American Economic Review 2025, 115(1): 1–42
https://doi.org/10.1257/aer.20240138

Five Facts about MPCs: Evidence from a Randomized Experiment[†]

By Johannes Boehm, Etienne Fize, and Xavier Jaravel*

We present five facts from an experiment on the marginal propensity to consume (MPC) out of transitory transfers: (1) the one-month MPC on a cash-like transfer is 23 percent; (2) it is substantially higher (61 percent) on a transfer administered via a card where remaining funds expire after three weeks, inconsistent with money fungibility; (3) the consumption response is concentrated in the first three weeks; (4) MPCs vary with household characteristics but are high even for the liquid wealthy; (5) unconditional MPC distribution exhibits large variation. Our findings inform the design of stimulus policies and pose challenges to existing macroeconomic models. (*JEL* D12, D91, E21, G51, I38)

The marginal propensity of households to consume out of a transitory income shock (MPC) plays a central role in both macroeconomic models and stimulus policies. It determines the partial equilibrium response to such shocks and has broader implications for general equilibrium responses, notably for how monetary and fiscal authorities can boost demand through direct stimulus transfers (e.g., Kaplan, Moll, and Violante 2018; Auclert, Rognlie, and Straub 2023). Despite extensive research, estimates of MPC out of transfers that are relevant for fiscal policy remain debated due to limitations arising from the source of variation used for causal identification (e.g., Parker et al. 2013; Orchard, Ramey, and Wieland 2023b; Borusyak, Jaravel, and Spiess 2023). Furthermore, the recent pandemic-induced downturn

* Boehm: Geneva Graduate Institute and CEPR (email: johannes.boehm@graduateinstitute.ch); Fize: Paris School of Economics, Institut des politiques publiques (email: etienne.fize@ipp.eu); Jaravel: LSE and CEPR (email: x.jaravel@lse.ac.uk). Yuriy Gorodnichenko was the coeditor for this article. We dedicate this paper to the memory of Philippe Martin, without whom this project would not have been possible. For helpful comments, we thank Adrien Auclert, Florin Bilbiie, Corina Boar, Michael Boutros, Peter Ganong, Joe Hazell, Chen Lian, Philippe Martin, Ben Moll, Sarah Mouabbi, Emi Nakamura, Alan Olivi, Jonathan Parker, Valerie Ramey, Xavier Ragot, Ricardo Reis, Jesse Shapiro, Johannes Spinnewijn, Ludwig Straub, Gianluca Violante, and Christian Wolf, as well as seminar and conference participants at the Bank of England, the Bank of France, the Bank of Israel, the Center for Economic Studies at the US Census Bureau, the CEPR public economics symposium, the Federal Reserve Board, George Washington University, HEC-Paris, the Hebrew University of Jerusalem, the Institute for Fiscal Studies, the London School of Economics, the Paris School of Economics, the SITE conference, the Stockholm School of Economics, University College London, and the University of Munich. We are grateful to Emmanuel Giguel, Julien Fournel, and the teams at Euro-Information and Crédit Mutuel Alliance Fédérale for their work on the implementation of the experiment. Boehm gratefully acknowledges financial support from the Agence Nationale de la Recherche under grant agreement ANR-20-CE26-0006 and from the Banque de France/Sciences Po partnership. Jaravel gratefully acknowledges financial support from UK Research and Innovation under grant agreement EP/X016056/1. We are indebted to Antoine Ding and Tyler Woodbury for excellent research assistance. This study is registered as AEARCTR-0009296 in the AEA RCT Registry (Boehm, Fize, and Jaravel 2022). IRB Approval Number 2021-018 was granted by the Research Ethics Committee of the Paris Institute of Political Studies on 2021-03-02.

† Go to https://doi.org/10.1257/aer.20240138 to visit the article page for additional materials and author disclosure statement(s).

has seen an increased variety of large-scale stimulus policies, using prepaid cards or time-limited consumption vouchers (including California, Milan, and Seoul in 2020; Hong Kong and Northern Ireland in 2021; and Thailand in 2023), raising questions about whether the way a stimulus payment is administered might affect economic outcomes. Finding scalable ways of raising MPCs can be an important policy objective given recent estimates of low MPCs out of standard tax rebates in the United States.[1]

In this paper, we estimate MPCs by running a randomized experiment allocating transfers at random across households. We use high-frequency bank data to measure households' overall consumption response and its heterogeneity across households.[2] Going beyond standard estimation of MPCs, we examine whether transfers with time limits or negative interest rates may yield larger MPCs in order to inform policymakers about the role of stimulus design in determining the consumption response to transfers.

Our experiment is designed with scalability and generalizability in mind. We randomly provide stimulus transfers to a sample of about one thousand French individuals who are representative of the adult French population and for whom we observe detailed financial transactions and consumption expenditure data through bank records. The experiment was launched in May 2022 at a time when interest rates were still at zero in the eurozone. Our baseline treatment evaluates the consumption response to a simple one-off money transfer in the form of a debit card with a balance of €300. We compare the total consumption spending of treated households, on both the prepaid cards and their regular bank accounts, with those of a large sample of about ninety thousand untreated households. In further treatment groups, we investigate two potential ways of increasing the households' overall consumption response by assigning a negative interest rate on the transferred wealth: either by making the card expire after three weeks—at which point any remaining balance is lost to the household—or with a weekly deduction of an amount close to 10 percent of the remaining balance on the card. While households in all treatment groups are free to spend the transfer however they want, we make the interest payments potentially binding by preventing cash withdrawals from the cards. We also assign an additional framing treatment where participants are asked to "spend soon, on French products, and on things [they] would have otherwise not purchased." Using this experimental setup, we establish five facts about MPCs. We then discuss why these facts are informative for macroeconomic models and for the design of stimulus policies.

We start by estimating MPCs depending on the card type, establishing our first two key facts. We find that participants in the baseline treatment group (without an expiry date or negative interest rate) increase their total consumption expenditure after receiving the card, with an average MPC of 23 percent over one month

[1] See Parker et al. (2022b) and Parker et al. (2022a) for MPCs out of the 2020 stimulus payment; Borusyak, Jaravel and Spiess (2023) and Orchard, Ramey and Wieland (2023b) for the 2008 stimulus payments; and Orchard, Ramey and Wieland (2023a) for the 2001 stimulus.

[2] A vast literature has examined MPCs out of various shocks, including typical income shocks (Ganong et al. 2020), lottery winnings (Fagereng, Holm, and Natvik 2021; Golosov et al. 2024), and recurring lump-sum payments (Kueng 2018). Instead, we study one-time transfers comparable to those deployed to stimulate the economy during an economic downturn.

(Fact 1). We then establish the main finding of the paper, showing that implementation design matters: the MPC is substantially higher for treatment groups where any remaining balance becomes unusable after three weeks, at 61 percent (or 70 percent when conditioning on take-up), or where remaining balances are subject to the 10 percent negative interest rate every week, at 35 percent (Fact 2). In contrast, we do not find any significant effect of the additional "framing" treatment paragraph in the letter, suggesting that our results are unlikely to be driven by experimenter demand effects.

We examine the possibility that the faster spending for cards with an expiry date or negative interest rates could induce detrimental consequences for these households due to behavioral internalities. We find no such evidence: these households do not incur more volatile nondurable consumption in later periods, and they are not more likely to make purchases that could entail adverse health consequences, such as tobacco or gambling.

We next analyze the dynamics of the consumption response—the path of intertemporal MPCs (iMPCs) (Auclert, Rognlie, and Straub 2023; Angeletos, Lian, and Wolf 2023)—yielding our third key fact. We find that, for all treatment cards, the additional spending occurs immediately after the onset of the experiment. Specifically, the increase in consumption is much larger early on in the first weeks following the transfer (Fact 3). We observe that the consumption response is concentrated early on even for nondurables.

To understand the spending behavior of the participants upon receiving the treatment, we administer a survey among participants, and we analyze the bank data to assess potential changes in the composition of expenditures. Recipients are well aware that they spend less on their main account (thereby having an MPC below one), and they mention precautionary saving as a key motive. They use the card they receive primarily to cover running expenses, but some also report purchasing a "treat," or making a large expenditure earlier. Treated households have similar expenditure shares on most consumption categories as control households but purchase relatively more clothing and household equipment. Treated households also spend slightly more on durables and on imported goods.

We then turn to MPC heterogeneity, establishing our fourth and fifth key facts. We find that there is significant MPC heterogeneity by observed household characteristics, including for liquid wealth, current income, and proxies for permanent income, gender, and age (Fact 4). The most important source of heterogeneity we document is about gender: the average MPC of men is about twice as high as for women. We also find that households with lower income and households with lower average pretreatment consumption levels (our proxy for permanent income) have higher MPCs. Liquid wealth does play a role in explaining MPC heterogeneity, albeit a limited one, and MPCs remain high even for households whose liquid wealth exceeds twice their monthly income. Finally, we find that MPCs appear to increase with age although differences across age groups are relatively noisy. A set of LASSO regressions confirms that the most important predictors of MPC heterogeneity are treatment group memberships and demographic characteristics like gender.

Going beyond heterogeneity that is associated with observed characteristics, we estimate the full unconditional distribution of MPCs across households. Due to our experimental setting, we know that the distribution of error terms is identical in the treatment and control groups. Under the assumption that the treatment effect is independent from the error term,[3] this fact allows us to use statistical deconvolution techniques to estimate the full unconditional distribution of treatment effects. Applying this methodology, we find large unconditional heterogeneity in consumption responses following the transfer (Fact 5). In the baseline treatment group where households receive a cash-like transfer, a quarter of households increase their consumption expenditure over a four-week horizon by less than 13 percent of the transfer, and a quarter increase their consumption by more than 48 percent. In contrast, in the treatment group where cards expire after three weeks, three quarters of recipients increase their four-week consumption by more than 52 percent of the transfer amount. These results again highlight the power of implementation design choices to shift MPCs.

Finally, we discuss the implications of these five facts about MPCs for macroeconomic models and for policy. While our MPC estimates do not speak to general equilibrium effects, they are informative about key building blocks of modern macroeconomic models. Our findings contrast with the predictions of the canonical implementation of the benchmark two-asset Heterogeneous-Agents New Keynesian (HANK) model in three ways. First, the magnitude and dynamics of the MPC are difficult to reconcile with HANK models. In all our treatment groups the entire spending response is concentrated in the first weeks (up to three weeks), while the MPC response is much more long lived according to HANK (Kaplan and Violante 2014; Kaplan, Moll, and Violante 2018; Auclert, Rognlie, and Straub 2023). For example, in Kaplan, Moll, and Violante (2018), the MPC out of a \$300 transfer is 17 percent over a quarter and increases to about 32 percent over a year. Instead, with our baseline treatment (without an expiry date or negative rates), we obtain a larger MPC in the first month, at 23 percent, but no further increases in spending in later periods.[4] Laibson, Maxted, and Moll (2022) note that durables require special treatment when analyzing the dynamic reponse of spending since the effective consumption derived from durables occurs over a long period rather than at the time of purchase. However, the concentrated spending response we estimate is not driven by durables. Second, in HANK the MPC is strongly correlated with the level of liquid assets that agents hold; while we do find some heterogeneity of MPCs for groups with different levels of liquid asset holdings, we find that average MPCs are also high for households that have moderate or high levels of liquid asset holdings. Third, our estimates of the unconditional distributions of MPCs reveal that MPCs are high for a large majority of the population, in contrast to standard calibrations of the HANK model where high MPCs are concentrated among a subset of agents who have low liquid wealth and hit their borrowing constraints. Assessing whether alternative calibrations or extensions of the HANK model can match our five MPC facts is an important direction for future research.

[3] We provide empirical support for this assumption through auxiliary tests in Section III.

[4] Our finding is consistent with quasi-experimental evidence on MPCs using the 2008 US tax rebates. Using scanner data to document high-frequency spending responses to tax rebates, Borusyak, Jaravel, and Spiess (2023) show that they are concentrated in the first month after the rebate.

We also show that our results are difficult to reconcile with agents being rational and treating money as fungible. A rational agent who treats money as fungible should first "use up" the treatment card to avoid potentially losing money (through the negative interest rate or expiry) before using their normal debit or credit card. The transferred amount of €300 is well below the normal three-week consumption expenditure of most households, suggesting that the expiry date in group 2 should in principle not bind (and therefore not affect behavior) for most households. Moreover, we observe that households in the treatment groups with an expiry date or a negative interest rate frequently make payments with other means before exhausting the transfer card. Our results thus echo a literature documenting the nonfungibility of money (Hastings and Shapiro 2013, 2018; Baugh et al. 2021; Geng, Shi, and Song 2022; Gelman and Roussanov 2023) and deliver three lessons for behavioral models. First, models of consumption that rely on present bias in preferences (e.g., Laibson 1997; Maxted 2020; Laibson, Maxted, and Moll 2021) are able to explain why the consumption response to the transfer is concentrated early on but cannot explain the difference in the magnitude of responses between the treatment groups. Indeed, under such preferences consumers in all three groups should be present biased, but the negative interest rate and the expiry date would remain nonbinding constraints given that it should be costless for agents to substitute current account spending for prepaid card spending. Second, while implementations of "spender-saver" models (Campbell and Mankiw 1989) can be made to feature consumption responses that are concentrated very early on, they would also imply strongly bimodal distributions of MPCs, which we do not find. Third, our finding that households consume more when presented with an urgent spending need (in the form of the negative interest rate or expiry date) is consistent with theories where the salience of treatments affects economic choices by drawing attention away from other considerations (Bordalo, Gennaioli, and Shleifer 2012, 2013; Ilut and Valchev 2023). We note that other empirical studies of consumption behavior following income shocks are consistent with theories of salience (Kueng 2018; Baugh et al. 2021).

Our new facts have two implications for policy. First, the large difference in MPCs across treatment groups shows that changing the design of transfers can be a powerful way to increase the MPC. The treatment we find to have the highest MPC takes a particularly simple form: a debit card that features an expiry date, a feature that consumers know from gift vouchers. Second, our estimates of MPC heterogeneity have implications for the targeting of transfers by observable household characteristics.[5] We find that it is possible, based on simple observable characteristics like income, to find household populations with significantly higher MPCs than average. However, the change in MPC obtained by targeting is smaller than when using a card with an expiry date. We conclude that implementation design choices are a more powerful tool, compared with targeting, to increase the recipients' average MPC.

[5] Aguiar, Bils, and Boar (2023) discuss the targeting of individuals in a model where differences in MPCs originate from preference heterogeneity. Gelman (2021) highlights the importance of discount factor heterogeneity in explaining MPC heterogeneity.

Related Literature.—A unique feature of our setting is to use an experiment to analyze how the MPC varies with implementation design choices, comparing the effects of standard transfers to transfers featuring an expiry date or a negative interest rate. More broadly, our results contribute to a vast literature that seeks to estimate marginal propensities to consume.

While there is a very large literature on MPCs (see, e.g., Jappelli and Pistaferri 2010 for a survey), only a relatively small subset of papers analyzes MPCs out of unanticipated and transitory transfers, which are most informative for stimulus policies and macroeconomic models. The literature has taken two types of approaches. First, a range of studies analyze the staggered disbursement of tax rebates. The seminal papers analyzing staggered tax rebates in the United States (Johnson, Parker, and Souleles 2006; Parker et al. 2013; Broda and Parker 2014) found large MPCs of 50 percent to 90 percent over a quarter, which are commonly used to discipline macro models. However, the staggered difference-in-differences design raises an identification challenge: a recent literature finds that using difference-in-differences estimators that are robust to treatment effect heterogeneity yields much smaller MPCs of about 25 percent over a quarter (Borusyak, Jaravel, and Spiess 2023; Orchard, Ramey, and Wieland 2023b). The staggered disbursement leveraged in these studies means that even control households were expecting to receive the rebates at some point. In contrast, our stimulus transfers are entirely unanticipated, and we compare the consumption response of treated households to consumption of households that were entirely untreated. Second, the literature has studied the impact of lottery wins (Fagereng, Holm, and Natvik 2021; Golosov et al. 2024), inferring the consumption response from income and wealth data rather than using direct consumption measures like we do.[6] While these studies benefit from much larger samples than ours, a limitation is that they cannot analyze the consumption response by expenditure type (e.g., durables versus nondurables).

A much broader literature measures the spending responses to anticipated shocks (e.g., Gelman et al. 2014; Kueng 2018; Olafsson and Pagel 2018; McDowall 2020; Gelman 2022; Baugh et al. 2021), typical income shocks with a persistent component (Ganong et al. 2020), and unanticipated permanent shocks (Gelman et al. 2023). Relative to these papers, we focus on unanticipated transitory transfers that are similar to standard stimulus transfers.

