

Five Facts about MPCs: Evidence from a Randomized Experiment*

Johannes Boehm

Sciences Po & CEPR

`johannes.boehm@sciencespo.fr`

Etienne Fize

PSE & IPP

`etienne.fize@ipp.eu`

Xavier Jaravel

LSE & CEPR

`x.jaravel@lse.ac.uk`

November 8, 2023

Abstract

We conduct a randomized controlled trial to study the consumption response of French households to unanticipated one-time money transfers of 300 Euros. Using prepaid debit cards, we consider three implementation designs: (i) a transfer without restrictions; (ii) a transfer where any unspent value expires after three weeks; (iii) a transfer subject to a 10% negative interest rate every week. We observe the participants' main bank accounts, such that we can compute the impact of the transfer on their overall spending. We establish five facts about MPCs in this setting. First, we find that participants in the baseline treatment group have an average marginal propensity to consume (MPC) of 23% over one month. Second, we find that implementation design matters: the one-month MPC is substantially higher for treatment groups where any remaining balance becomes unusable after three weeks (61%) or where remaining balances are subject to the 10% negative interest rate every week (35%). Third, we document that the cumulative consumption responses are concentrated in the first weeks following the transfer and are flat thereafter. Fourth, we find that there is significant MPC heterogeneity by observed household characteristics, including by liquid wealth, current income, proxies for permanent income, and gender; the MPC remains high even for agents with high liquid wealth. Fifth, we estimate the unconditional distribution of MPCs across households and find that a large fraction of households have high MPCs. These facts are difficult to reconcile with the consumption response in standard Heterogeneous Agent New Keynesian models, which is long-lived and driven by households with little liquid wealth. Furthermore, we observe that households in the treatment groups with a short expiry date or a negative interest rate frequently use other means of payment while still having a sufficient balance on the prepaid card to cover their expenses, indicating that participants see money as non-fungible. Our finding that households consume more when presented with an urgent spending need lends support to theories where the salience of treatments affects economic choices. We conclude that implementation design and the targeting of transfers can greatly alter the effectiveness of stimulus policies.

Keywords: Marginal Propensity to Consume, Randomized Controlled Trial, Helicopter Money.

JEL: D12, E21, C93, D91.

*For helpful comments, we thank Adrien Auclert, Corina Boar, Michael Boutros, Peter Ganong, 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, 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 U.S. Census Bureau, the CEPR public economics symposium, the Federal Reserve Board, George Washington University, HEC, the Hebrew University of Jerusalem, the Institute for Fiscal Studies, the London School of Economics, the Paris School of Economics, the SITE conference, Stockholm School of Economics, UCL 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 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. We are grateful to Antoine Ding and Tyler Woodbury for excellent research assistance. AEA preregistration: <https://doi.org/10.1257/rct.9296-1.0>.

1 Introduction

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 et al., 2018](#), [Auclert et al., 2023b](#)). Despite extensive research, estimates of MPC out of transfers remain debated due to limitations arising from data availability and the sources of variation used for causal identification (e.g., [Parker et al., 2013](#), [Orchard et al., 2023b](#), [Borusyak et al., 2023](#)). In this paper, we overcome these challenges 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.¹ Going beyond standard estimation of MPC, we examine whether transfers with negative interest rates may yield larger MPCs, motivated by the fact that several countries have implemented large-scale household transfers with features akin to negative interest rates, e.g. time-limited consumption vouchers.²

Our experiment is designed with scalability and generalizability in mind. We randomly provide stimulus transfers to a sample of about 1,000 French individuals that is representative of the adult French population, and from which 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 euro zone. 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 Euros. 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 90,000 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.³ 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

¹A vast literature has examined MPCs out of various shocks, including typical income shocks ([Ganong and Noel, 2019](#)), lottery winnings ([Fagereng et al., 2021](#), [Golosov et al., 2021](#)), and recurring lump-sum payments ([Kuang, 2018](#)). Instead, we study one-time transfers comparable to those deployed to simulate the economy during an economic downturn.

²There 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 and Northern Ireland in 2021.

³Our primary motivation to study these treatments is their policy relevance and potential to raise MPCs. Several countries have disbursed time-limited prepaid cards to increase aggregate demand (e.g., [Hong Kong](#) and [Northern Ireland](#) in 2021), but the MPC in these programs is not known. 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 (see [Parker et al. \(2022b\)](#) and [Parker et al. \(2022a\)](#) for MPCs out of the 2020 stimulus payment, [Borusyak et al. \(2023\)](#) and [Orchard et al. \(2023b\)](#) for the 2008 stimulus payments, and [Orchard et al. \(2023a\)](#) for the 2001 stimulus).

