payment usage


	Wave 7 - 12 (06/11/20-07/21/20)	Wave 22 - 33 (01/06/21-07/05/21)
	Mostly spend	Mostly spend
Household income before taxes in 2019		
(Reference: \$25,000 - \$34,999)		
Less than \$25,000	0.0428 (0.0071)	0.0250 (0.0081)
\$35,000 - \$49,999	-0.0286 (0.0076)	0.0037 (0.0065)
\$50,000 - \$74,999	-0.0855 (0.0077)	0.0166 (0.0067)
\$75,000 - \$99,999	-0.1324 (0.0096)	0.0215 (0.0074)
\$100,000 - \$149,999	-0.1820 (0.0082)	0.0291 (0.0073)
\$150,000 - \$199,999	-0.1968 (0.0103)	0.0406 (0.0111)
\$200,000 and above	-0.2093 (0.0157)	0.0476 (0.0139)
Job loss status/expectation in households		
(Reference: Had none and expect none)		
Had losses in the past 4 weeks	0.1417 (0.0052)	-0.0185 (0.0047)
Expect losses in the next 4 weeks	0.1373 (0.0102)	-0.0068 (0.0136)
Had losses and expect losses	-0.0736 (0.0117)	0.0085 (0.0139)
Education		
(Reference: Graduate)		
Less than high school	0.0914 (0.0058)	-0.0589 (0.0055)
Some college/Associate's	0.0700 (0.0058)	-0.0496 (0.0051)
Bachelor's	0.0066 (0.0068)	-0.0177 (0.0053)
Housing tenure		
(Reference: Owner/occupied without rent)		
Mortgage	0.0243 (0.0050)	-0.0256 (0.0039)
Rent	0.0400 (0.0060)	-0.0204 (0.0050)

(table continued on next page)

Note: This table reports the ordinary least squares estimates of equation (1). The dependent variable is the indicator of individual respondents reporting “mostly spend.” The baseline category for each discrete covariate is in parenthesis. Standard errors (in parentheses) are clustered at the state level.

	Wave 7 - 12 (06/11/20-07/21/20)	Wave 22 - 33 (01/06/21-07/05/21)
	Mostly spend	Mostly spend
Marital status		
(Reference: Not married)		
Currently married	-0.0049 (0.0043)	-0.0006 (0.0038)
Hispanic origin		
(Reference: Not hispanic)		
Hispanic	0.0163 (0.0063)	-0.0348 (0.0061)
Race		
(Reference: White)		
Black	0.0345 (0.0066)	-0.0230 (0.0053)
Asian	0.0454 (0.0116)	0.0637 (0.0079)
Other	0.0106 (0.0076)	-0.0060 (0.0072)
Gender		
(Reference: Male)		
Female	0.0019 (0.0034)	-0.0533 (0.0039)
Other demographic status		
Age	0.0084 (0.0007)	-0.0017 (0.0007)
Age ²	-0.0001 (7.24e-06)	3.49e-5 (7.08e-06)
Household size	0.0051 (0.0016)	0.0041 (0.0018)
Number of children in households	0.0199 (0.0023)	0.0023 (0.0026)
State × Week fixed effects		
	Yes	Yes
Observations	316,343	205,369
R-squared	0.1323	0.0313


Note: This table reports the ordinary least squares estimates of equation (1). The dependent variable is the indicator of individual respondents reporting “mostly spend.” The baseline category for each discrete covariate is in parenthesis. Standard errors (in parentheses) are clustered at the state level.



Note: b1 = (less than \$25,000), b2 = (\$25,000 - \$34,999), b3 = (\$35,000 - \$49,999), b4 = (\$50,000 - \$74,999), b5 = (\$75,000 - \$99,999), b6 = (\$100,000 - \$149,999), b7 = (\$150,000 - \$199,999), b8 = (\$200,000 and above). This figure plots the proportion of the respondents reporting “mostly spend” by income among households who received or were expected to receive a payment for each state and the District of Columbia.

Figure 4

FOR ONLINE APPENDIX Conditional Average Spending Rate by 2019 Household Income before Tax



Note: Shaded area represents recession as determined by the National Bureau of Economic Research (NBER).
Source: Board of Governors of the Federal Reserve System (US), Percent Change of Total Consumer Credit [TOTALSLAR], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/TOTALSLAR>, May 4, 2023.

Figure 5
FOR ONLINE APPENDIX Annualized Monthly Percent Change in Consumer Credit