Four other strands of the literature provide MPC estimates. First, a growing literature seeks to estimate the distribution of MPCs (Misra and Surico 2014; Lewis, Melcangi, and Pilossoph 2019). A key advantage of our experimental setup is that we can use deconvolution methods to estimate the distribution of treatment effects. Second, another strand of the literature uses theory-informed moment conditions to identify MPCs (Blundell, Pistaferri, and Preston 2008; Commault 2022a). Third, a number of papers have elicited MPCs from surveys where respondents are asked how they would respond to a hypothetical transfer (Shapiro and Slemrod 2003, 2009; Jappelli and Pistaferri 2014; Bunn et al. 2018; Parker and Souleles 2019; Fuster, Kaplan, and Zafar 2021; Commault 2022b). Fourth, a

[6] To impute consumption, Fagereng, Holm, and Natvik (2021) use observed changes in assets and liabilities, while Golosov et al. (2024) leverage estimates of capital gains.

small literature estimates the spending response to consumption vouchers (Hsieh, Shimizutani, and Hori 2010; Kan, Peng, and Wang 2017; Xing et al. 2023; Geng, Shi, and Song 2022; Ding et al. 2024; Chan and Kan 2024). Online Appendix A describes large-scale policies using time-limited consumption vouchers that were deployed in the wake of the COVID-19 pandemic.

The advantages and drawbacks of randomized experiments have been discussed at length in the development literature (see, e.g., Banerjee and Duflo (2009); Banerjee, Duflo, and Kremer (2016); Deaton (2010); and Deaton and Cartwright (2018)). The advantages of experiments are at the core of our contributions: we can evaluate the role of several different stimulus policy designs in the exact same setting thereby cleanly isolating the role of stimulus design (Fact 2); moreover, the assumptions required to identify the *distribution* of treatment effects (Fact 5) are much more likely to be satisfied than in observational settings. However, as any experiment, our analysis is subject to several potential limitations. First, estimates obtained in randomized experiments may suffer from experimenter demand effects, or Hawthorne effects. The finding that our framing treatment does not lead to a significantly different response suggests that experimenter demand effects are not driving our results (De Quidt, Haushofer, and Roth 2018). Second, empirical estimates obtained in a specific setting may not be externally valid. While this concern affects experimental and observational studies alike, the fact that our experimental MPC estimates are close to the observational estimates of Borusyak, Jaravel, and Spiess (2023) and Orchard, Ramey, and Wieland (2023b), which were obtained in a completely different time and setting, helps to alleviate such concerns. Third, due to the limited size of our experimental sample, the standard errors of our estimates are larger than in some quasi-experimental studies (e.g., Ganong et al. 2020), while they are more precise than in others (e.g., Parker et al. 2013). One advantage of our experimental settings for inference is that we can use Fisher's exact test to assess the statistical significance of our estimates—which we use to probe the robustness of our main finding, Fact 2. Finally, our experimental estimates are only informative about a limited number of data moments (e.g., we did not randomize stimulus size); they should therefore be viewed as complementing the existing literature only along specific dimensions.

More broadly, this paper demonstrates the possibility of evaluating and improving the design of macroeconomic policy tools through experimental means.[7] Our experimental stimulus transfers are designed with scalability and real-world implementability in mind. Moreover, our causal estimates help to distinguish between classes of macroeconomic models (Nakamura and Steinsson 2018).

Outline.—The remainder of the paper is organized as follows: Section I presents the data and experimental design; Section II presents our main MPC estimates, establishing our first three key facts; Section III documents MPC heterogeneity, leading to our fourth and fifth key facts; Section IV uses our five facts to draw lessons for macroeconomic models and stimulus policies.

[7] For an experimental study of the consumption response to credit expansions, see Aydin (2022).

I. Data and Experimental Design

A. *Dataset*

Our analysis is made possible by running an experiment on a panel of households for which we have access to comprehensive, detailed financial transactions data.[8] This panel of households has been constructed to be representative of the overall French population. For ethical and operational reasons we restricted this sample before randomly drawing treatment assignments. We describe both the larger and the restricted samples in turn, as well as the content of the data.

The Bank Data and the Experimental Sample.—Our data come from the French banking group Crédit Mutuel Alliance Fédérale (Boehm, Fize, and Jaravel 2025).[9] We start with a panel of households that were drawn by the bank in June 2020 and that are representative of the French population in terms of location, age, and socioeconomic characteristics as shown by Bounie et al. (2020) and Bonnet et al. (2023) who compare the bank sample to official statistics (see online Appendix B.1). The data provide socioeconomic information about the individuals in the household, transaction-level information for transaction accounts, transaction-level information for all payment cards linked to the accounts, debt and balances on nontransaction accounts at the monthly frequency, as well as information about real estate assets at a much lower frequency. Card transactions include information on the Merchant Category Code (MCC) of the vendor, which we describe in online Appendix B.2, along with additional information on the data.

We use a subset of the full household panel for our experiment based on eligibility criteria defined at the individual level. To be eligible, individuals must be between 25 and 75 years of age, must have a known residential address, and should not be deemed by the bank to be financially fragile.[10] In order to obtain a population where we are able to measure spending well, we retain only individuals who are part of a household that, according to the bank's records, do not hold accounts with another bank.[11] We also exclude those individuals who have been using their debit card infrequently in the months prior to the experiment, suggesting that they may predominantly use cash. After applying all conditions, we obtain a balanced sample of 85,700 unique households who have at least one or at most two eligible persons in the household. We also drop 40 households that were treated in a pilot of the experiment. Purchase transactions and assets/liabilities are available at the level

[8] See Baker and Kueng (2022) for a survey of recent research that uses financial transactions data.

[9] Crédit Mutuel Alliance Fédérale made de-identified data available to us on a secure server protecting customer privacy. The bank aims to contribute to the public good and policy debates by facilitating economics research. This is part of Crédit Mutuel Alliance Fédérale's mission as an *entreprise à mission*, a French legal framework in which businesses pursue certain societal goals. The total cost of the experiment, including the transfers and a fee to cover the operational cost, was financed by the researchers through a grant from the French national agency for research (Agence nationale de la recherche).

[10] About 1 percent of households are deemed financially fragile by the bank.

[11] Note that it is relatively rare for households to use multiple banks in France. According to the 2017 French Wealth Survey, 55 percent of French households only have one bank (i.e., no household member uses a different bank). Furthermore, 75 percent of French households have checking accounts at a single bank. Finally, on average, 81 percent of households' total assets are held in their main bank.

of the household. Some specifications will investigate response heterogeneity by characteristics of the *individual* who has been drawn to receive the transfer.

Variable Definitions.—Our main outcome variable is weekly consumption expenditure of the household, defined as the sum of all (credit and debit) card purchases and cash withdrawals of the household between Tuesday and the subsequent Monday at midnight.[12] We winsorize weekly consumption spending using the nontreatment cards at the ninety-ninth percentile of the distribution, which is €1,940, before adding treatment card expenditures to arrive at total weekly consumption expenditure. The results are not sensitive to this winsorization step, as described below. Furthermore, wire transfers and direct debit are not included in our baseline consumption measure, but we analyze an expanded consumption measure including these outflows in robustness checks. Online Appendix B.3 provides more detail on variable definitions.

Our estimated consumption responses are thus marginal propensities to spend (see Laibson, Maxted, and Moll (2022) on the difference between marginal propensities to spend and notional MPCs). For the purpose of studying heterogeneity in consumption responses with respect to observable characteristics, we define time-invariant household characteristics as the average of the corresponding end-of-month characteristic in the six months prior to the treatment (November 2021 to April 2022) per capita.

Summary Statistics.—Table 1 shows summary statistics of the main variables. The table illustrates the richness of the bank data and the large heterogeneity in observable characteristics. Online Appendix Tables D1 and D2 provide additional summary statistics.

B. *Experimental Design*

Treatment Arms.—From the set of eligible individuals, we randomly draw 915 participants over three treatment groups (Boehm, Fize, and Jaravel 2022).

Treatment group 1 (G1, $N = 379$) participants receive a MasterCard debit card linked to a new transactions account with an initial balance of €300. The card expires and becomes unusable five months after it has been sent, after which the participants receive any unspent balance wired to their main transactions account. Prior to this date the participants are unable to transfer funds from or to the newly created transactions account except by means of making purchases with the associated debit card. Notably, participants are unable to withdraw cash from those accounts. Otherwise, the participants are free to spend the account balance wherever MasterCard is accepted (i.e., in stores or online). Participants can monitor the remaining balance through their mobile phone bank app where the account appears alongside their other bank accounts. Note that for purchases made in stores, merchants may be willing to split transactions into several payments, thus allowing the

[12] We choose this interval to line up with the negative interest payments of group 3, which take place on Mondays at midnight. One exception to the construction of weekly aggregates is that we assign Monday, May 2 (the first day when participants use the card), to the subsequent week. The first post-treatment week is therefore comprised of eight days; this feature does not create any challenge for the estimation of MPCs as we use week fixed effects, as described below.

Table 1—Summary Statistics

	Observations	Mean	SD
Age of eligible household member	85,700.00	47.03	12.92
Number of eligible household members	85,700.00	1.15	0.36
Average monthly incoming transfers, six months prior	85,685.00	2,654.04	1,439.56
Average monthly incoming salaries, social allowance, pensions, benefits, six months prior	80,034.00	2,109.55	4,968.86
Dummy: has received unemployment benefits within the last six months	85,685.00	0.14	0.35
Average current account balance, six months prior	85,698.00	4,448.51	19,976.04
Average liquid savings, six months prior	85,698.00	16,896.51	34,466.19
Average value of life insurance assets, six months prior	85,698.00	5,867.43	32,466.19
Average net illiquid wealth, six months prior	85,700.00	64,746.97	185,159.99
Average net liquid wealth, one month prior	85,700.00	19,265.55	46,067.00
Average total debt, six months prior	85,698.00	33,300.8	55,007.44
Average consumer debt, six months prior	85,698.00	2,388.27	5,194.01
Average mortgage debt, six months prior	85,698.00	30,872.90	54,286.18
Number of adult members in the household	85,698.00	1.53	0.50
Number of children in the household	85,698.00	0.61	0.96
Average monthly consumption expenditures (cash, card payments), one year prior	85,698.00	1,205.52	658.29
Average monthly direct debits, debt payments, subscriptions, one year prior	85,698.00	631.29	1,154.20
Average monthly outgoing transfers, one year prior	85,698.00	316.24	639.88
Average total monthly consumption (broad measure)	85,698.00	2,153.05	1,736.50
Weekly consumption expenditure (cash and cards), total	2,571,000.00	417.66	435.02
Weekly consumption expenditure (broad measure), excluding treatment cards	2,571,000.00	727.63	1,992.72

Notes: This table reports summary statistics for our main analysis sample. The broad measure of consumption includes the total of cash withdrawals, card spending, automatic debits, and wire transfers.

participant to purchase an item above €300 by combining the balance available on the treatment card with funds from their regular bank account.

Treatment group 2 (G2, $N = 268$) participants receive the same type of account and card as G1 participants except that the card expires after three weeks; any remaining balance on the account after three weeks is *not* wired to their main checking account, but is deducted from the account and lost to the participants.

Treatment group 3 (G3, $N = 268$) participants receive the same type of account and card as G1 participants except that an "interest" payment is deducted at a weekly frequency. We approximate a 10 percent negative interest rate by decreasing the remaining balance on the account (i) by €30 if the remaining euro balance is in the interval $(200, 300]$, (ii) by €20 if the remaining euro balance is in the interval $[100, 200]$, and (iii) by €10 if the remaining balance is below €100. If the remaining balance is below €10, the entire remaining balance is deducted. The card and account remain active until the balance has reached zero. The deduction rule that we apply has the advantage of being quite similar to a weekly negative rate of 10 percent while still remaining easy to explain and understand.

Orthogonal to the treatment group status, half of the all treated participants (stratified across treatment groups) were additionally treated with a framing treatment where they were encouraged to spend the money quickly on local goods or services and "on items they would not have purchased otherwise, so that the overall increase of [their] spending and its impact on the French economy is maximized" (our translation; for the original see online Appendix C.1).

Dividing the recipient households into these three treatment arms allows us to estimate the extent to which transfer design choices might shift the MPC out of one-time transitory transfers. In particular, we can learn about the role of negative interest rates with the cards in group 3 (close to 10 percent a week) and group 2 (as the expiry amounts to a 100 percent negative rate). While our experiment directly estimates the impact of transfer design choices on the MPC, it is not meant to capture how households would consume in a setting where they face a negative interest rate on their main bank account.

Timeline.—Our experiment took place between May and October 2022. On Wednesday, April 27, the cards (which from now on we will call "treatment cards" or "prepaid cards" to distinguish them from the households' other means of payment), accompanying instructions and explanations (see online Appendix C.1), and pin codes are sent by post to the residential addresses of the selected individuals,[13] with expected arrival on or around Monday, May 2. In the meantime, the bank advisers of the treatment group individuals contact their clients by phone as well as through a banking app, explaining that they have been selected to participate in an academic study and explaining the terms of the cards according to the treatment arm. Participants are informed that they can opt out from the study (in which case they would be unable to use the money they are set to receive) although nobody expressed a desire to do so. The fact that the bank advisors contacted the clients helps alleviate potential concern about participants' mistrust. Another letter with instructions, serving as a reminder, is sent to all participants on Wednesday, May 11.

On Monday, May 9, treatment group 3 participants experience the first weekly deduction for any remaining balance. The second deduction for this group occurs on Monday, May 16, and so on every week from then onward. For treatment group 2 participants, the card expires on Tuesday, May 24. An online survey is sent to all participants in the middle of June, which we use to better understand the spending behavior of the participants. Finally, treatment group 1 cards expire on October 3, and the remaining balances are transferred to the participants' main bank accounts.

Take-Up.—Participants started using the card from May 2 onward. Among the 915 treated households, 830 used the treatment card at least once before October 6. Eighty-five participants chose not to use the card, possibly for economic reasons (e.g., group 1 participants can save by not using the card and getting the remaining balance transferred to their account) or for operational reasons (e.g., in group 2 some participants may have missed the deadline). We do not exclude these households from the sample in our main results but investigate in robustness checks how MPC estimates change when conditioning on take-up.[14] The fact that participants receive a reminder and are informed about the remaining balance through their phone app reduces the likelihood that they forget that they have available funds on the treatment card.

[13] The pin codes for the treatment cards are set by the bank to be the same as each participant's main debit card.

[14] Two participants filed a complaint stating that they did not receive a working card in time (they were issued a replacement several weeks later); we exclude them from the analysis.

Randomization Tests.—We implement statistical tests to assess the validity of the randomization protocol. Since the randomization was done at the level of the individual, but spending is observed at the level of the household, households with multiple eligible members will be over-represented in the treatment. We therefore conduct all our analysis within bins of households that have the same number of eligible members E (which we will refer to as "household size"), always comparing households with one treated member (we do not have households with multiple treated individuals) to households of the same size with no treated individuals. For the sake of brevity, we will refer to households with one treated individual as "treated households" and those without as "control households."

Online Appendix Figure D1 shows the results of randomization tests where we regress a treatment dummy on a set of standardized household characteristics and a set of dummies for the number of eligible individuals within the household. The coefficients on the household characteristics are all small and not statistically significant, indicating that the means of these characteristics are similar across treated and untreated households (within bins for the number of eligible individuals). These result confirm the validity of the experimental design.

II. Main MPC Estimates: Facts 1, 2, and 3

In this section, we report our main MPC estimates. We first consider all treatment cards at once (Section IIA). We then report estimates by card types (Section IIB), establishing our first three key facts about MPCs in this subsection. We also describe the participants' spending behavior by analyzing the composition of expenditure as well as auxiliary survey data (Section IIC). Finally, we report MPC estimates by framing group (Section IID).

A. *Pooled MPC Estimates*

We first present MPC estimates for all treatment cards, first presenting evidence from the raw data and then turning to a regression framework.

MPC Estimates from Raw Data.—Panels A and B of Figure 1 present the MPC estimates from raw data. Panel A first documents the timing of purchases that treated households make using the treatment card alone. The figure shows that average spending increases rapidly and reaches about €250 after two months; that is, participants spend 84 percent of the transfer within two months. However, this direct spending response may be offset by reduced spending in the households' main bank accounts.