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 marginal propensity to consume of 23 percent over one month (Fact #1). Moreover, we find that implementation design matters: the MPC is substantially higher for treatment groups where any remaining balance becomes unusable after three weeks, at 61% (or 70%, when conditioning on take-up), or where remaining balances are subject to the 10% negative interest rate every week, at 35% (Fact #2).

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 also do not find any significant effect of the additional “framing” treatment paragraph in the letter.

We next analyze the dynamics of the consumption response – the path of intertemporal MPCs, or iMPCs (Auclert et al., 2023b) – 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 consumption responses are concentrated in the first weeks following the transfer and are flat thereafter (Fact #3). We observe that the consumption response is concentrated early on even for non-durables.

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, 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 pre-treatment 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 (gender, age).

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,⁴ this fact allows us to use statistical deconvolution

⁴We provide empirical support for this assumption through auxiliary tests in Section 4.

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 4-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 4-week consumption by more than 52 percent of the transfer amount. These results again highlight the power of the choice of implementation design 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 is 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 et al., 2018, Auclert et al., 2023b). For example, in Kaplan et al. (2018), the MPC out of a \$300 transfer is 17% over a quarter and increases to about 32% 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%, but no further increases in spending in later periods.⁵ Laibson et al. (2022) note that durables require special treatment when analyzing the dynamic response 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.

We also show that our results are difficult to reconcile with agents being rational and treating money as fungible. 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. The transferred amount of 300 euros is well below the normal 3-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 non-fungibility of money (Hastings and Shapiro, 2013, 2018, Baugh et al., 2021, Geng et al., 2022, Gelman

⁵Our finding is consistent with quasi-experimental evidence on MPCs using the 2008 U.S. tax rebates. Using scanner data to document high-frequency spending responses to tax rebates, Borusyak et al. (2023) show that they are concentrated in the first month after the rebate.

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 et al., 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 non-binding 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 bi-modal 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 et al., 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 show that the design of transfers is very important to maximize MPCs. 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.⁶ 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 is smaller than by 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.

Related literature. This paper contributes to a vast literature that seeks to estimate marginal propensities to consume. The main MPC estimates used to discipline macro models are based on the staggered disbursement of tax rebates. The seminal papers analyzing staggered tax rebates in the United States (Johnson et al., 2006, Parker et al., 2013, Broda and Parker, 2014) found large MPCs, of about 50% over a quarter. However, these estimates are subject to two limitations. First, due to noise in survey data, it is challenging to reject small MPCs, to estimate the monthly dynamics of MPCs, as well as to estimate MPC heterogeneity across households. Second, 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% over a quarter (Borusyak et al., 2023, Orchard et al., 2023b). In this paper, we overcome these limitations by directly running an experiment, using high-frequency bank data with rich household covariates to measure the dynamics of the spending response as well as its heterogeneity across households. The MPC estimates for our baseline treatment are close to those obtained by Borusyak et al. (2023) and Orchard et al. (2023b).

Several papers use high-quality administrative data to measure the spending responses to a variety of economic shocks, including typical income shocks (Ganong et al., 2020), lottery wins (Fagereng et al.,

⁶Aguiar et al. (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.

2021, Golosov et al., 2021), or recurring payments (Kueng, 2018). Relative to these papers, we focus on transitory transfers that are similar to standard stimulus transfers. Since the marginal propensity to consume depends on the source of the shock, its persistence, and its magnitude, to draw lessons for macroeconomic stabilization policies it is important to directly study the spending response to typical 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 et al., 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 et al., 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 et al., 2021, Commault, 2022b). Fourth, a small literature estimates the spending response to consumption vouchers (Hsieh et al., 2010, Kan et al., 2017, Xing et al., 2023, Geng et al., 2022, Ding et al., 2023). A unique feature of our setting is to use an experiment to analyze how the marginal propensity to consume varies with implementation design choices, comparing the effects of standard transfers to transfers featuring an expiry date or a negative interest rate.

More broadly, this paper demonstrates the possibility of evaluating and improving the design of macroeconomic policy tools through experimental means. 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 2 presents the data and experimental design; Section 3 presents our main MPC estimates, establishing our first three key facts; Section 4 documents MPC heterogeneity, leading to our fourth and fifth key facts; Section 5 uses our five facts to draw lessons for macroeconomic models and stimulus policies.

2 Data and Experimental Design

In this section, we describe our dataset, our main variables, as well as the experimental design.

2.1 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. 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 comes from the French banking group Cr dit Mutuel Alliance F d rale.⁷ 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 socio-economic characteristics, as shown by Bounie et al. (2020) and Bonnet et al. (2023), who compare the bank sample to official statistics (see Appendix A.1). The data provide socioeconomic information about the household, transaction-level information for transaction accounts, transaction-level information for all payment cards linked to the accounts, debt and balances on non-transaction 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 Appendix A.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, have a known residential address, should not have accounts with another bank (according to the bank’s records), and are not deemed by the bank to be financially fragile.⁸ In order to obtain a population where we are able to measure spending well, we also exclude those individuals that have been using their debit card infrequently in the months prior to their 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. Random sampling for participation in the experiment is implemented at the individual – rather than household – level so that we can later on analyze heterogeneity based on individual characteristics, e.g. gender.