To assess the magnitude of potential substitution effects, we plot the level of spending in each week in the treatment and control groups. Given that treatment was assigned at the level of eligible individuals, we implement one adjustment to the raw data, reweighing participants by the propensity score (i.e., so that the number of eligible individuals within households is the same across the control and treatment groups). Panel B of Figure 1 shows clear graphical evidence that treated households spend more upon receipt of the transfer, but the extra spending is short-lived, lasting about three to four weeks. A month after the start of the experiment, there

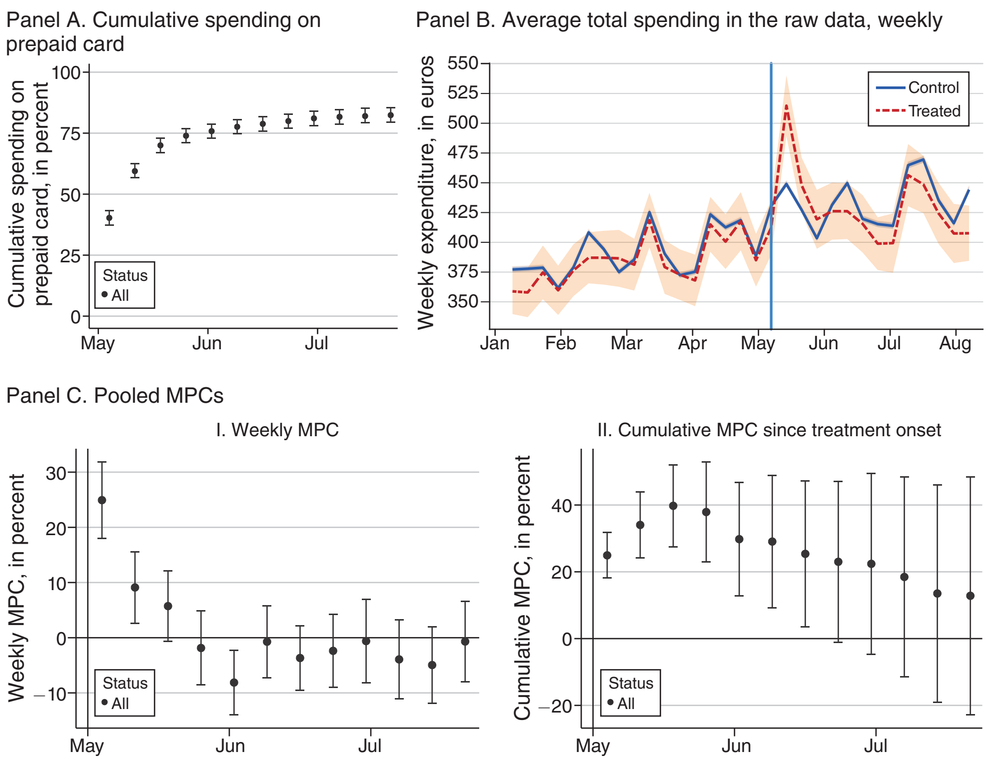

Figure 1. Pooled MPC Estimates

Notes: Panels A and B report the treatment effects in the raw data, plotting cumulative spending on the prepaid card in panel A for treated households and average weekly spending for control and treated households in panel B. The 95 percent confidence intervals for mean weekly spending are reported as shaded regions in panel B. Panel C reports the regression-based MPC estimates. Panel C.I reports the weekly estimates, while panel C.II depicts the cumulative effects. Ninety-five percent confidence intervals are reported, clustering the data at the household level.

is no evidence for any difference in spending patterns between treated and control households. Thus, the response is concentrated in the very short run, with little intertemporal substitution. This panel also shows that spending displays significant seasonality, which the control group allows us to address along with any other time variation in consumption.

Next, we move to a regression framework to provide more precise estimates of MPCs.

MPC Estimates from Regression Specification

Specification: Our baseline econometric specification to estimate consumption responses is a standard two-way fixed effect linear model:

$$Y_{it} = \sum_{\tau=0}^{\tilde{T}} \beta_\tau \mathbf{1}\{\tau \text{ } weeks\ since\ i\ treated\}_{it} + \alpha_i + \alpha_{tE} + \varepsilon_{it}, \tag{1}$$

where Y_{it} is the outcome variable (usually consumption spending of household i in week t), the dummy $\mathbf{1}\{\tau \text{ weeks since } i \text{ treated}\}_{it}$ is one if and only if i contains a treated individual and week t is τ weeks after the first treatment week (the week of May 2), α_i are household fixed effects, and α_{tE} are fixed effects for "week by number of eligible individuals within the household." Given that treatment is assigned at random across eligible individuals, we only need to control for α_{tE} to achieve identification,[15] but we also include household fixed effects to reduce noise. Standard errors are clustered at the household level.

Given that a control group of untreated households is available, our two-way specification is not subject to the "negative weights" issue analyzed in recent work on difference-in-differences design (e.g., De Chaisemartin and Haultfoeuille 2020; Borusyak, Jaravel, and Spiess 2023).

Results: The results are reported in panels C.I and C.II of Figure 1. Panel C.I reports the estimates for the β_τ coefficients at a weekly frequency after treatment. The panel shows that, on average, participants' spending increases by €75 in the first week, €25 in the second week, and €20 in the third week. The estimates are close to zero in the following weeks, indicating that the spending burst is concentrated in the short run.[16]

Panel C.II shows the point estimates and standard errors of the cumulative sum since the start of treatment. The point estimate for the cumulative average effect after four weeks is €113, corresponding to an MPC of $113/300 = 38$ percent. As in other papers that compare recipients of transfers with nonrecipients to estimate MPCs, estimates over longer horizons are becoming increasingly less precise as the variance of cumulative consumption increases for both treatment and control groups over time due to the presence of idiosyncratic shocks. We find that the decrease in point estimates over longer horizons in Panel C.ii of Figure 1 results from a small number of positive outliers in weekly preperiod consumption expenditures that push up the household fixed effects and make subsequent expenditures appear small in comparison. Online Appendix Figure D2 shows FGLS estimates that downweigh household groups with higher preperiod consumption volatility, and these estimates are flat over the corresponding horizon.

To investigate the dynamics of the cumulative MPC over longer horizons, we estimate a specification analogous to (1) except that to increase parsimony we use a seventh-order polynomial to model the weekly MPC response after treatment. Figure 2 plots the results over two quarters, showing that the cumulative MPC increases very quickly in the first few weeks and remains stable from the first month onward. The point estimates are statistically significant at the 95 percent level for the first five months. If spending increased at the same rate as during the first three weeks, the cumulative MPC would reach €300 around mid-June, which we reject in Figure 2. This analysis thus confirms that the increase in spending is concentrated in the short run.[17]

[15] In practice, the estimates remain similar when we do not include this control.

[16] Consumption appears to fall in the first week of June, raising the possibility of intertemporal substitution. We discuss cumulative MPC estimates below, finding little evidence of intertemporal substitution with our preferred feasible generalized least square (FGLS) estimates.

[17] Accordingly, in much of the analysis that follows, we choose the one-month MPC as our main focus.

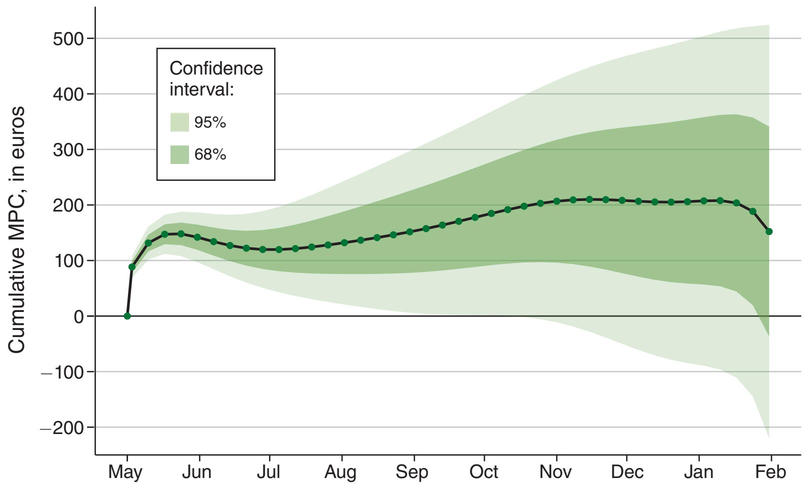

Figure 2. Long-Term MPC Estimates

Notes: In this figure, we run a specification analogous to (1) except that to reduce noise we use a seventh-order polynomial to model the weekly MPC response after treatment: $Y_{it} = \sum_{k=1}^{8} \beta_{\tau}^{k-1} \cdot \tau_{it}^{k-1} + \alpha_i + \alpha_{tE} + \varepsilon_{it}$. We still view this specification as nonparametric estimation of the MPC given the flexibility of the seventh-order polynomial. To reduce noise further, we use the same FGLS procedure as in online Appendix Figure D2. The figure reports the cumulative MPC and both the 95 percent and 68 percent confidence intervals clustered at the household level.

Finally, it is worth noting that we of course lack the statistical power to detect long-term changes in consumption of a few euros per month. For instance, a permanent income consumer would consume the annuity value of the €300 transfer: with a 2 percent interest rate, this corresponds to an increase in consumption of about €6 per year, or 0.50 cents per month, which we cannot detect in the data. Our empirical results are therefore not inconsistent with theoretical reasoning based on an intertemporal budget constraint and a transversality condition implying a long-term MPC of one.

Additional robustness results are reported in online Appendix E.1, analyzing observed savings at the bank, alternative specifications, take-up, alternative consumption measures, reweighting, winsorization, and showing bootstrapped confidence intervals.

B. *MPC Estimates by Card Type*

Next, we analyze MPC by card type and establish the key result of the paper: the MPC is larger when treatment cards have negative interest rates.

The estimates are reported in Figure 3 separately for the three treatment groups. Panel A shows estimates of the β_{τ} for households in group 1 with no restrictions. Card 1 leads households to increase their weekly consumption spending in the two weeks after treatment by about €40; the point estimates thereafter are close to zero and not significant. Panel B shows that households in group 2—which receive a card that expires after three weeks—increase their weekly consumption significantly for the first three weeks after treatment, by about €65 in the first week and by about

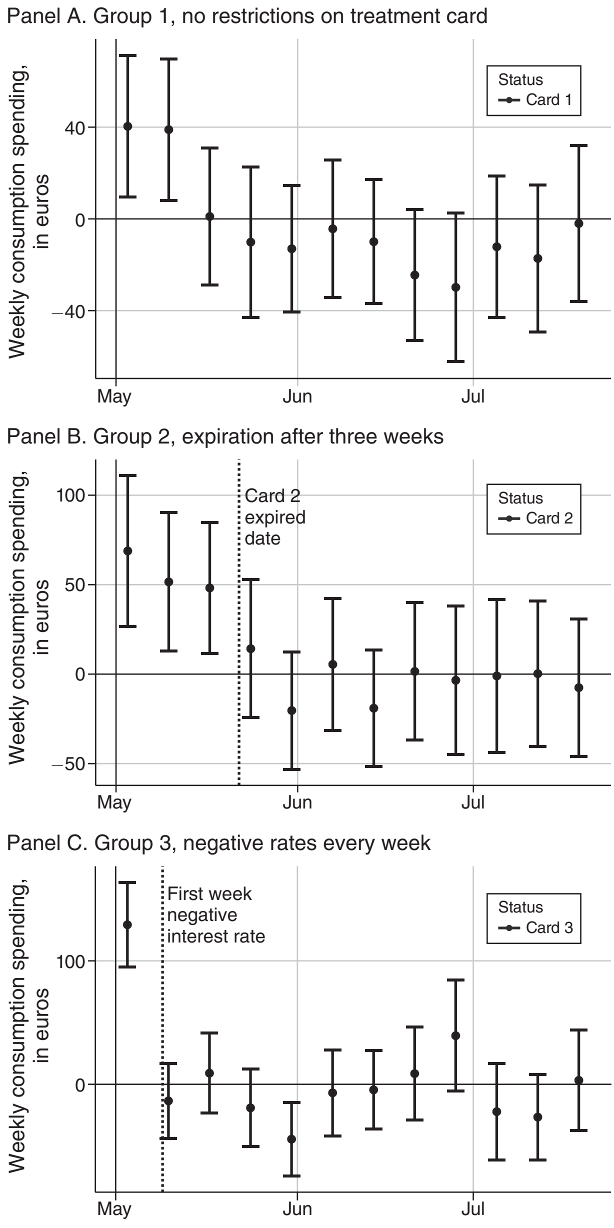

Figure 3. MPC by Card Type, Weekly

Notes: This figure reports MPC estimates depending on the card type. Panel A reports the weekly estimates for group 1, panel B for group 2, and panel C for group 3. Card 1 has no restrictions, while card 2 expires three weeks after the onset of the experiment, and card 3 applies a negative interest rate on the remaining balance every Monday at 11:59 PM. Ninety-five percent confidence intervals are reported, clustering the standard errors at the household level.

€50 in the second and third weeks. There is no sign of intertemporal substitution as estimates hover around zero after the third week. Finally, panel C shows the response for households in group 3—with the negative interest rates—which increase their

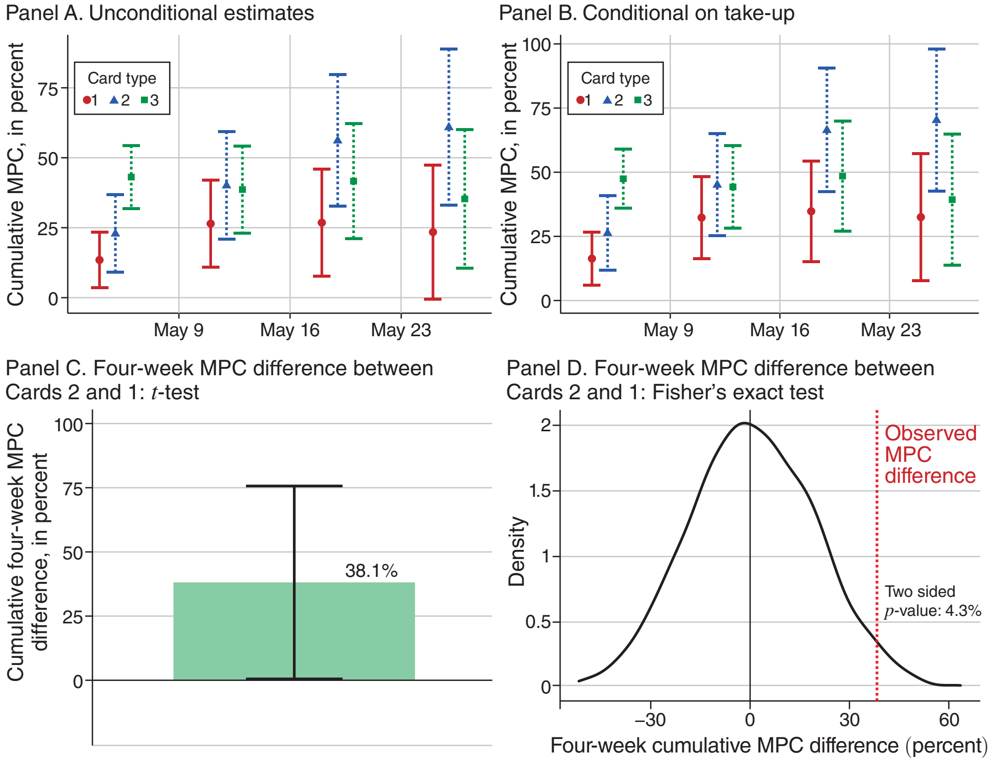

Figure 4. MPC Estimates by Treatment Group

Notes: Panels A and B of this figure report cumulative MPC estimates depending on the card type. Card 1 has no restrictions, while card 2 expires three weeks after the onset of the experiment, and card 3 applies a negative interest rate on the remaining balance every Monday at 11:59 PM. Panel A includes treated households that do not use the card in the treatment groups; panel B does not. Ninety-five percent confidence intervals are reported, clustering the data at the household level. Panels C and D report statistical tests to estimate the difference in cumulative MPCs after four weeks for cards 1 and 2. We estimate specification (1) with OLS and compare the sum of the coefficients for the first four weeks for cards 1 and 2. Panel C reports the point estimate and 95 percent confidence interval, while panel D reports the result obtained with randomization statistical inference (Fisher's exact test). The distribution shown in panel D is the distribution of estimated MPC differences where placebo-treated households are drawn from the population of untreated households.

spending immediately in the first week of the experiment by about €130 but not thereafter.[18]

Panels A and B of Figure 4 report the cumulative spending response. The figure shows that the cumulative MPC for group 1 is much lower than for groups 2 and 3. After four weeks, the cumulative MPC for group 1 is €70 (23 percent), compared with €183 (61 percent) for group 2 and €106 (35 percent) for group 3. Panel B of Figure 4 shows MPC estimates conditional on using the card (at some point in the entire sample) to make purchases. The point estimate for the average four-week MPC

[18] The time pattern for this group suggest that most households understood correctly the terms of the card as they tried to spend it down before the date of the first negative interest rate. The terms were explained in simple language in the instruction letter, including simple examples and explicitly stating that if the household spent the full amount of €300 before the date of first deduction, no money would be lost. We cannot however know for certain that all households understood the terms correctly, especially among those who did not spend down the card.

of group 2 participants is about 10 percentage points higher, at 70 percent. In both figures the consumption response to the stimulus transfer is substantially higher for group 2 compared to group 1, indicating that stimulus design choices can affect the MPC.

Since the confidence intervals in panels A and B of Figure 4 overlap, it is important to assess formally whether there is a statistically significant difference by card type. Focusing on the cumulative MPC after four weeks, panels C and D report a significant difference between card 2 and card 1, which we view as the key result of this paper. Both conventional *t*-tests and exact randomization tests show a significant difference at the 5 percent confidence level. Online Appendix Figure D3 reports the results for cards 1 and 3 for which we do not find statistically significant differences after four weeks. Note, however, that heterogeneity in the iMPCs is present: an *F*-test of the hypothesis of the same path of MPC across all treatment groups for the first twelve weeks is rejected with a *p*-value of 2.5 percent.[19]

Online Appendix Table D3 reports the differences in cumulative MPCs by card type after four, eight, and twelve weeks using both OLS and FGLS. Comparing card 2 and card 1, we find that the difference in point estimates grows larger after two and three months but is no longer statistically significant at conventional levels, with *p*-values around 0.11–0.15. Furthermore, we do not find statistically significant differences between cards 1 and 3 or cards 2 and 3 at any horizon, although the difference in point estimates grows larger with time. Thus, an important direction for future research would be to scale up sample sizes to estimate more precisely differences by card types, especially at long horizons.

Additional results are reported in online Appendix E.2, analyzing long-term MPC dynamics, spending on the prepaid card, predictors of take-up or lost funds, and showing bootstrapped confidence intervals.

Taking Stock.—We can now summarize our first three key facts about the magnitude of the MPC by card type—including our main result, Fact 2—and its time profile:

FACT 1: *The average one-month MPC on a cash-like transfer is 23 percent.*

FACT 2: *The design of the stimulus transfer can substantially affect the MPC: the average one-month MPC out of a prepaid card whose remaining balance expires after three weeks is 61 percent.*

FACT 3: *The increase in consumption is much larger early on in the first two to three weeks.*

[19]There is a statistically significant difference between card 1 and card 2 despite the overlap in the 95 percent confidence interval in Figure 4 because of the covariance between estimators. Specifically, estimating equation (1) by OLS, we find that the correlation between the cumulative four-week MPCs of card 2 and card 1 is 0.62. Intuitively, given the skewness of consumption expenditures, large purchases in the control group during the treatment period can have a sizable impact on the point estimates for all treatment cards.

C. *Understanding the Spending Response*

To better understand the participants' spending behaviors, we combine two approaches: survey questions to the treatment group[20] and an analysis of the spending categories for treatment cards and linked bank accounts. We first analyze the patterns for all cards and then study the three types of cards in turn.

All Cards.—The results for all cards are presented in Figure 5. Our analysis delivers three takeaways. First, survey responses show that participants are well aware that they spend less on their main account and use the treatment card to substitute for regular spending. They mention precautionary savings as a key motive for the money they saved out of the transfer (panel A of Figure 5), and they report that they use the treatment card primarily to cover running expenses (panel B).

Second, we use the treatment card and the bank data to analyze the composition of expenditures. For each transaction our data contain the four-digit MCC that is associated with the vendor. Panel C of Figure 5 shows that treated households spend more on clothing and household equipment (furniture, consumer electronics, etc.). Panel D breaks down the purchases on the treatment cards by category, confirming the importance of spending on clothing and electronics.

Next, we examine whether there are significant differences in terms of spending on durables.[21] Extending the product classification from Ganong and Noel (2019), we classify MCC codes into one of four spending categories used by the French National Statistical Institute (INSEE): nondurables (including food and drink, fuel, and items that depreciate quickly), semidurables (including, notably, apparel, footwear, and other textiles), durables (furniture, electronics, and durable household equipment, as well as leisure items and cars), and services. Online Appendix Table D4 provides examples of products belonging to each of these categories. We also build a crosswalk to the French input-output table to assess the import content of households' consumption baskets. We find that the spending share on durables increases to 10 percent upon treatment, relative to about 7 percent in the control group, as reported in panel E of Figure 5. Online Appendix Figure D4 reports the cumulative MPC response separately for durables and nondurables, showing a more sustained increase in spending on nondurables over time.

Finally, we analyze the propensity to spend on imports. We measure imports by mapping the MCC codes to the French input-output tables, which provide import penetration rates across categories. We find that treated households make purchases that have on average a higher import content, resulting in an increase in the weighted average import share from 7 percent in the control group to 9 percent in the treatment group (panel F). While these differences are statistically significant, they are modest from an economic perspective.

[20] The survey was administered via the implementation partner's web platform. The survey response rate is 46 percent.

[21] Of course, at a high frequency (e.g., weekly) it is difficult to draw distinctions between durables and nondurables—to a certain extent, all products can be viewed as durables.

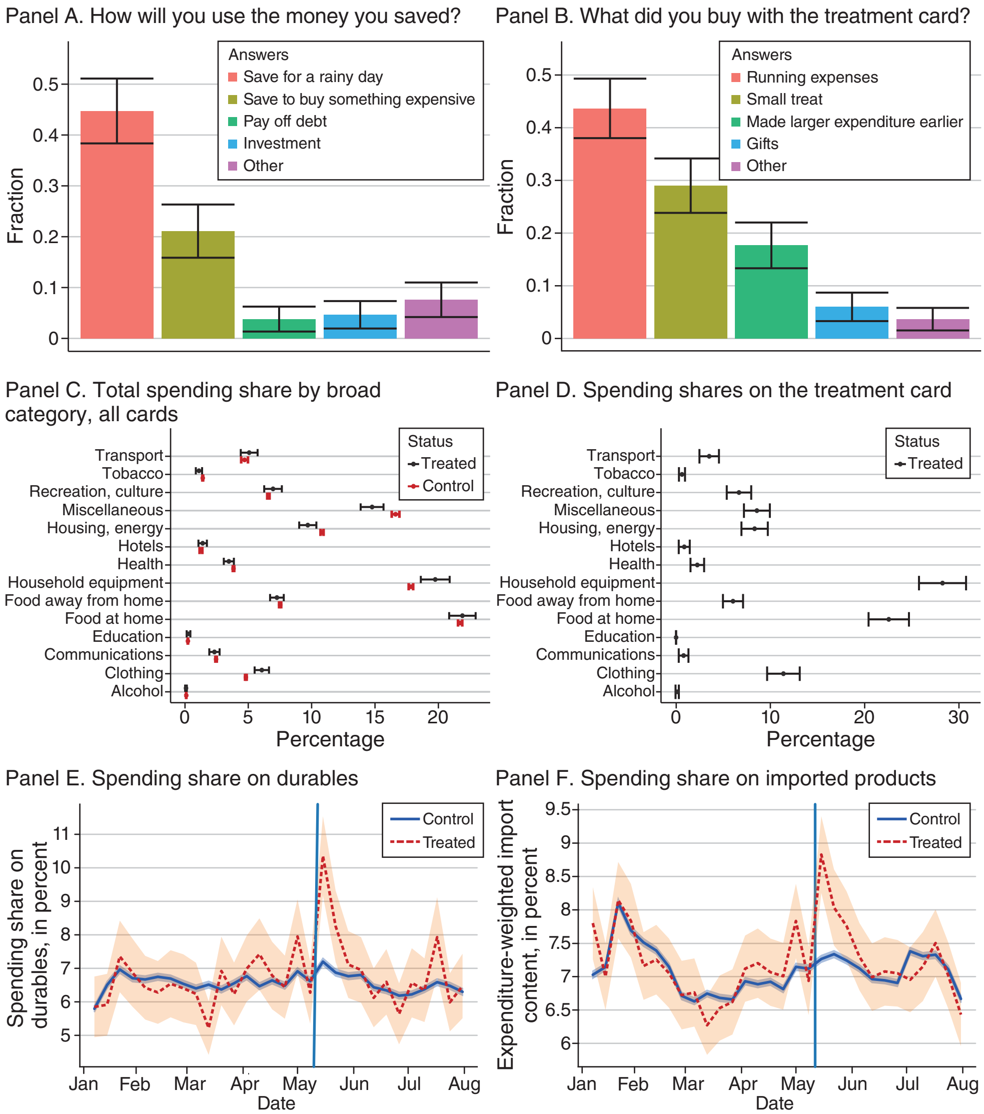

Figure 5. Understanding Participants' Spending Behavior, All Groups

Notes: Panels A and B of this figure report the answers of participants to survey questions. The other panels use the bank data to document the expenditure patterns of the treatment and control groups across product categories. Panel C shows expenditure shares in the total expenditure basket, panel D shows expenditure shares using the treatment cards only. Panel E shows the weekly average expenditure share on durables (as defined in the BEA personal consumption expenditures classification) for treatment and control groups; panel F shows the import content of households' expenditure baskets. The import content for each household's consumption basket is the expenditure-weighted industry-level import content. The industry-level import content has been constructed using INSEE's input-output tables for France.

By Card Type.—We repeat the analysis by card type in online Appendix Figure D5. We first rely on the survey results and find that households in groups 2 and 3 report that they are less likely to cover running expenses and more likely to make large purchases earlier, consistent with the higher MPC estimates in the data.

Table 2—Decomposition of Four-Week MPC and Expenditure Shares by Type of Expenditure (Percent)

	Decomposition of four-week MPC			Expenditure shares over four weeks	
	Card 1 (1)	Card 2 (2)	Card 3 (3)	Treatment group (4)	Control group (5)
Nondurables	18.8	9.2	26.6	29.5	30.5
Semidurables	49.0	30.9	28.6	19.8	17.9
Durables	24.1	19.1	49.8	9.6	7.8
Services	13.6	54.2	11.6	36.6	38.9
Not categorized	−5.6	−13.5	−16.9	4.5	4.7

Notes: This table reports the average four-week MPC on the row category divided by the total four-week MPC (any consumption expenditure), for treatment groups 1, 2, and 3. Columns 4 and 5 report expenditure shares for the treatment and control groups over the four weeks following treatment. The last row refers to expenditures that cannot be classified into the four main product categories, for example, cash. Columns do not sum up exactly to 100 percent because of rounding. When computing the estimates for this table, we do not winsorize consumption expenditures so that we can decompose the total consumption response across product categories.

Second, we decompose the consumption expenditure increase by durability. Table 2 shows the estimated fraction of the expenditure increase on each of the four durability categories by dividing the four-week cumulative point estimate of a regression of consumption expenditure in the row category on time-since-treatment dummies (and fixed effects as in the baseline specification) by the corresponding four-week cumulative point estimate in a regression of all consumption expenditures (the baseline specification).

The results show that households that receive treatment card 2 channel a substantial fraction of the additional expenditure into personal services, whereas card 3 households see a disproportionate increase in durables purchases. Comparing the three types of prepaid cards, Table 2 shows that the short-run spending response is not driven by spending on durables as card 3 alone features a substantial increase on this category. Additional results on the composition of spending by card type are reported in online Appendix E.2, analyzing the response of durables over a longer horizon and using an alternative classification of products into durables and nondurables.

A potential concern is that the higher overall spending response with cards 2 and 3 might come at the expense of the "quality" of spending. For example, a recent study by Jaroszewicz et al. (2022) finds that unconditional cash transfers taking place during COVID-19 sometimes had detrimental effects on recipients' self-reported measures of well-being in a sample of about 5,200 US households living in poverty. Other papers have documented that consumption opportunities may lead to a "consumption binge" with the potential to reduce welfare in the long run (Garber et al. 2022). We study survey and spending outcomes to understand whether our transfers could have caused harm to some participants.

We first examine whether the spending share on goods that can be deemed to have "negative externalities" (drinking, tobacco, gambling, and lottery products) differs across treatment arms. Online Appendix Figure D5 shows that there is no significant difference across treatment groups. Second, we analyze whether participants of groups 2 and 3 experience a fall in nondurable consumption or higher volatility, which could be caused by an initial consumption binge. We reject this hypothesis: participants in groups 2 and 3 spend more on nondurables in the short run and experience no fall in the longer run (see online Appendix E.2). Finally, we use the survey to elicit the subjective impact of the transfer. In response to the question: "Has the transfer of the €300 card increased your happiness?" only eight out of 391 respondents (or 2 percent) report that the transfer has not at all increased their happiness. Ninety-two percent of respondents respond that the transfer has either "very strongly" or "somewhat" increased their happiness. We therefore conclude that it is very unlikely that our implementation design choices have caused harm, while they led to a large increase in MPCs.

D. *MPC Estimates by Framing Group and Experimenter Demand Effects*

Finally, we evaluate whether households have different average MPCs depending on the framing of the intervention. This framing treatment may be of substantive interest insofar as policymakers can frame household transfers through public discourse or in official letters to households. The results are also informative about experimenter demand effects, or Hawthorne effects. Indeed, a potential concern about our experiment is that some households may feel compelled to act according to what they perceive to be the goal of the experiment. However, our framing treatment makes it possible to assess whether Hawthorne effects are likely to drive our results since only the participants in the framing group are explicitly told that they are expected to spend quickly and increase their total spending instead of covering running expenses.[22]

We find that households that received the additional framing treatment with a paragraph encouraging them to spend the money quickly on local goods or services have very similar average consumption expenditures overall (difference in MPC $<$ 10 percentage points and not statistically significant) as households that did not receive the framing treatment (online Appendix Figure D6). We also examine whether the composition of expenditures varies across groups, finding no difference. For example, spending on imports is similar across the two framing groups (online Appendix Figure D7).

Thus, we conclude that implementation design choices are powerful tools to increase the MPC, while written framing treatments are unlikely to have large effects. These results also show that Hawthorne effects are unlikely to drive our results.

III. MPC Heterogeneity across Households: Facts 4 and 5

We now turn to the analysis of MPC heterogeneity across households. Estimating MPC heterogeneity is key both for policy—as policymakers may wish to target certain

[22] See the translation of the letter in online Appendix C.1. All participants are told that "The objective of this initiative is to study, within the framework of a policy aimed at promoting economic recovery, people's spending behaviors when a sum of money is distributed to them for free." Participants in the framing group are also told the following: "Although you are free to use the amount of €300 as you wish, we invite you to: spend the money as quickly as possible; buy products made in France and services that support local employment, as the objective of this transfer is to stimulate the French economy by encouraging the consumption of products made in France; purchase products or services that you wouldn't normally buy (other than your regular expenses) to increase your total spending and thereby contribute to the economic recovery, rather than covering expenses that were already planned."

households to maximize the aggregate MPC—and for macroeconomics models —as MPC heterogeneity is a useful moment to assess the accuracy of the predictions and potentially falsify certain models. We first document MPC heterogeneity by observable household characteristics (Section IIIA), establishing our fourth key fact about MPCs. Finally, we present estimates of the unconditional distribution of MPC across households with a deconvolution approach (Section IIIB), our fifth key fact about MPCs. We discuss the implications of our findings in Section IV.

A. *Heterogeneity by Observable Household Characteristics*

To examine the importance of various observable household characteristics in predicting treatment effect heterogeneity, we first use a simple OLS specification and then turn to a machine learning (LASSO) analysis.

OLS Analysis.—We first estimate differences in the MPC for households with different characteristics. Specifically, we estimate specifications of the form:

$$(2)\quad Y_{it} = \sum_{q=1}^{4} \sum_{\tau=0}^{\tilde{T}} \beta_{\tau}^{q} \mathbf{1}\{\tau \text{ weeks since } i \text{ treated}\}_{it} \mathbf{1}\left\{X_i \in Q_q^X\right\}_i + \alpha_i + \alpha_{tEQ_q^x} + \varepsilon_{it},$$

where Q_1^X to Q_4^X are the quartiles of the distribution of the time-invariant household characteristic X.

We consider six characteristics: net liquid wealth, net illiquid wealth, average pretreatment consumption (as a proxy for permanent income), income, age, and gender. The first four variables are motivated by macroeconomic models, which make predictions about heterogeneity in the MPC by net wealth (e.g., Kaplan and Violante 2014) and by current or permanent income (e.g., Straub 2019); we further discuss the relationship between our findings and these models in Section IV. In addition, we consider age and gender as these characteristics are easily observed and could in principle be used to target transfers toward certain populations.

The variables are built as follows. Net liquid wealth corresponds to the sum of household-level current accounts and net liquid saving deposits at the bank (factoring out short-term debt, such as consumer debt). We measure this variable at the onset of the experiment, on the first day of May, to capture the liquid funds actually available to households. Net illiquid wealth captures the sum of illiquid savings, asset level, and mortgage debt for the household at the bank level. Average pretreatment consumption is measured as the average monthly consumption expenditure in the year prior to treatment at the household level. Finally, we define income at the household level as the sum of all incoming transfers.[23] Except for net liquid wealth, which is measured at the beginning of the experiment, and average preperiod consumption, which is computed as an average over a year, these variables are averages over the monthly levels in the six months prior to the experiment. Regarding age and gender, the characteristics pertain to the eligible household member.[24]

[23] Inflows above €15,000 are trimmed out at the household level.