Variable definitions. Our main outcome variables 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.⁹ We winsorize weekly consumption spending using the non-treatment cards at the 99th percentile of the distribution, which is 1940 euros, 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. Appendix A.3 provides more detail on variable definitions.

Our estimated consumption responses are thus marginal propensities to spend (see Laibson et al., 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

⁷Cr dit Mutuel Alliance F d rale 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 Cr dit Mutuel Alliance F d rale’s mission as an “*entreprise   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*”).

⁸About 1% of households are deemed financially fragile by the bank.

⁹We 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 2 May (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 I Summary Statistics

	<i>N</i>	Mean	S.D.
Age of eligible household member	85,700.00	47.03	12.92
Number of eligible household members	85,700.00	1.15	0.36
Avg. monthly incoming transfers, 6m prior	85,685.00	2,654.04	1,439.56
Avg. monthly incoming salaries, social allowance, pensions, benefits, 6m prior	80,034.00	2,109.55	4,968.86
Dummy: has received unemployment benefits within 6m prior	85,685.00	0.14	0.35
Avg. current account balance, 6m prior	85,698.00	4,448.51	19,976.04
Avg. liquid savings, 6m prior	85,698.00	16,896.51	34,466.19
Avg. value of life insurance assets, 6m prior	85,698.00	5,867.43	32,466.19
Avg. net illiquid wealth, 6m prior	85,700.00	64,746.97	185,159.99
Avg. net liquid wealth, 1m prior	85,700.00	19,265.55	46,067.00.
Avg. total debt, 6m prior	85,698.00	33,300.8	55,007.44
Avg. consumer debt, 6m prior	85,698.00	2,388.27	5,194.01
Avg. mortgage debt, 6m 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
Avg. monthly consumption expenditures (cash, card payments), 1 year prior	85,690.00	1,205.62	658.22
Avg. monthly outgoing transfers (direct debits, debt payments, Subscriptions), 6m prior	85,635.00	518.12	1,001.24
Avg. total monthly consumption (broad measure)	85,700.00	2,039.91	1,563.41
Weekly consumption expenditure (cash and cards), total	2,571,000.00	417.66	435.02
Weekly consumption expenditure (broad measure), excl. treatment cards	2,571,000.00	619.50	1,315.05

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.

in the 6 months prior to the treatment (November 2021 to April 2022), per capita.

Summary statistics. Table I shows summary statistics of the main variables. The table illustrates the richness of the bank data and the large heterogeneity in observable characteristics. Appendix Tables A1 and A2 provide additional summary statistics.

2.2 Experimental Design

Treatment arms. From the set of eligible individuals, we randomly draw 915 participants over three treatment groups.

Treatment Group 1 (G1, $N = 379$) participants receive a MasterCard debit card linked to a new transactions account with an initial balance of 300 Euros. The card expires and becomes unusable six months after it has been sent, at which point 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 participant to purchase an item above 300 Euros

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% negative interest rate by decreasing the remaining balance on the account (i) by 30 euros if the remaining euro balance is in the interval $(200, 300]$, (ii) by 20 euros if the remaining euro balance is in the interval $[100, 200]$, and (iii) by 10 euro if the remaining balance is below 100 euros. If the remaining balance is below 10 euros, 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” (transl., for the original see Appendix B.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% a week) and group 2 (as the expiry amounts to a 100% 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 27th the cards (which from now on we will call “treatment cards” or “prepaid cards” to distinguish them from the households’ other means of payment) and accompanying instructions and explanations (see Appendix B.1), as well as pin codes were sent by post to the residential addresses of the selected individuals,¹⁰ with expected arrival on or around Monday May 2nd. 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 16th, and so on every week from then onward. For Treatment Group 2 participants, the card expires on Tuesday May 24th.

¹⁰The pin codes for the treatment cards are set by the bank to be the same as each participant’s main debit card.

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 3rd, and the remaining balances are transferred to the participants’ main bank accounts.

Take-up. Participants started using the card from May 2nd onward. Among the 915 treated households, 830 used the treatment card at least once before 6 October. 85 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.¹¹ 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.

2.3 Randomization Tests


We now present 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”.

Figure 1 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 results confirm the validity of the experimental design.

¹¹Two 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.

Figure 1 Randomization Tests



Notes: This figure reports the randomization tests for participation in the experiment, regressing a dummy for participation status on several household characteristics. We control for the number of eligible members in the households. The probability of being treated in the sample is 1%.