[24] When there are multiple eligible household members in the control group, we pick one of the eligible members at random and use their characteristics. For treated households, age and gender are taken from the selected

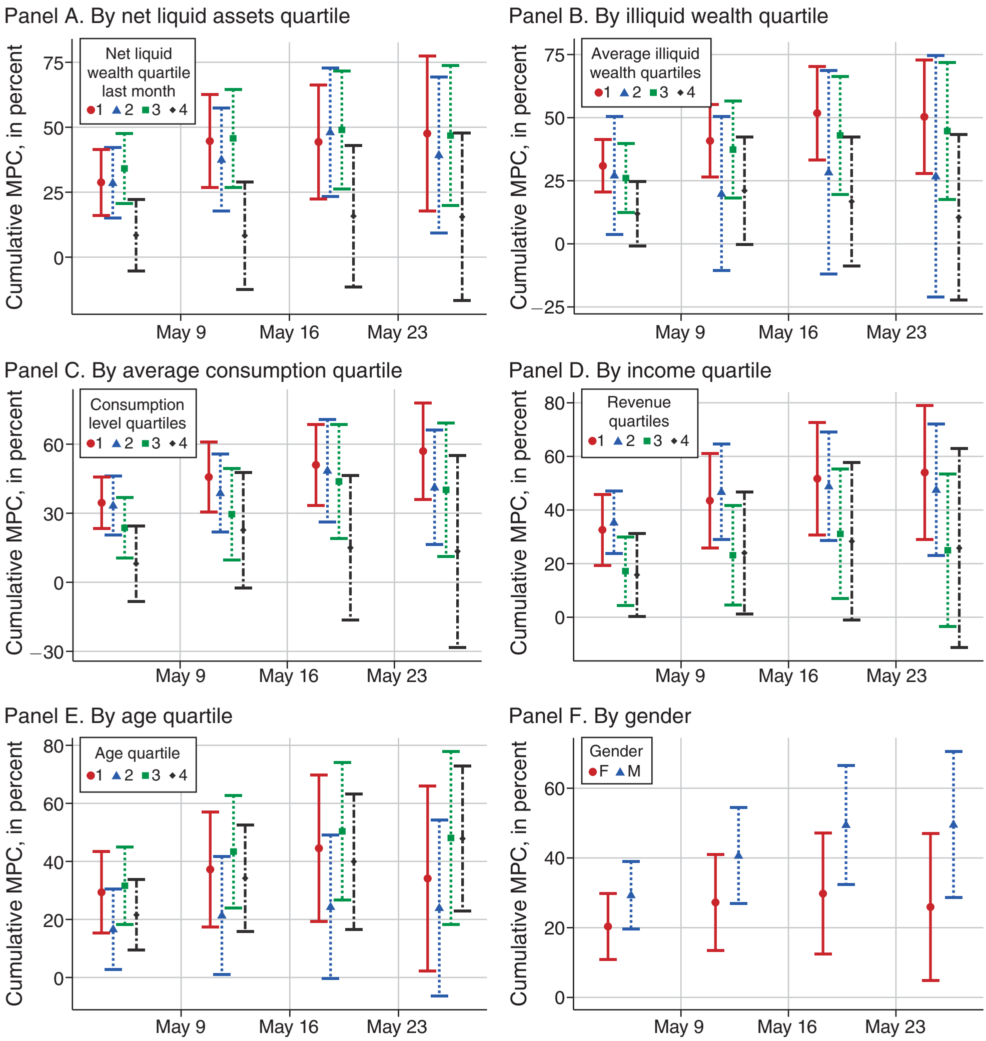

Figure 6. MPC Heterogeneity by Observable Household Characteristics

Notes: This figure reports MPC estimates depending on observable household characteristics. We document heterogeneity in turn by net liquid wealth, illiquid wealth, average consumption prior to the experiment (as a proxy for permanent income), income, age, and gender. Ninety-five percent confidence intervals, with standard errors clustered at the household level, are reported in all panels.

Figure 6 reports the results, plotting cumulative MPCs across household groups. We first consider the role of liquid and illiquid wealth in panels A and B. Panel A shows that MPCs fall with the level of net liquid wealth. Although the standard errors are sizable, there appears to be a negative relationship between the level of liquid wealth and the MPC. We obtain similar results when liquid wealth is measured as an average over six months prior to the experiment rather than at the beginning of

individual. We are unaware of prior results on MPC heterogeneity by gender. As a result, heterogeneity regressions by gender were not specified in the pre-analysis plan.

the experiment. Panel B turns to illiquid wealth, depicting a negative relationship between MPCs and illiquid wealth quartiles. Despite these negative relationships, online Appendix Figure D8 shows that the MPC remains high even for households that have substantial liquid wealth, a fact we will use later on when drawing implications of our findings for consumption models. Online Appendix Figure D8 also shows that the results are similar when using current account funds alone as the measure of liquid wealth.

Next, we turn to income in panels C and D. We first consider our proxy for permanent income: average consumption prior to the experiment. Panel C shows that MPCs tend to be lower for households with higher levels of consumption prior to the experiment. Similarly, panel D reports that MPCs fall with household income.

Finally, panels E and F consider in turn age and gender. Panel E shows that older households appear to have larger MPCs. Turning to gender, panel F shows that women have a much lower MPC than men. After a month, the cumulative MPC is close to 50 percent for men and only about half as large for women.

In sum, income and gender constitute the strongest sources of observable MPC heterogeneity among the predictors we consider. Online Appendix E.3 reports additional results on liquidity and gender, FGLS specifications, and the statistical precision of the estimates.

LASSO Analysis.—We now turn to a set of regressions that attempt to uncover which household characteristics are most relevant for explaining MPC heterogeneity after four weeks. We implement specification (2) with all six variables (divided into quartiles when relevant) included jointly, as well as some additional variables (unemployment, local area characteristics, and household size). In order to avoid overfitting, we estimate the coefficients using a LASSO estimator for varying levels of the regularization parameter. Although these results do not isolate causal links, they reveal which variables are the most important predictors of MPC heterogeneity.

Figure 7 shows the results of our estimates on the entire sample of treatment and control group participants. We find that the most important variables to predict treatment effect heterogeneity are demographic characteristics—specifically, gender, high-age dummies, household size, and the location characteristic (urban versus rural; the omitted category is semi-urban)—as well as the dummy that captures the top quartile of average past consumption (our proxy for permanent income) and the third quartile of liquid wealth. Conditional on these variables, other characteristics contribute little to predicting treatment effect heterogeneity. Perhaps surprisingly, income and wealth (whether liquid or illiquid) have little predictive power to explaining MPC heterogeneity. The results also clearly show how the variables that capture variation in the treatment design—the treatment group dummies—stand out in explaining treatment effect heterogeneity.

Online Appendix E.3 reports complementary results, repeating the LASSO analysis with treatment group 1 alone as well as at a longer horizon.

While leading macroeconomic models highlight the role of liquid and illiquid wealth as key predictors of treatment effect heterogeneity, our LASSO analysis shows that other predictors are more powerful. We further discuss the implications of these results for household targeting in Section IVB.

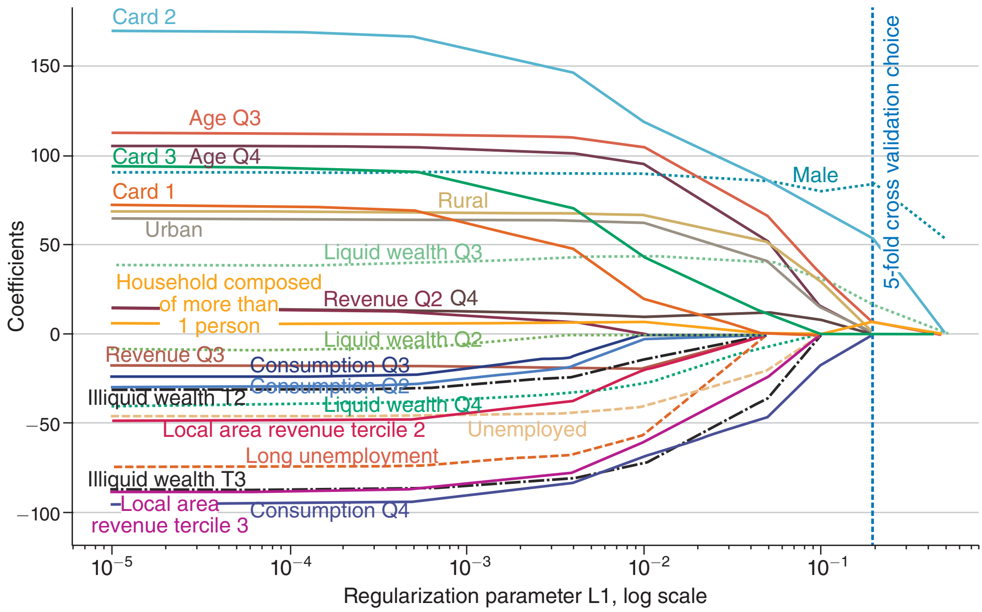

Figure 7. LASSO Estimates of Four-Week MPC Heterogeneity

Notes: The figure shows LASSO estimates of coefficients of interactions of the respective characteristic with a treatment dummy in specification (2) for varying regularization parameters (horizontal axis). We predict the cumulative MPC after four weeks. The dashed vertical line shows the regularization parameter chosen by five-fold cross validation.

Takeaways.—The OLS and LASSO results together establish our fourth key fact:

FACT 4: *MPCs vary with observed household characteristics, notably by gender and proxies for permanent income, and are high even for the liquid wealthy.*

B. *Unconditional Distributions of MPCs*

We now proceed to estimating the unconditional distributions of MPCs across households regardless of observable household characteristics. In an experimental setting like ours, the entire distribution of the outcome variable is different from that of the control group only because of the treatment effect. When the treatment effect is independent from the error term, we can therefore recover the full distribution of the treatment effect using statistical deconvolution techniques.

Setting, Identification, and Estimation.—We consider the model

$$Y_{it} = \sum_{\tau=0}^{\tilde{T}} \beta_\tau \mathbf{1}\{\tau \textit{ weeks since i treated}\}_{it} + \alpha_i + \alpha_{tE} + \varepsilon_{it},$$

where now, in contrast to the previously studied model, we assume that the β_τ are stochastic, with $\beta_\tau \sim F_\tau$. We further assume that the β_τ are independent from ε_{it}; we discuss and test this key assumption at the end of this section. As before, the treatment dummies are independent from the errors ε_{it} as well as from the β_τ due to the experimental design. We seek to recover the distribution of $\sum_{\tau=0}^{\tilde{T}} \beta_\tau$, which corresponds to the $\tilde{T}$-period marginal propensities to consume. Under the assumptions stated above, the distributions F_τ and therefore the distribution of the $\tilde{T}$-period MPC are identified under no parametric assumption.

The model thus takes the same form as a classic measurement error model (see Schennach 2016 for a survey), and the distribution of the β_τ can be estimated using a deconvolution method: we first estimate the distribution of ε_{it} from the population of untreated households, and we then deconvolve that distribution from the distribution of the dependent variable of the treated at time of treatment. Intuitively, apart from the fixed effects α_i, α_{tE}, the distribution of outcome variables for treated and untreated households differ *only* because of the presence of the treatment effect terms β_τ. Under the assumption that the β_τ are independent from ε_{it}, we can recover the distribution F_τ. Note that the assumption that the treatment and control groups have identical distributions of error terms ε_{it} would be difficult to defend in a nonexperimental setting. To avoid confounding treatment effects with the potentially different tails of households of different sizes, we conduct the exercises in this subsection on households with one eligible member only.

We implement this approach in a two-step procedure. In the first step, we estimate α_i and α_{tE} from the set of observations (i,t) where either i is not in the treatment group or i is in the treatment group but has not been treated yet, in the spirit of Borusyak, Jaravel, and Spiess (2023).[25] In the second step, we construct cumulatives of de-meaned weekly consumption expenditure:

$$C_{it}^{\tilde{T}} = \sum_{\tau=0}^{\tilde{T}} \left(Y_{it} - \hat{\alpha}_i - \hat{\alpha}_{tE}\right).$$

We estimate the distributions of $\sum_{\tau=0}^{\tilde{T}} \beta_\tau$ through deconvolution, constraining the distribution of the estimand to have only positive support. This constraint is motivated by the fact that we find no evidence for a fall in consumption anywhere in the distribution, as shown in online Appendix Figure D9, which reports the quantile treatment effects (i.e., the differences in the quantiles of the distributions of $C_{it}^{\tilde{T}}$ for treated and untreated households) over 4-week, 8-week, and 12-week horizons. The figure shows that the left tail of the distributions of cumulative de-meaned consumption is the same for treated and control, implying that the treatment effect distributions do not have mass on the negative part of the real line.[26]

We use the flexible quadratic-programming-based estimation procedure proposed by Yang et al. (2020), which, compared to standard Fourier-based methods, has the advantage that it also allows the density to be restricted to be nonnegative on its support and to integrate to one—restrictions that we also impose. Since deconvolution estimates often suffer from from oscillating densities in the tails, Yang et al. (2020)

[25] We first estimate household effects α_i from all pretreatment observations; then, conditional on these estimates, we estimate α_{tE} from control group observations. We choose this sequential procedure to avoid asymmetries across treatment and control groups in the precision of $\hat{\alpha}_i$; that is, we ensure that we have the same number of observations to estimate household fixed effects in treatment and control groups.

[26] We emphasize, however, that this nonnegativity constraint is not necessary to successfully apply the deconvolution procedure and that our results are very similar when not applying this constraint, as discussed below.

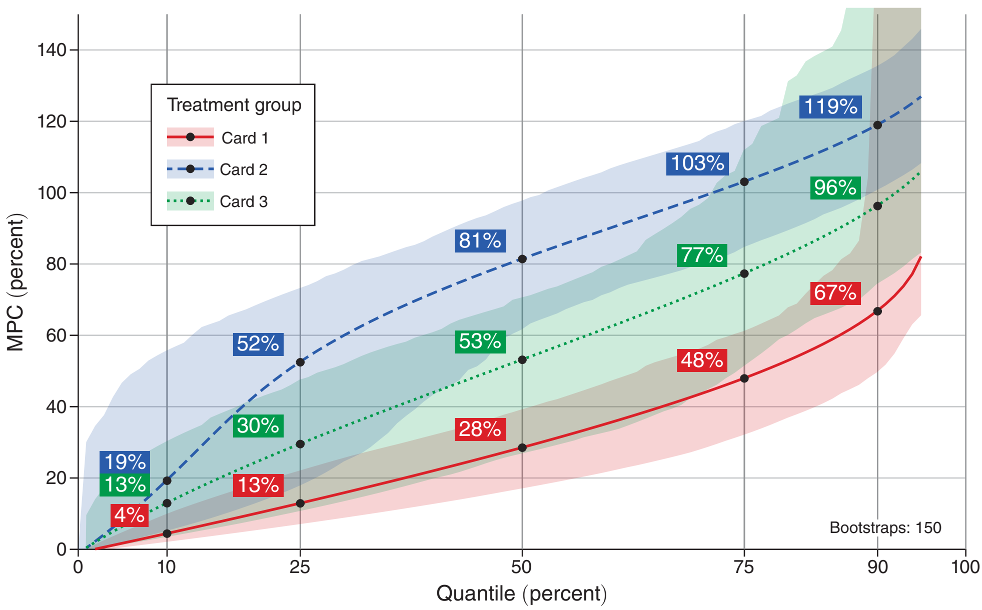

Figure 8. Household-Level Quantiles of the Four-Week MPC Distribution

Notes: This figure reports the quantiles of the distribution of four-week treatment effects by treatment group. Shaded regions are delineated by the tenth and ninetieth percent quantile of the bootstrapped simulated distribution of the corresponding moment.

recommend regularizing the density estimates through a penalty term. We follow this suggestion and penalize oscillations by adding a weighted finite-difference estimate of the second derivative of the density with a small penalty weight ($\lambda = 10^{-5}$). Since deconvolution estimators tend to be sensitive to noise in the tails of outcome distributions, we winsorize weekly consumption expenditure at the 90th percentile. We obtain standard errors for the estimated quantiles of the treatment effect distribution through a bootstrap of the entire two-step procedure.

Results.—Figure 8 shows estimates for the distribution of four-week MPCs by treatment group. The median MPCs are close to but slightly higher than the average treatment effect estimates we obtained in Section IIB for each card type, at 28 percent, 81 percent, and 53 percent for groups 1, 2, and 3, respectively. The estimates show a substantial heterogeneity in the propensity to consume out of the transfer, with the bottom quartile having a four-week MPC of less than 13 percent (card 1), 52 percent (card 2), and 30 percent (card 3), while the top quartile has MPCs above 48 percent (card 1), 103 percent (card 2), and 77 percent (card 3). The distribution of estimated treatment effects for group 2 first-order stochastically dominates the distribution of group 1.

These results establish our fifth key fact:

FACT 5: *There is substantial heterogeneity in the unconditional MPC out of a windfall transfer, and a large fraction of households has a high MPC.*

Fact 5 relates to existing papers that estimate the distribution of MPCs. Misra and Surico (2014) compare spending distributions of US households around the 2001 and 2008 tax rebates using quantile regressions and data from the Consumer Expenditure Survey. In contrast to our results, they find that significant shares of households experience negative treatment effects. The fact that the left tails of the spending distributions of treated versus untreated households are very similar (online Appendix Figure D9) is difficult to reconcile with negative MPCs in our data. Lewis, Melcangi, and Pilossoph (2019) use clustering-based Gaussian Mixture linear models to estimate MPC heterogeneity following the 2008 tax rebates. In their model the distribution of MPCs is parameterized to be discrete: MPCs are fixed but vary across groups; group memberships and MPCs for each group are identified through parametric assumptions on the error terms. In contrast to their approach, our treatment effect distribution is identified and estimated entirely nonparametrically. Similarly to us, they find that MPCs are ranging from close to zero to above one hundred percent and that most of the variation in MPCs is unexplained by observed characteristics. Our results do not, however, indicate the presence of discrete mass points in the distribution of unconditional MPCs.

We report additional results in online Appendix E.4, analyzing a model linear in log-consumption, dropping nonnegativity constraints, pooling treatment cards 2 and 3, and estimating quantile treatment effects.

Robustness.—The key assumption for identification of the treatment effect distribution is that the error term is independent from the treatment effect distribution. We now assess the plausibility of this assumption with several robustness checks.

First, the independence assumption may be violated if, for example, certain subgroups of the population (say, poorer households) that have a higher average treatment effect also happen to have systematically different errors terms ε_{it} shortly after the experiment was conducted (for example, because of calendar events such as bank holidays, which are common in May in France where poorer households may increase spending less than others). While we also cannot directly test the independence assumption, since both the treatment effect and the error terms are unobserved, we can perform a falsification test. We conduct an exercise where in the first step of the estimation procedure we project consumption on household fixed effects and week fixed effects interacted with (a, i, c, l, g) fixed effects, where a, i, c, and l are age, income, consumption, and liquid assets quartile bins and g is a gender dummy (instead of projecting it on just household and week fixed effects). The resulting estimates of the treatment effect distribution, shown in online Appendix Figure D10, remain virtually unchanged. Therefore, for our results to be biased, unobservable predictors of MPC heterogeneity should be much more strongly correlated with unobserved shocks ε_{it} than observable predictors. This sensitivity test, in the spirit of Altonji, Elder, and Taber (2005) and Oster (2019), lends support to our baseline estimates.

Second, another possible scenario that would violate the assumption of independence of treatment effect and error term is that some households may have more volatile consumption than others and may also have a systematically different MPC. For example, during bank holidays the volatility of consumption may be higher for high-income households who may have lower MPCs. In order to investigate this potential concern, we split households into two groups depending on whether they are above or below the median variance of weekly consumption expenditure measured in the pretreatment period. We perform the deconvolution exercise separately by treatment group on each of those samples. The estimates, shown in online Appendix Figure D11, are very similar across high- and low-variance groups for each card, indicating that MPC distributions are unlikely to be very different for groups of households with different higher moments of the error term.

IV. Implications

We now discuss the implication of our five facts about MPCs, for both macroeconomic models and stimulus policies.

A. *Implications for Models*

Our experiment is not designed to test any particular model of consumption but instead to robustly estimate moments of consumption responses to transfers that are scalable and therefore relevant for policy. Nonetheless, it is worth discussing which models of consumption can be reconciled with our findings.

Benchmark Rational Models.—We first compare our findings with the predictions of canonical "rational" models. In the HANK model of Kaplan, Moll, and Violante (2018), high average MPCs arise because of precautionary savings in the presence of borrowing constraints. In their baseline calibration, matching moments of the liquid and illiquid wealth distributions and income processes, the simulated consumption response to a one-off lump-sum transfer is long-lived (see Figure 2 in Kaplan, Moll, and Violante (2018)): the estimated MPC is about 17 percent over a quarter (for a $300 transfer), about 25 percent over two quarters, and 32 percent over one year; furthermore, the high MPCs are driven by households with low levels of liquid wealth. The MPC is long-lived in the benchmark HANK model because agents (rationally) increase spending whenever they hit their borrowing constraints, which happens gradually over time as some agents experience negative idiosyncratic income shocks. Over the first two quarters, the increase in the aggregate cumulative MPC is driven by constrained households who deplete the rebate in full at this horizon. Afterwards, the aggregate MPC increases more slowly due to the population of unconstrained agents who consume the annuity value of the transfer.[27]

Our findings stand in contrast with the predictions of the canonical implementation of the benchmark HANK model in three ways. First, even in our treatment group 1—the group that receives a transfer that is most similar to a cash transfer—the entire spending response we find is concentrated in the first two weeks after the transfer (panel A of Figure 3).[28] In contrast, as previously mentioned, the

[27] See Achdou et al. (2022) for a characterization of how cumulative MPCs vary with the time horizon in the Aiyagari–Bewley–Huggett model.

[28] Consistent with our experimental finding, Borusyak, Jaravel, and Spiess (2023) document that the consumption response to tax rebates is concentrated in the first two to three weeks after the tax rebate. Likewise, Baugh et al. (2021) find that households spend a significant part of the tax refunds they receive on consumption in the month after receiving the refund. In contrast, analyzing lottery winnings in Norway worth $9,200 on

MPC response is much more long-lived in HANK and in canonical buffer-stock saving models (Kaplan, Moll, and Violante 2018; Auclert, Rognlie, and Straub 2023). While spending on durables could in principle explain a short-run spending burst in a standard model (Laibson, Maxted, and Moll 2022), we find that the response is also concentrated in the short run for nondurables. We compare our MPC estimates to the standard calibration of the HANK model more formally in online Appendix Figure D12, documenting that the rate of decay of MPCs we estimate is one order of magnitude higher than in HANK.

Second, in HANK the simulated MPC is strongly correlated with the level of liquid assets that agents hold. While we do find some heterogeneity of MPCs for groups with different levels of liquid asset holdings, we find that average MPCs are also high for households that have moderate or high levels of liquid asset holdings (Figure 6). In online Appendix Figure D8, we show that the MPC remains high even for households that hold liquid wealth above twice their monthly income. These findings echo results from the literature that finds high MPCs even for agents with high liquid wealth, including Kueng (2018) in response to anticipated payouts from the Alaska Permanent Fund; Olafsson and Pagel (2018) in response to regular and irregular income transfers in Iceland; Fagereng, Holm, and Natvik (2021) among lottery winners in Norway; and Baugh et al. (2021) in response to expected tax refunds in the United States.[29]

Third, our estimates of the unconditional distributions of MPCs reveal that MPCs are high for a large majority of the population (Figure 8). For group 1 participants, our estimates indicate that half of the population has a one-month MPC above 28 percent and three quarters have an MPC above 13 percent. The fraction of households with a high MPC is thus higher than in benchmark HANK models.[30]

Furthermore, our finding that MPCs are higher for households with lower average past consumption (our proxy for permanent income) stands in contrast with standard macroeconomic models featuring homothetic preferences where the MPC is independent of permanent income. Straub (2019) extends the canonical precautionary savings model to include nonhomothetic preferences, allowing for MPCs that vary with permanent income, consistent with our findings.

Finally, it is instructive to compare our estimates to those used in standard calibrations of HANK models. In Table 3, we summarize the estimates of the consumption response to the 2008 tax rebates in the United States, and we contrast them with our estimates. Following Laibson, Maxted, and Moll (2022), we now draw a distinction between the observed marginal propensity to spend (denoted

average, Fagereng, Holm, and Natvik (2021) estimate a more long-lived MPC response: there is a large consumption response in the first year followed by gradually declining MPCs over several years. This finding could stem from the fact that the lottery winnings are on average larger than tax rebates or tax refunds. The size of the shock matters for the dynamics of the consumption response: for example, Boutros (2023) studies a structural behavioral model in which the planning horizon of the households depends endogenously on the amount of the transitory income shock, such that a larger shock is endogenously smoothed over a longer time horizon.

[29] Stephens and Unayama (2011), Parker (2017), Ganong and Noel (2019), and McDowall (2020) also find that highly liquid households still have elevated MPCs. In contrast, studying the consumption response to typical month-to-month fluctuations in labor income, Ganong et al. (2020) find an MPC close to zero for households with high liquid wealth.

[30] For example, in Kaplan and Violante's (2022) calibrated two-asset model, the 40 percent of households with the highest MPCs (the hand-to-mouth) have an average MPC of 28 percent (see their Figure 4). In our Figure 8, 50 percent of households have an MPC above 28 percent.

Table 3—First-Quarter MPX and MPC Estimates for Calibration of Macroeconomic Models

	Parker et al. (2013) (1)	Broda and Parker (2014) (2)	Borusyak, Jaravel, and Spiess (2023) (3)	Orchard, Ramey, and Wieland (2023b) (4)	This paper, treatment group 1 (5)
Total MPX (%)	52.3 to 91.1	50.8 to 74.8	24.8 to 36.6	28	23
Nondurable MPX (%)	12.8 to 30.8	14.1 to 20.8	6.9 to 10.2	0	6.6
Notional MPC (%)	16.3 to 28.5	15.9 to 23.4	7.8 to 11.4	8.8	7.2

Notes: This table reports the first-quarter MPX and MPC in studies of the 2008 tax rebates in the United States (columns 1–4) and for treatment card 1 participants in our experiment (column 5). The first row reports the MPX on all goods and services, while the second row focuses on nondurables alone. The third row follows the methodology of Laibson, Maxted, and Moll (2022) and reports the model-consistent ("notional") MPC that can be used as a target for macroeconomic models, equal to the total MPX divided by 3.2. The range of estimates in column 1 corresponds to different household samples (see Tables 2 and 3 of Parker et al. (2013)). The range of estimates in columns 2 and 3 corresponds to the lowest and highest values among the three rescaling methods used by Broda and Parker (2014) and Borusyak, Jaravel, and Spiess (2023) to extrapolate the spending response they observe for consumer-packaged goods to broader samples. The estimates in the first two rows of column 4 are taken from Tables 3 and 5 of Orchard, Ramey, and Wieland (2023b). We compare our estimates to a larger set of papers in online Appendix Figure D13.

MPX) and the model-consistent, or "notional," MPC that should be used as a target for macroeconomic models.[31] Columns 1 and 2 summarize the results of Parker et al. (2013) and Broda and Parker (2014), which are the typical targeted moments in fiscal policy models. For example, Kaplan and Violante (2014); Kaplan, Moll, and Violante (2018); and Auclert, Bardóczy, and Rognlie (2023) calibrate their heterogeneous agent models to match an MPC of 25 percent on the nondurables component of consumption expenditures. Columns 3 and 4 report the estimates of Borusyak, Jaravel, and Spiess (2023) and Orchard, Ramey, and Wieland (2023b). Applying event study estimators that are robust to treatment effect heterogeneity in the same samples as Broda and Parker (2014) and Parker et al. (2013), they obtain smaller MPC estimates. Our experimental estimates for treatment group 1, reported in column 5, are close to the bottom of the range of estimates from Borusyak, Jaravel, and Spiess (2023) and Orchard, Ramey, and Wieland (2023b). The last row of Table 3 reports the notional MPC that should be used as a target for macroeconomic models, following the methodology of Laibson, Maxted, and Moll (2022). In sum, when studying standard fiscal transfers, macro models should target the notional MPCs reported in columns 3 to 5, which are about half as large as in the commonly used estimates from the seminal studies summarized in columns 1 and 2.

Assessing whether suitable calibrations or modifications of the HANK model can match the facts summarized above is an important direction for future research.[32] A potential avenue is to augment standard consumption models with

[31] The notional MPC accounts for the fact that spending on durables corresponds to a consumption flow over several periods, which can be consistent with consumption smoothing even though expenditures are front-loaded. Before Laibson, Maxted, and Moll (2022) showed that the notional MPC is the relevant target, state-of-the-art macroeconomic models targeted nondurable MPX estimates.

[32] See Wolf (2023) for a characterization of the shape that iMPCs in HANK models can take, and the extent to which they can be well-approximated by simple models with occasionally binding borrowing constraints (as in, e.g., Farhi and Werning 2019).

certain behavioral frictions. For example, in recent work Boutros (2023) and Lian (2021) develop structural behavioral models in which high-liquidity households have large MPCs because of behavioral biases. Consistent with this line of work, some results of our experiment are difficult to reconcile with agents being rational and treating money as fungible, which we discuss next.

Behavioral Models.—Our motivation to turn to behavioral models is that the difference in MPCs between households assigned to group 1 or groups 2–3 rejects standard rational models where agents treat money as fungible. Indeed, when we consider only transactions below €300 (which can be made with the treatment card), we find that 88 percent of households in group 2 spent at least €300 on the main bank account in the three weeks before the expiry date of card 2. This indicates that it should be costless for a vast majority of households to substitute current account spending for prepaid card spending. In other words, under the rational benchmark, we expect that the three-week expiry date for most households in treatment group 2 should not be a binding constraint; that is, their MPC should be similar to households in treatment group 1 in contrast with our findings.

Online Appendix Figure D14 shows for each day the fraction of households in treatment groups 2 and 3 that would have had a high enough balance on the treatment card to cover the day's expenditures (as measured by their spending on nontreatment cards) but for some reason did not use the card. A nonnegligible share of households in groups 2 and 3 has a high enough remaining balance on the treatment card to cover the day's expenditures but chooses to use their regular debit or credit card instead to make purchases. In the first few days of the experiment this ratio may be high because some households had not opened their mail and therefore had not started to use the card. But even after more than a week into the experiment, the ratio remains above 15 percent. Online Appendix Figure D15 shows that the patterns are the same in a restricted sample of households with a single adult and no children, ruling out the possibility that this phenomenon is driven by multiperson households of whom only one has access to the treatment card.[33]

These facts are hard to reconcile with rational households that treat money as fungible. Indeed, a rational agent that treats money as fungible should first "use up" the treatment card to avoid potentially losing money (through the negative interest rate or expiry) before using their normal debit or credit card. Thus, our results echo a literature in economics (Hastings and Shapiro 2013, 2018; Gelman

[33] Households may need to incur some (e.g., time) costs to use new means of payment for some transactions, for example automatic payments like utility bills. A model where households face small costs to adjust their means of payment could in principle generate the pattern in online Appendix Figure D14 as households might wait to pay the switching cost. However, in France most automatic payments are made through wire transfers or direct debit, which are not included in our baseline consumption measure: these transactions do not drive the patterns in online Appendix Figure D14. Furthermore, a model with small costs cannot explain why the MPC in groups 2 and 3 is larger. Indeed, cards with an expiry date or negative rates may spur households to pay the small adjustment costs more quickly and cover their automatic payments with the prepaid cards, but this would result in a *lower* MPC for these groups since these expenses were already planned. Alternatively, a model where people might forget about using the card or might be distracted by certain life events could generate the patterns in online Appendix Figure D14. Forgetfulness may be plausible given that the amount (€300) is a small fraction of households' lifetime income.

and Roussanov 2023; Chan and Kan 2024) and in sociology (e.g., Zelizer 1989) that emphasizes the nonfungibility of money.[34]

With these patterns in mind, our findings deliver three lessons for behavioral models. First, models of consumption that rely on present bias in preferences (e.g. Laibson 1997; Maxted 2020; Laibson, Maxted, and Moll 2021; Gelman 2022) are able to explain why the consumption response to the transfer is concentrated early on but cannot explain the difference in the magnitude of responses between the treatment groups. Indeed, under such preferences consumers in all three groups should be present biased, but the negative interest rate and the expiry date would remain nonbinding constraints given that it should be costless for agents to substitute current account spending for prepaid card spending. Thus, present bias does not appear to be the key friction explaining our findings.[35]

Second, another class of behavioral models that have been used for macro policy analysis are models that feature two sets of agents, "savers" and "spenders," who have low and, respectively, high MPCs (Campbell and Mankiw (1989); for the Two-Agent New Keynesian models, see the review in Galí (2018)). While implementations of such models can be made to feature consumption responses that are concentrated very early on, they would also imply strongly bimodal distributions of MPCs, which we do not find (Figure 8). Furthermore, like other models of present bias, this type of model cannot account for the difference in spending patterns by card type.