3 Main MPC Estimates: Facts #1, #2 & #3

In this section, we report our main MPC estimates. We first consider all treatment cards at once (Subsection 3.1). We then report estimates by card types (Subsection 3.2), 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 (Subsection 3.3). Finally, we report MPC estimates by framing group (Subsection 3.4).

3.1 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.


3.1.1 MPC Estimates from Raw Data

Figure 2 presents 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 spending increases rapidly and reaches about 250 euros after two months, i.e. they spend 84% 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 2 shows clear graphical evidence that treated households spend more upon

Figure 2 Spending Behavior in the Raw Data

A. Cumulative Spending on Prepaid Card



B. Average Total Spending, Weekly



Notes: This figure reports 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% confidence intervals for mean weekly spending are reported as shaded regions in panel B.

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 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.

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

3.1.2 MPC Estimates from Regression Specification

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

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

where Y_{it} is the outcome variable, usually consumption spending of household i in week t , the dummy $1(\tau \text{ weeks since } i \text{ treated})_{it}$ 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,¹² 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 et al., 2023](#)).


Results. The results are reported in Figure 3. Panel A report the estimates for the β_{τ} coefficients at a weekly frequency after treatment. The picture shows that, on average, participants’ spending increases by 75 euros in the first week, 25 euros in the second week, and 20 euros in the third week. The estimates are not significant in any of the following weeks.

Panel B 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 euros, corresponding to a marginal propensity to consume of $113/300 = 38$ percent. The figure suggests that the cumulative MPC is slightly reduced going forward, with a point estimate falling to about 45 euros after three months, or a cumulative MPC of 15 percent. As in other papers that compare recipients of transfers with non-recipients 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 horizon results from a small number of positive outliers in weekly pre-period consumption expenditures, which push up the household fixed effects and make subsequent expenditures appear small in comparison. Appendix Figure A1 shows FGLS estimates that effectively downweigh households with higher consumption volatility, and these estimates are flat over the corresponding horizon.

We implement two complementary analyses in the Online Appendix suggesting that the cumulative

¹²In practice, the estimates remain similar when we don’t include this control.

Figure 3 Pooled MPC Estimates



Notes: This figure reports our main MPC estimates. Panel A reports the weekly estimates, while panel B depicts the cumulative effects. 95% confidence intervals are reported, clustering the data at the household level.

MPC remains stable over long horizons after the initial spending burst. First, we estimate a specification analogous to (1), except that to reduce noise we use a seventh-order polynomial to model the weekly MPC response after treatment. Appendix Figure A2 plots the results over two quarters, showing that the cumulative MPC remains stable from the first month onward. The point estimates are statistically significant at the 95% level for the first five months. Furthermore, as another piece of evidence that households effectively save the rest of the transfer after the initial spending burst, we implement a simple test leveraging the intertemporal budget constraint. We estimate specification (1) with a different outcome, adding to our consumption measures the observed savings at the bank in checking accounts and liquid or illiquid savings accounts. Appendix Figure A3 shows that the cumulative sum of consumption and savings increases by 300 euros immediately at the start of the experiment and hovers around this value for the following quarters, confirming that we correctly measure all flows and that the consumption response is concentrated in the short run only. Appendix Figure A4 documents the increase in savings in liquid accounts for treated households, with a cumulative increase of about 100 euros after a month.

Although these results suggest that the cumulative MPC is stable after the first month, we of course lack the statistical power to detect long-term changes in consumption of a few euros per month. A permanent income consumer would consume the annuity value of the 300 euro transfer: with a 2% interest rate, this corresponds to an increase in consumption of about 6 euros 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 results are reported in the appendix. First, Appendix Figure A5 shows that the results are similar when leads are included in specification (1), with no sign of pre-trends. Third, Appendix Figure A6 documents the characteristics of the households who chose not to use the treatment card. Fourth, we obtain similar results with alternative consumption measures including direct debit transactions and wire transfers (Appendix Figure A7). Fifth, Appendix Figure A8 shows that the estimates remain very similar with a reweighing scheme to ensure that the sample is representative of the French population by

local area, income, gender, and age. Finally, Appendix Figure A9 reports the results without winsorizing consumption expenditures, which again yields very similar estimates.

3.2 MPC Estimates By Card Type

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

The estimates are reported in Figure 4 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 household to increase their weekly consumption spending in the two weeks after treatment by about 40 euros; the point estimates thereafter are close to zero and not significant. Panel B shows that households in Group 2—who receive a card that expires after three weeks—increase their weekly consumption significantly for the first three weeks after treatment, by about 65 euros in the first week, and by about 50 euros 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 –, who increase their spending immediately in the first week of the experiment, by about 130 euros, but not thereafter.