Third, our results are consistent with models of salience where small but highly prominent features of the choice set distract the attention of decision makers and distort their choices (Bordalo, Gennaioli, and Shleifer 2012, 2013). In particular, salience can lead households to engage in "mental accounting" (e.g., Shefrin and Thaler 1988; Thaler 1990; McDowall 2020; Baugh et al. 2021; Boutros 2023). In online Appendix F, we formalize a stylized model of mental accounting that could explain the key empirical patterns we observe for the three treatment groups. In this model, the agent faces a tradeoff when spending the prepaid card on unplanned "windfall consumption" (e.g., going to a fancy restaurant, going out more frequently than usual, purchasing a treat, etc.). On the one hand, the agent incurs a cognitive dissonance cost if they spend the prepaid card on (planned) regular consumption rather than on an unplanned treat because of a mental account mechanism: the prepaid card is perceived by the agent to be "special money" meant to be spent on extra consumption, like in the sociology literature (Zelizer 1989, 2017). On the other hand, purchasing treats requires incurring search costs while using the prepaid card to cover running expenses does not. Resolving this tradeoff in the model, we show that the spending response is concentrated in the short run for all cards and is largest for group 2 followed by group 3 and finally group 1. Intuitively, prepaid cards with an expiry date or a negative interest rate spur the agent to incur the search costs

[34] In contrast to work that has found that the labeling of cash transfers has an impact on spending patterns (Beatty et al. 2014; Benhassine et al. 2015), in Section IID we detected no significant effect of framing on the magnitude and composition of expenditures.

[35] Loss aversion is another bias that has been widely studied (e.g., Tversky and Kahneman 1992). While households might exhibit loss aversion, it does not appear to be the key friction in our setting because loss aversion does not imply that participants would not treat money as fungible: they could easily avoid any loss—from the expiry date or negative interest rates—by using the prepaid card to cover running expenses.

more quickly as long as these costs are not too large. This need to take action in the short run is salient and can lead to groups 2 and 3 having higher MPCs than group 1. When the search costs are higher (e.g., when a decision must be made within a week to avoid a negative interest rate, as in group 3), the agent is more likely to cover regular consumption than to purchase unplanned windfall consumption goods and services, implying a lower MPC than with a longer expiry date (as in group 2).[36]

Another potential mechanism through which the difference in MPCs between groups 2 and 3 could be explained is dual reasoning (Ilut and Valchev 2023). Agents are confronted with different decision problems and can make decisions either rapidly and intuitively by projecting on past deliberations ("system 1 thinking"), or by carefully considering their choices, which leads to better outcomes but which is also cognitively costly ("system 2 thinking"). Situations where people receive a means of payment that they have a certain time frame to spend, such as in group 2, are familiar to many from gift vouchers and gift cards and may lead recipients to behave similarly to how they behaved in such situations (through system 1 thinking). In contrast, the situation where the participant receives a means of payment that rapidly loses value is unfamiliar to most, resulting in careful deliberations (activating system 2 thinking) to avoid the loss of value and, more often than not, the purchase of goods that they would have purchased anyway, implying a lower MPC in group 3 than in group 2. Note that we observe that group 3 participants, triggered by the salience of the one-week ultimatum before they lose money, on average spend more using the treatment card than group 2 on each day of the first week.[37]

B. *Implications for Macroeconomic Stabilization Policies*

Our results have two immediate implications for policy. First, the difference in MPCs across treatment groups 1 and 2 (23 percent versus 61 percent) shows that the design of stimulus transfers can help increase MPCs. Note that because some money ends up being returned in treatment designs 2 and 3, the average consumption stimulus per euro actually spent is larger than the MPC estimates reported above. In online Appendix Figure D16 we plot the MPC for groups 2 and 3 corrected by the fraction of the money that is being returned in the form of interest payments (group 3) or remaining balance upon card expiry (group 2), which is about 16 percent for both

[36] Analyzing the consumption response to tax refunds and tax payments, Baugh et al. (2021) highlight that the estimates in their study are most consistent with a mental accounting life-cycle model following Shefrin and Thaler (1988). They find that households increase spending when they receive an anticipated tax refund and that these same households completely smooth consumption when making anticipated tax payments, implying that they have the liquidity to smooth consumption through refunds. Thus, households spend out of tax refunds by choice rather than due to liquidity constraints, consistent with mental accounting. Anticipated tax refunds are part of the "future income" mental account and are not smoothed while tax payments are part of the "current income" mental account, which leads to consumption smoothing.

[37] An alternative potential mechanism, suggested by an anonymous referee, is the presence of small "real" frictions combined with potentially small benefits from fully rational consumption smoothing behavior. In the presence of "real" frictions (e.g., the hassle cost of using a new card), setting an expiry deadline or a negative interest rate can endogenously lead households to spend down their prepaid cards more quickly. If these households are fully rational, they should treat money as fungible and smooth consumption perfectly by reducing spending on their main account as they spend down the prepaid card (i.e., they should not have a higher overall MPC). However, there may be settings in which households are willing not to smooth consumption (e.g., because the amount on the card is small relative to lifetime consumption; i.e., the utility loss from not smoothing is small). In such cases, the real frictions could lead to higher spending and explain why cards 2 and 3 induce higher MPCs than card 1.

groups. The resulting effective stimulus at the four-week horizon is about 75 cents per euro of net transfer for group 2 and about 40 cents per euro for group 3.

The external validity of our experimental estimates and their broad applicability to high-income countries appears plausible given that (i) we used a representative sample of the French population and (ii) our estimates for group 1 are very similar to those obtained when studying the 2008 tax rebate response in the United States with robust estimators (Borusyak, Jaravel, and Spiess 2023; Orchard, Ramey, and Wieland 2023b). Our intervention was deliberately designed to be scalable to the macro level, and we note that there are several examples of large-scale stimulus policies using prepaid cards or time-limited consumption vouchers, including Japan in 1999; Taiwan in 2009; California, Milan, and Seoul in 2020; and Hong Kong in 2021. Finding ways of raising MPCs is of particular importance given recent estimates of the MPCs out of standard tax rebates in the United States, which are relatively low as found by Parker et al. (2022b) and Parker et al. (2022a) for the 2020 stimulus payment and by Borusyak, Jaravel, and Spiess (2023) and Orchard, Ramey, and Wieland (2023b) for the 2008 stimulus payments. Using prepaid cards with negative rates or expiration dates to raise MPCs is therefore a promising avenue for stimulus policies going forward, which could potentially be implemented by central banks using central bank digital currencies.[38] It is also worth noting that short-term interest rates were close to zero at the time when our experiment was implemented (see online Appendix Figure D17), indicating the possible potency of particular types of stimulus policies even in a liquidity trap.[39]

Second, our estimates of MPC heterogeneity have implications for the targeting of transfers by observable household characteristics. We documented in Section IIIA that many household characteristics can be used to predict heterogeneity in MPCs. Thus, transfers could be targeted to the households with the highest MPC. While liquidity is difficult to observe, other predictors are readily accessible to policymakers. To assess the extent to which the average MPC of transfer recipients could be increased by targeting, we conduct a simple exercise: we use the specification from Section IIIA with two sets of characteristics that policymakers might be able to observe as regressors and estimate the distribution of MPCs. We estimate the parameters using LASSO to avoid overfitting in a sample consisting of control group households and households receiving treatment card 1. By plotting the estimated distribution of treatment effects we can thus assess the extent to which household targeting can help increase the MPC for a standard transfer without negative rates or an expiry date.

Figure 9 plots the distribution of predicted MPCs. In the first exercise, reported in panel A, we run LASSO with age quartiles, income quartiles, and unemployment

[38] If transfers with expiry dates were used repeatedly, one could worry that households may start viewing these transfers as more fungible with their main bank account and thus have a lower MPC. However, existing evidence suggest that mental accounting continues to operate even for repeated transitory shocks as found by Baugh et al. (2021) for tax rebates, by Hastings and Shapiro (2013) for gasoline purchases, and by Hastings and Shapiro (2018) for food stamps.

[39] Other types of stimulus policies at the zero lower bound potentially include VAT rate cuts; see, for example, Correia et al. (2013). Note that our experiment cannot be used to learn about the impact of broader changes in interest rates on consumption behavior or about the intertemporal elasticity of substitution for two reasons: (i) the "negative interest rates" are only applied to the treatment card (i.e., to a very small fraction of households' assets); (ii) if households treated money as fungible, the negative rates on the treatment cards would not affect their budget sets.

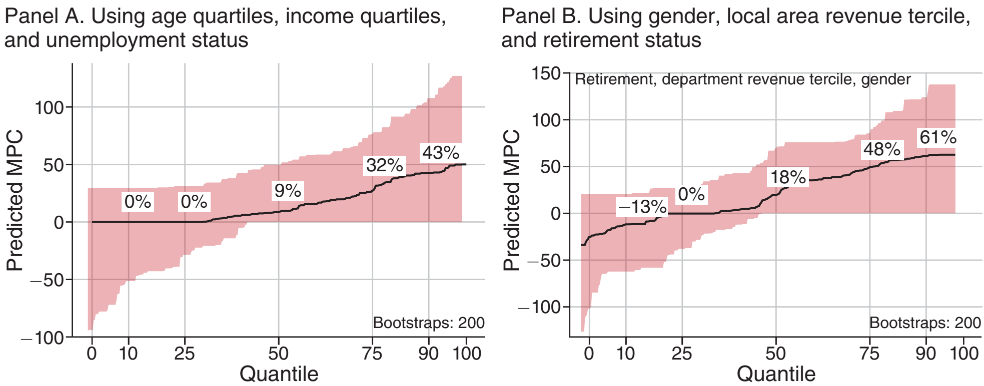

Figure 9. Predicted MPC Heterogeneity

Notes: This figure shows the distributions of the predicted MPC heterogeneity using different sets of characteristics as predictors of the treatment effect in a LASSO specification. Panel A uses age quartiles, income quartiles, and unemployment status as features in the LASSO specification, while panel B uses gender, local area revenue tercile, and retirement status as features. The sample is restricted to treatment card 1. The 95 percent confidence intervals are obtained by bootstrap and shown as shaded regions.

status. This panel shows that by using these observables it is possible to identify households with substantially above-average MPCs. For example, 10 percent of households are predicted to have an MPC above 43 percent. In the second exercise, shown in panel B, we predict the MPC distribution using gender, local area revenue tercile, and retirement status. The top 10 percent of households have an MPC of 61 percent. Targeting can therefore be a relatively powerful tool to increase the average MPC of recipients although it is not as potent as changing the design of the treatment card. For example, treatment card 2 with an expiry date yields an average MPC across *all* participants of 61 percent. Thus, our estimates highlight that implementation design choices are a more powerful tool compared to targeting to increase the recipients' average MPC.[40] In addition, targeting may raise political economy or fairness considerations that are avoided by providing a treatment card with an expiry date to all.

V. Conclusion

In this paper we presented five facts about MPCs obtained from a randomized experiment where we provide money transfers to a representative set of French households. These results inform the academic debate on models of consumption but are also directly relevant for the design of effective stimulus policies.

First, we found that the one-month MPC is 23 percent with a standard treatment card without negative interest rates. Second, we established our main result: the design of the transfer matters. The one-month MPC is higher when treatment cards feature a negative interest rate, at 61 percent when the remaining balance is reduced to zero after three weeks, and 35 percent when the remaining balance is reduced by approximately 10 percent every week. Third, the increase in consumption is much

[40] Of course the policymaker may have goals other than increasing the short-term MPC, for example to change the composition of spending or to smooth consumption over a longer time frame.

larger early on in the first two to three weeks after receiving the transfer. Fourth, heterogeneity in the MPCs that is explained by observed households characteristics is substantial, including by variables distinct from liquid wealth such as current income, proxies for permanent income, and gender. Fifth, the unconditional heterogeneity in MPCs is very large, and a large fraction of households have high MPCs.

These five facts are hard to reconcile with standard two-asset models of consumption. They point to the importance of behavioral features (e.g., salience) for macroeconomic models of the consumption response to transfers such that agents do not treat stimulus transfers as fungible with standard income sources. The "five facts about prices" of Nakamura and Steinsson (2008) called for a reevaluation of menu cost models; much in the same spirit, our five facts about MPCs provide moments that can help discipline consumption and macro models.[41]

From a policy perspective, our findings indicate that implementation design, and to a lesser extent household targeting, are key tools to manipulate MPCs and increase the effectiveness of stimulus. Prepaid cards with negative interest rates or an expiry date deliver much larger MPC than standard fiscal stimulus and constitute a powerful tool to stimulate demand even when interest rates are low.

An important avenue for future research is to scale up sample sizes to obtain more precise estimates of MPCs at longer horizons and of the heterogeneity by observable household characteristics and treatment designs. Indeed, while our empirical results are supportive of little intertemporal substitution, the precision of our estimates is not high enough to rule it out. To guide the design of future experiments, online Appendix G presents power calculations informed by our data.

Another important direction for future work is to quantify the welfare effects of administering stimulus programs with cards featuring time limits or negative interest rates. Specifically, a fruitful task would be to compare the household-level welfare losses when using such cards (as agents receiving these cards do not smooth consumption as much)[42] to the welfare gains from the aggregate demand externalities that arise in general equilibrium.[43] We leave these and other extensions for future work.

REFERENCES

Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. 2022. "Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach." *Review of Economic Studies* 89 (1): 45–86.

Aguiar, Mark A., Mark Bils, and Corina Boar. 2023. "Who are the Hand-to-Mouth?" Unpublished.

[41] Some of our findings stand in contrast with those obtained by eliciting households' stated marginal propensities to consume in surveys. While stated MPCs in surveys are clustered around specific values like 0 percent, 50 percent, and 100 percent (e.g., Jappelli and Pistaferri (2014)), we find that actual MPCs are smoothly distributed (Fact 5). Similarly, stated MPCs out of the 2008 tax rebate in the United States are about twice as large as the MPCs estimated in the data: see Parker and Souleles (2019) for the stated MPCs and Borusyak, Jaravel, and Spiess (2023) and Orchard, Ramey, and Wieland (2023b) for the empirical estimates of the consumption response. It is therefore important to study actual behavior rather than reported MPCs.

[42] To get an extreme upper bound for these welfare losses, one can simply assume that the €300 transfer does not raise participants' welfare at all; that is, the welfare loss would be equal to 1.16 percent of annual consumption in our data. In general, the welfare loss depends on the curvature of the per-period utility function and on the amount of the transfer.

[43] As usual, the general equilibrium response could be limited if supply constraints were binding (see, e.g. Orchard, Ramey and Wieland (2023b)). Supply constraints may be more likely to bind if spending is very concentrated in the short run.

Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber. 2005. "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools." *Journal of Political Economy* 113 (1): 151–84.

Angeletos, George-Marios, Chen Lian, and Christian K. Wolf. 2023. "Can Deficits Finance Themselves?" Unpublished.

Auclert, Adrien, Bence Bardóczy, and Matthew Rognlie. 2023. "MPCs, MPEs, and Multipliers: A Trilemma for New Keynesian Models." *Review of Economics and Statistics* 105 (3): 700–12.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub. 2023. "The Intertemporal Keynesian Cross." NBER Working Paper 25020.

Aydin, Deniz. 2022. "Consumption Response to Credit Expansions: Evidence from Experimental Assignment of 45,307 Credit Lines." *American Economic Review* 112 (1): 1–40.

Baek, Seungjun, Seongeun Kim, Tae-hwan Rhee, and Wonmun Shin. 2023. "How Effective are Universal Payments for Raising Consumption? Evidence from a Natural Experiment." *Empirical Economics* 65 (5): 2181–2211.

Baker, Scott R., and Lorenz Kueng. 2022. "Household Financial Transaction Data." *Annual Review of Economics* 14: 47–67.

Banerjee, Abhijit, and Esther Duflo. 2009. "The Experimental Approach to Development Economics." *Annual Review of Economics* 1 (1): 151–78.

Banerjee, Abhijit, Esther Duflo, and Michael Kremer. 2016. "The Influence of Randomized Controlled Trials on Development Economics Research and on Development Policy." In *State of Economics, The State of the World*, edited by Kaushik Basu, David Rosenblatt, and Claudia Sepúlveda, 439–88. Cambridge, MA: MIT Press.

Baugh, Brian, Itzhak Ben-David, Hoonsuk Park, and Jonathan A. Parker. 2021. "Asymmetric Consumption Smoothing." *American Economic Review* 111 (1): 192–230.

Beatty, Timothy K.M., Laura Blow, Thomas F. Crossley, and Cormac O'Dea. 2014. "Cash by Any Other Name? Evidence on Labeling from the UK Winter Fuel Payment." *Journal of Public Economics* 118: 86–96.

Benhassine, Najy, Florencia Devoto, Esther Duflo, Pascaline Dupas, and Victor Pouliquen. 2015. "Turning a Shove into a Nudge? A "Labeled Cash Transfer" for Education." *American Economic Journal: Economic Policy* 7 (3): 86–125.