Figure 5 reports the cumulative spending response. The figure shows that the cumulative MPC for group 1 is much lower than for groups 2 and 3. After 4 weeks, the cumulative MPC for group 1 is 70 euros (23%), compared with 183 euros (61%) for group 2 and 106 euros (35%) for group 3. The estimates for MPCs over longer horizons become noisier due to idiosyncratic variation in consumption spending diluting the differences between treatment and control groups, but the point estimates for groups 2 and 3 remain large and above the spending response of group 1. Appendix Figure A1 reports FGLS estimates of cumulative MPCs by card, confirming that the spending response of groups 2 and 3 remain large and above that of group 1 at all horizons.

Panel B of Figure 5 shows MPC estimates conditional on using the card (at some point) to make purchases. The point estimate for the average 4-week MPC of group 2 participants is about 10 percentage points higher, at 70%. 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 strongly affect the MPC. The consumption response is concentrated early on and the MPC is not increasing with time.

These results establish our first three key facts, about the magnitude of the MPC by card type and its time profile:

Fact 1: The average one-month MPC on a cash-like transfer is 23%.


Fact 2: The design of the stimulus transfer can substantially affect the MPC. The average one-month MPC out of a transferred card whose remaining balance expires after three weeks is 61%. The average one-month MPC out of a transferred card whose remaining balance declines every week by about ten percent is 35%.

Fact 3: The consumption response to stimulus transfers is concentrated early on, in the first one to three weeks.


Additional results are reported in the Online Appendix. Appendix Figure A10 reports long-term MPC dynamics by card type over two quarters. We find that the cumulative MPC remains higher at longer horizons for Cards 2 and 3, compared to Card 1. Next, Appendix Figure A11 reports cumulative spending

Figure 4 MPC by Card Type, Weekly


A. Group 1, no restrictions on treatment card



B. Group 2, expiration after three weeks

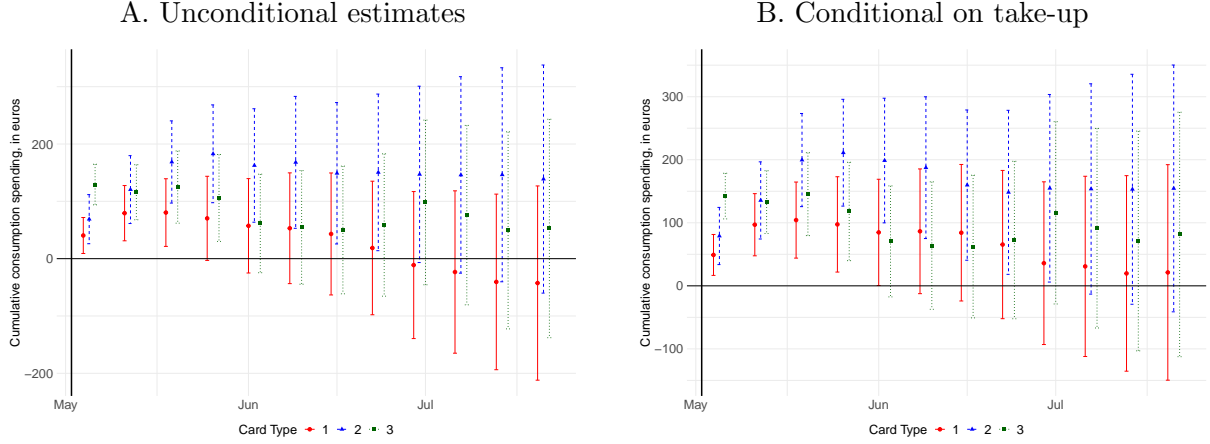


C. Group 3, negative rates every week



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 11:59pm. 95% confidence intervals are reported, clustering the standard errors at the household level.

Figure 5 MPC Estimates by Treatment Group



Notes: This figure reports 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 11:59pm. Panel A includes treated households that do not use the card in the treatment groups; panel B does not. 95% confidence intervals are reported, clustering the data at the household level.

on prepaid cards. In the short run, participants in Groups 2 and 3 spend down the prepaid card much faster than those in Group 1. But after three months, the total cumulative spending is approximately the same across the three card types. Thus, Group 1 participants don't have a lower overall MPC simply because they do not spend down the prepaid card. Rather, the substitution patterns differ depending on the card: Group 1 participants spend down about 85% of the prepaid card after three months but primarily cover expenses that they would have incurred otherwise as well. Appendix Figure A12 documents that socio-demographic characteristics are not significant predictors to identify the subset of households from Group 1 who did not use the card by July 1st, nor the subset of households from Groups 2 and 3 who lost significant funds. Finally, Appendix Figure A13 reports a formal test of MPC heterogeneity after 4 weeks, highlighting a statistically significant difference between Card 2 and Card 1.

3.3 Understanding the Spending Response

To better understand participant's spending behavior, we combine two approaches: survey questions to the treatment group,¹³ 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 6. 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 key motive for the money they saved out of the transfer (panel A of Figure 6), 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

¹³The survey was administered via the implementation partner's web platform. The survey response rate is 46% among the participants. We find that the average MPC is the same in this sub-sample as in the full sample.

each transaction, our data contains the 4-digit merchant category code (MCC) that is associated with the vendor. Panel C of Figure 6 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. 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), semi-durables (including, notably, apparel, footwear, and other textiles), durables (furniture, electronics, and durable household equipment, as well as leisure items and cars), and services. Appendix Table A3 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 6. Appendix Figure A14 reports the cumulative MPC response separately for durables and nondurables, showing a more sustained increase in spending on nondurables over time.

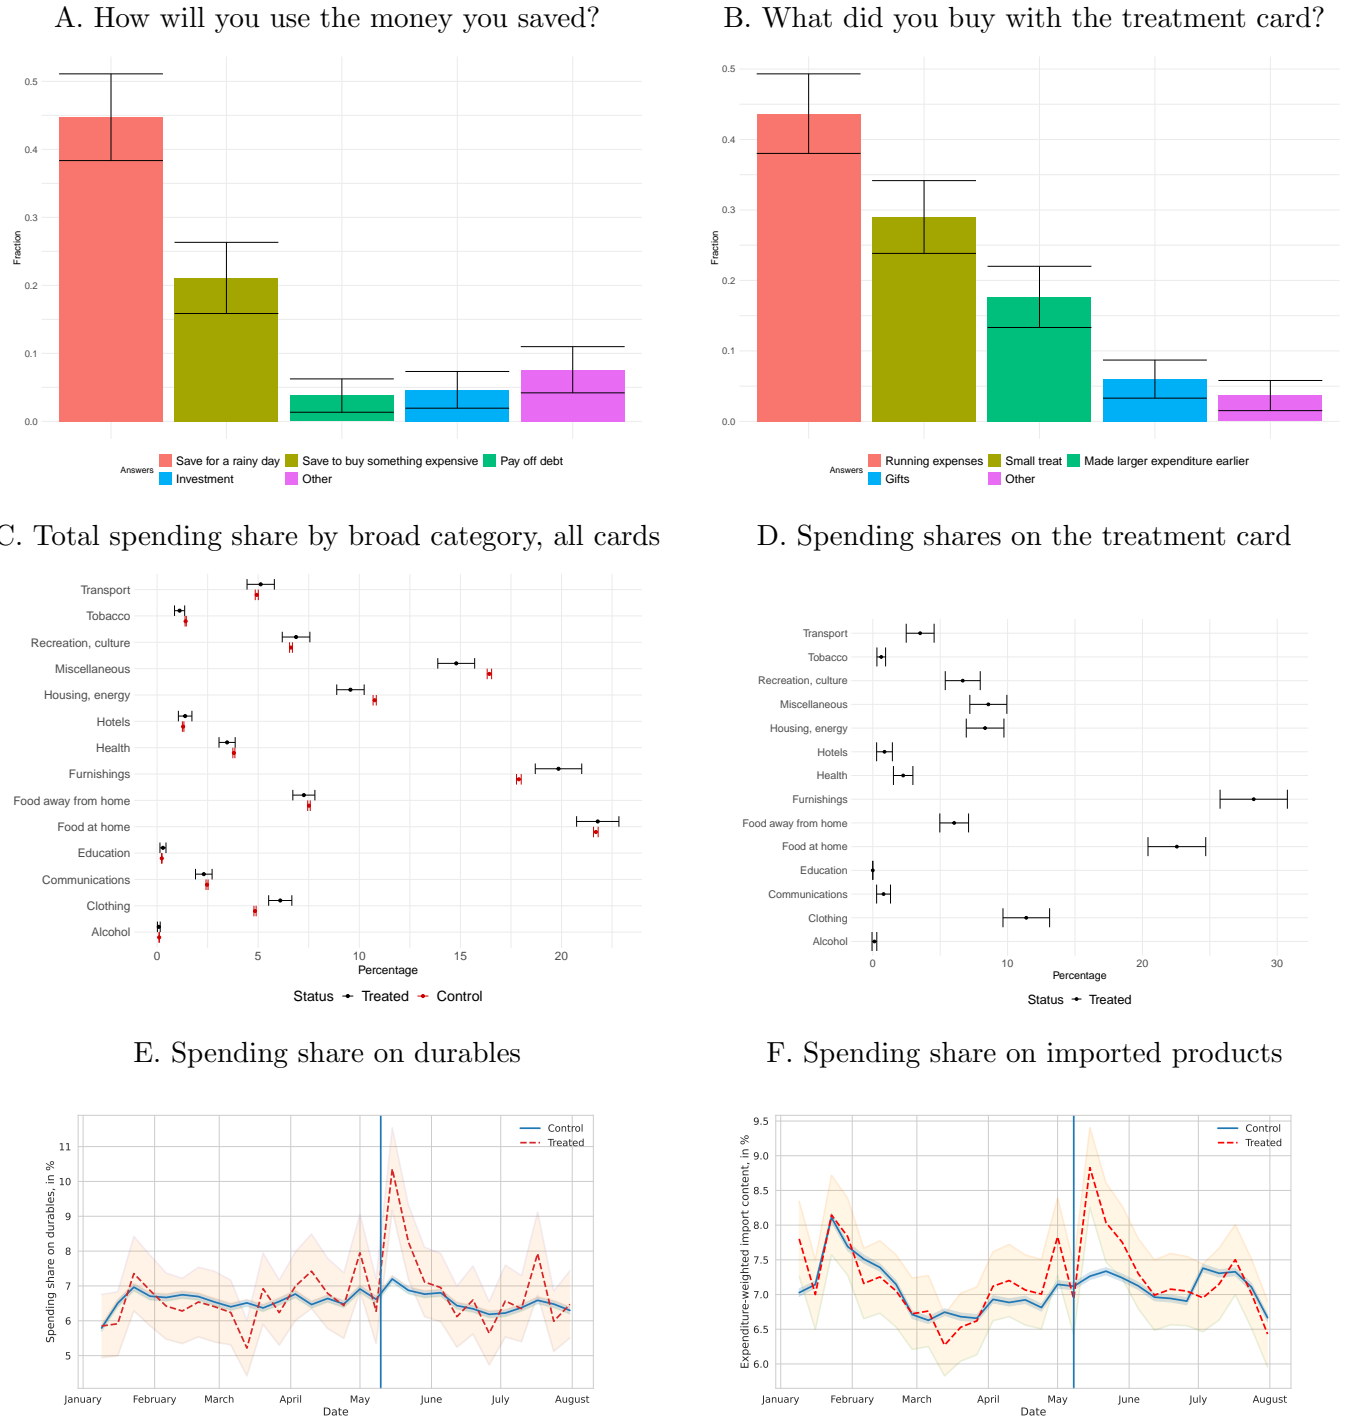
Finally, we analyze the propensity to spend on imports. We measure imports by mapping the MCC codes to the French input-output tables, which provide import penetration rates across categories. We find that treated households make purchases that have on average a higher import content, resulting in an increase in the weighted average import share from 7% in the control group to 9% in the treatment group (panel F). While these differences are statistically significant, they are modest from an economic perspective. For example, the 2 percentage point increase in the import share is too small to generate concerns that economic stimulus payment primarily increase aggregate demand for foreign trading partners.