Blundell, Richard, Luigi Pistaferri, and Ian Preston. 2008. "Consumption Inequality and Partial Insurance." *American Economic Review* 98 (5): 1887–1921.

Boehm, Johannes, Etienne Fize, and Xavier Jaravel. 2022. *Implementing Helicopter Money: A Randomized Policy Evaluation*. AEA RCT Registry. https://doi.org/10.1257/rct.9296-1.0.

Boehm, Johannes, Etienne Fize, and Xavier Jaravel. 2025. *Data and Code for: "Five Facts about MPCs: Evidence from a Randomized Experiment."* Nashville, TN: American Economic Association; distributed by Inter-university Consortium for Political and Social Research, Ann Arbor, MI. https://doi.org/10.3886/E209343V1.

Bonnet, Odran, Etienne Fize, Tristan Loisel, and Lionel Wilner. 2023. "How Does Fuel Demand Respond to Anticipated Price Changes? quasi-experimental Evidence Based on High-Frequency Data." Unpublished.

Bonnet, Odran, Tom Olivia, and Théo Roudil-Valentin. 2021. "En 2020, la chute de la consommation a alimenté l'épargne, faisant progresser notamment les hauts patrimoines financiers: quelques résultats de l'exploitation de données bancaires." Unpublished.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2012. "Salience Theory of Choice under Risk." *Quarterly Journal of Economics* 127 (3): 1243–85.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2013. "Salience and Consumer Choice." *Journal of Political Economy* 121 (5): 803–43.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess. 2023. "Revisiting Event Study Designs: Robust and Efficient Estimation." Unpublished.

Bounie, David, Youssouf Camara, Etienne Fize, John Galbraith, Camille Landais, Chloe Lavest, Tatiana Pazem, and Baptiste Savatier. 2020. "Con- sumption Dynamics in the Covid Crisis: Real Time Insights from French Transaction Bank Data." *Covid Economics* 59: 1–39.

Boutros, Michael. 2023. "Windfall Income Shocks with Finite Planning Horizons." Unpublished.

Broda, Christian, and Jonathan A. Parker. 2014. "The Economic Stimulus Payments of 2008 and the Aggregate Demand for Consumption." *Journal of Monetary Economics* 68: S20–36.

Bunn, Philip, Jeanne Le Roux, Kate Reinold, and Paolo Surico. 2018. "The Consumption Response to Positive and Negative Income Shocks." *Journal of Monetary Economics* 96: 1–15.

Campbell, John Y., and N. Gregory Mankiw. 1989. "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence." *NBER Macroeconomics Annual* 4 (1): 185–216.

Chan, Marc, and Kamhon Kan. 2024. “Consumption Sensitivity to Stimulus Payments and Income Nonfungibility.” Unpublished.

Commault, Jeanne. 2022a. “Does Consumption Respond to Transitory Shocks? Reconciling Natural Experiments and Semistructural Methods.” *American Economic Journal: Macroeconomics* 14 (2): 96–122.

Commault, Jeanne. 2022b. “How Do Persistent Earnings Affect the Response of Consumption to Transitory Shocks?” Unpublished.

Correia, Isabel, Emmanuel Farhi, Juan Pablo Nicolini, and Pedro Teles. 2013. “Unconventional Fiscal Policy at the Zero Bound.” *American Economic Review* 103 (4): 1172–1211.

Deaton, Angus. 2010. “Instruments, Randomization, and Learning about Development.” *Journal of Economic Literature* 48 (2): 424–55.

Deaton, Angus, and Nancy Cartwright. 2018. “Understanding and Misunderstanding Randomized Controlled Trials.” *Social Science and Medicine* 210: 2–21.

De Chaisemartin, Clement, and Xavier Haultfoeuille. 2020. “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects.” *American Economic Review* 110 (9): 2964–96.

De Quidt, Jonathan, Johannes Haushofer, and Christopher Roth. 2018. “Measuring and Bounding Experimenter Demand.” *American Economic Review* 108 (11): 3266–3302.

Ding, Jing, Lei Jiang, Lucy Msall, and Matthew J. Notowidigdo. 2024. “Consumer-Financed Fiscal Stimulus: Evidence from Digital Coupons in China.” NBER Working Paper 32376.

Ellison, Glenn, and Alexander Wolitzky. 2012. “A Search Cost Model of Obfuscation.” *RAND Journal of Economics* 43 (3): 417–41.

Evans, George W., and Garey Ramey. 1992. “Expectation Calculation and Macroeconomic Dynamics.” *American Economic Review* 82 (1): 207–24.

Fagereng, Andreas, Martin B. Holm, and Gisle J. Natvik. 2021. “MPC Heterogeneity and Household Balance Sheets.” *American Economic Journal: Macroeconomics* 13 (4): 1–54.

Farhi, Emmanuel, and Iván Werning. 2019. “Monetary Policy, Bounded Rationality, and Incomplete Markets.” *American Economic Review* 109 (11): 3887–3928.

Fuster, Andreas, Greg Kaplan, and Basit Zafar. 2021. “What Would You Do With $500? Spending Responses to Gains, Losses, News, and Loans.” *Review of Economic Studies* 88 (4): 1760–95.

Galí, Jordi. 2018. “The State of New Keynesian Economics: A Partial Assessment.” *Journal of Economic Perspectives* 32 (3): 87–112.

Ganong, Peter, and Pascal Noel. 2019. “Consumer Spending During Unemployment: Positive and Normative Implications.” *American Economic Review* 109 (7): 2383–2424.

Ganong, Peter, Damon Jones, Pascal J. Noel, Fiona E. Greig, Diana Farrell, and Chris Wheat. 2020. “Wealth, Race, and Consumption Smoothing of Typical Income Shocks.” NBER Working Paper 27552.

Garber, Gabriel, Atif R. Mian, Jacopo Ponticelli, and Amir Sufi. 2022. “Consumption Smoothing or Consumption Binging? The Effects of Government-led Consumer Credit Expansion in Brazil.” Unpublished.

Gelman, Michael. 2021. “What Drives Heterogeneity in the Marginal Propensity to Consume? Temporary Shocks vs Persistent Characteristics.” *Journal of Monetary Economics* 117: 521–42.

Gelman, Michael. 2022. “The Self-Constrained Hand-to-Mouth.” *Review of Economics and Statistics* 104 (5): 1096–1109.

Gelman, Michael, and Nikolai Roussanov. 2023. “Managing Mental Accounts: Payment Cards and Consumption Expenditures.” Unpublished.

Gelman, Michael, Shachar Kariv, Matthew D. Shapiro, Dan Silverman, and Steven Tadelis. 2014. “Harnessing Naturally Occurring Data to Measure the Response of Spending to Income.” *Science* 345 (6193): 212–15.

Gelman, Michael, Yuriy Gorodnichenko, Shachar Kariv, Dmitri Koustas, Matthew D. Shapiro, Dan Silverman, and Steven Tadelis. 2023. “The Response of Consumer Spending to Changes in Gasoline Prices.” *American Economic Journal: Macroeconomics* 15 (2): 129–60.

Geng, Hao, Ce Shi, and Michael Zheng Song. 2022. “Evaluating Hong Kong Consumption Voucher Scheme.” Unpublished.

Golosov, Mikhail, Michael Graber, Magne Mogstad, and David Novgorodsky. 2024. “How Americans Respond to Idiosyncratic and Exogenous Changes in Household Wealth and Unearned Income.” *Quarterly Journal of Economics* 139 (2): 1321–95.

Hastings, Justine, and Jesse M. Shapiro. 2018. “How are SNAP Benefits Spent? Evidence from a Retail Panel.” *American Economic Review* 108 (12): 3493–3540.

Hastings, Justine S., and Jesse M. Shapiro. 2013. “Fungibility and Consumer Choice: Evidence from Commodity Price Shocks.” *Quarterly Journal of Economics* 128 (4): 1449–98.

Hsieh, Chang-Tai, Satoshi Shimizutani, and Masahiro Hori. 2010. “Did Japan’s Shopping Coupon Program Increase Spending?” *Journal of Public Economics* 94 (7- 8): 523–29.

Ilut, Cosmin, and Rosen Valchev. 2023. “Economic Agents as Imperfect Problem Solvers.” *Quarterly Journal of Economics* 138 (1): 313–62.

Jappelli, Tullio, and Luigi Pistaferri. 2010. “The Consumption Response to Income Changes.” *Annual Review of Economics* 2 (1): 479–506.

Jappelli, Tullio, and Luigi Pistaferri. 2014. “Fiscal Policy and MPC Heterogeneity.” *American Economic Journal: Macroeconomics* 6 (4): 107–36.

Jaroszewicz, Ania, Jon Jachimowicz, Oliver Hauser, and Julian Jamison. 2022. “How Effective is (More) Money? Randomizing Unconditional Cash Transfer Amounts in the US.” Unpublished.

Johnson, David S., Jonathan A. Parker, and Nicholas S. Souleles. 2006. “Household Expenditure and the Income Tax Rebates of 2001.” *American Economic Revie*w 96 (5): 1589–1610.

Kan, Kamhon, Shin-Kun Peng, and Ping Wang. 2017. “Understanding Consumption Behavior: Evidence from Consumers’ Reaction to Shopping Vouchers.” *American Economic Journal: Economic Policy* 9 (1): 137–53.

Kaplan, Greg, and Giovanni L. Violante. 2014. “A Model of the Consumption Response to Fiscal Stimulus Payments.” *Econometrica* 82 (4): 1199–1239.

Kaplan, Greg, and Giovanni L. Violante. 2022. “The Marginal Propensity to Consume in Heterogeneous Agent Models.” *Annual Review of Economics* 14: 747–75.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante. 2018. “Monetary Policy According to HANK.” *American Economic Review* 108 (3): 697–743.

Kim, Moon Jung, and Soohyung Lee. 2021. “Can Stimulus Checks Boost an Economy under COVID-19? Evidence from South Korea.” *International Economic Journal* 35 (1): 1–12.

Kueng, Lorenz. 2018. “Excess Sensitivity of High-Income Consumers.” *Quarterly Journal of Economics* 133 (4): 1693–1751.

Ku, Inhoe, Sunyu Ham, and Heyjin Moon. 2023. “Means-Tested COVID-19 Stimulus Payment and Consumer Spending: Evidence from Card Transaction Data in South Korea.” *Economic Analysis and Policy* 78: 1359–71.

Laibson, David. 1997. “Golden Eggs and Hyperbolic Discounting.” *Quarterly Journal of Economics* 112 (2): 443–78.

Laibson, David, Peter Maxted, and Benjamin Moll. 2021. “Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy.” NBER Working Paper 29094.

Laibson, David, Peter Maxted, and Benjamin Moll. 2022. “A Simple Mapping from MPCs to MPXs.” NBER Working Paper 29664.

Lewis, Daniel J., Davide Melcangi, and Laura Pilossoph. 2019. *Latent Heterogeneity in the Marginal Propensity to Consume*. New York, NY: Federal Reserve Bank of New York.

Lian, Chen. 2021. “Mistakes in Future Consumption, High MPCs Now.” NBER Working Paper 29517.

Liu, Qiao, Qiaowei Shen, Zhenghua Li, and Shu Chen. 2021. “Stimulating Consumption at Low Budget: Evidence from a Large-Scale Policy Experiment Amid the COVID-19 Pandemic.” *Management Science* 67 (12): 7291–7307.

Maxted, Peter. 2020. “Present Bias in Consumption-Saving Models: A Tractable Continuous-Time Approach.” Unpublished.

McDowall, Robert A. 2020. “Consumption Behavior Across the Distribution of Liquid Assets.” Unpublished.

Misra, Kanishka, and Paolo Surico. 2014. “Consumption, Income Changes, and Heterogeneity: Evidence from Two Fiscal Stimulus Programs.” *American Economic Journal: Macroeconomics* 6 (4): 84–106.

Nakamura, Emi, and Jón Steinsson. 2008. “Five Facts about Prices: A Reevaluation of Menu Cost Models.” *Quarterly Journal of Economics* 123 (4): 1415–64.

Nakamura, Emi, and Jón Steinsson. 2018. “Identification in Macroeconomics.” *Journal of Economic Perspectives* 32 (3): 59–86.

Olafsson, Arna, and Michaela Pagel. 2018. “The Liquid Hand-to-Mouth: Evidence from Personal Finance Management Software.” *Review of Financial Studies* 31 (11): 4398–4446.

Orchard, Jacob, Valerie A. Ramey, and Johannes F. Wieland. 2023a. “Using Macro Counterfactuals to Assess Plausibility: An Illustration using the 2001 Rebate MPCs.” Unpublished.

Orchard, Jacob, Valerie A. Ramey, and Johannes Wieland. 2023b. “Micro MPCs and Macro Counterfactuals: The Case of the 2008 Rebates.” Unpublished.

Oster, Emily. 2019. “Unobservable Selection and Coefficient Stability: Theory and Evidence.” *Journal of Business and Economic Statistics* 37 (2): 187–204.

Parker, Jonathan A. 2017. “Why Don’t Households Smooth Consumption? Evidence from a $25 million Experiment.” *American Economic Journal: Macroeconomics* 9 (4): 153–83.

Parker, Jonathan A., and Nicholas S. Souleles. 2019. "Reported Effects Versus Revealed-Preference Estimates: Evidence from the Propensity to Spend Tax Rebates." *American Economic Review: Insights* 1 (3): 273–90.
Parker, Jonathan A., Jake Schild, Laura Erhard, and David Johnson. 2022a. "Economic Impact Payments and Household Spending During the Pandemic." *Brookings Papers on Economic Activity* 52 (2): 81–130.
Parker, Jonathan A., Jake Schild, Laura Erhard, and David Johnson. 2022b. "Household Spending Responses to the Economic Impact Payments of 2020: Evidence from the Consumer Expenditure Survey." NBER Working Paper 30596.
Parker, Jonathan A., Nicholas S. Souleles, David S. Johnson, and Robert McClelland. 2013. "Consumer Spending and the Economic Stimulus Payments of 2008." *American Economic Review* 103 (6): 2530–53.
Schennach, Susanne M. 2016. "Recent Advances in the Measurement Error Literature." *Annual Review of Economics* 8: 341–77.
Shapiro, Matthew D., and Joel Slemrod. 2003. "Consumer Response to Tax Rebates." *American Economic Review* 93 (1): 381–96.
Shapiro, Matthew D., and Joel Slemrod. 2009. "Did the 2008 Tax Rebates Stimulate Spending?" *American Economic Review* 99 (2): 374–79.
Shefrin, Hersh M., and Richard H. Thaler. 1988. "The Behavioral Life-Cycle Hypothesis." *Economic Inquiry* 26 (4): 609–43.
Stephens, Melvin, and Takashi Unayama. 2011. "The Consumption Response to Seasonal Income: Evidence from Japanese Public Pension Benefits." *American Economic Journal: Applied Economics* 3 (4): 86–118.
Straub, Ludwig. 2019. "Consumption, Savings, and the Distribution of Permanent Income." Unpublished.
Thaler, Richard H. 1990. "Anomalies: Saving, Fungibility, and Mental Accounts." *Journal of Economic Perspectives* 4 (1): 193–205.
Tversky, Amos, and Daniel Kahneman. 1992. "Advances in Prospect Theory: Cumulative Representation of Uncertainty." *Journal of Risk and Uncertainty* 5 (4): 297–323.
Wolf, Christian K. 2023. "Interest Rate Cuts vs. Stimulus Payments: An Equivalence Result." Unpublished.
Woo, Seokjin, Sangmin Aum, Dohyung Kim, Heyjin Moon, and Soohyung Lee. 2021. "Consumption Response to Seoul's COVID-19 Shopping Coupons: Evidence from Consumer Data." Unpublished.
Wu, Di, Harikesh Nair, and Tong Geng. 2020. "Consumption Vouchers During Covid-19: Evidence from E-commerce." Unpublished.
Xing, Jianwei, Eric Zou, Zhentao Yin, Yong Wang, and Zhenhua Li. 2023. ""'Quick Response" Economic Stimulus: The Effect of Small-Value Digital Coupons on Spending." *American Economic Journal: Macroeconomics* 15 (4): 249–304.
Yang, Ran, Daniel W. Apley, Jeremy Staum, and David Ruppert. 2020. "Density Deconvolution with Additive Measurement Errors Using Quadratic Programming." *Journal of Computational and Graphical Statistics* 29 (3): 580–91.
Youth Training Statistics and Research Branch. 2022. *Analysis of the Northern Ireland High Street Scheme*. Belfast, Ireland: Northern Ireland Department for the Economy.
Zelizer, Viviana A. 1989. "The Social Meaning of Money: Special Monies." *American Journal of Sociology* 95 (2): 342–77.
Zelizer, Viviana A. 2017. *The Social Meaning of Money: Pin Money, Paychecks, Poor Relief, and Other Currencies*. Princeton, NJ: Princeton University Press.