By card type. We now repeat the analysis by card type. We first rely on the survey results, reporting the answers in panels A and B of Figure 7. We find that households in Groups 2 and 3 report that they are less likely to cover running expenses (panel A) and more likely to make large purchases earlier (panel B), consistent with the higher MPC estimates in the data.

Second, we decompose the consumption expenditure increase by durability. Table II shows the estimated fraction of the expenditure increase on each of the four durability categories by dividing the 4-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 4-week cumulative point estimate in a regression of all consumption expenditures (the baseline specification).

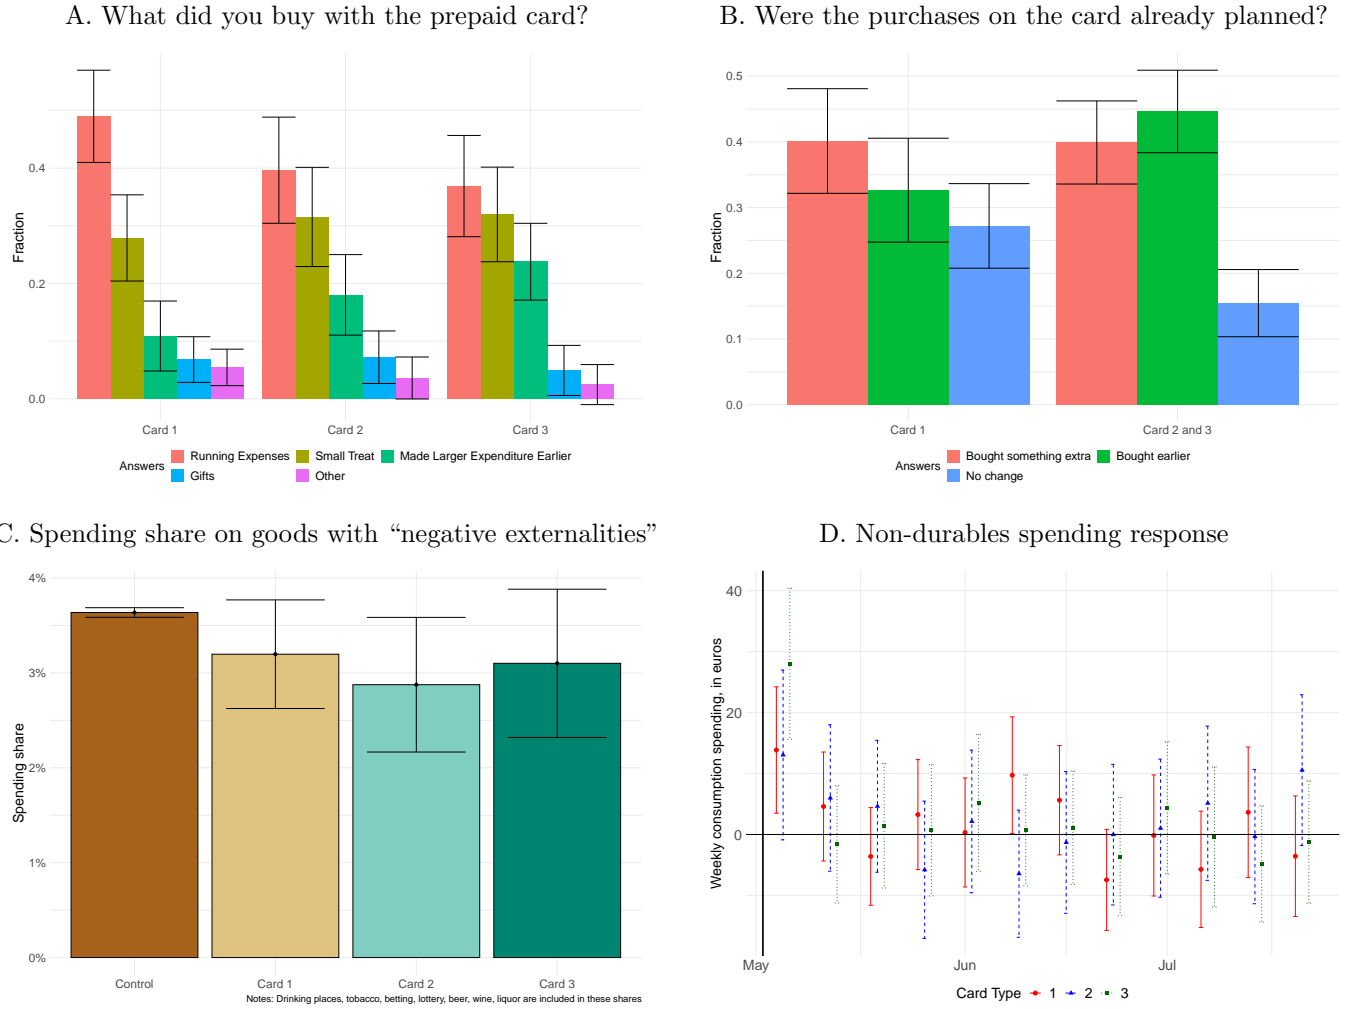
The results show that the increase in consumption expenditures goes to a larger extent into the purchase of semi-durable and durable goods. Households that receive treatment Card 2 channel a substantial fraction of the additional expenditure into personal services, whereas Card 3 households see a disproportionate increase in durables purchases. Considering all three cards together, Table II 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. Appendix Figure A15 shows the MPCs on these categories in the months following treatment. Appendix Figure A16 shows the response of durables over a longer time horizon for each card type, highlighting that the higher spending on durables for Card 3 is sustained in the following

Figure 6 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 PCE 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.

Figure 7 Understanding Participants' Spending Behavior by Card Type



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 depending on the prepaid card type. Panel C reports the spending share on treatment cards for the treatment groups, considering products that may have negative externalities (drinking, tobacco, betting, lottery). Panel D documents the spending response for non-durables by card.

quarters..

Table II Decomposition of 4-week MPC and Expenditure Shares by Type of Expenditure

	Decomposition of 4-week MPC			Expenditure shares over 4 weeks	
	Card 1 (1)	Card 2 (2)	Card 3 (3)	Treatment group (4)	Control group (5)
Nondurables	18.8%	9.2%	26.6%	28.8%	30.5%
Semi-durables	49.0%	30.9%	28.6%	19.9%	17.9%
Durables	24.1%	19.1%	49.8%	10.2%	7.8%
Services	13.6%	54.2%	11.6%	37.0%	38.9%
Not categorized	-5.6%	-13.5%	-16.9%	4.0%	4.7%

Notes: This table reports the average 4-week MPC on the row category divided by the total 4-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 4 weeks following treatment. The last row refers to expenditures that cannot be classified into the four main product categories, for example cash. Rows do not sum up exactly to 100% 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.

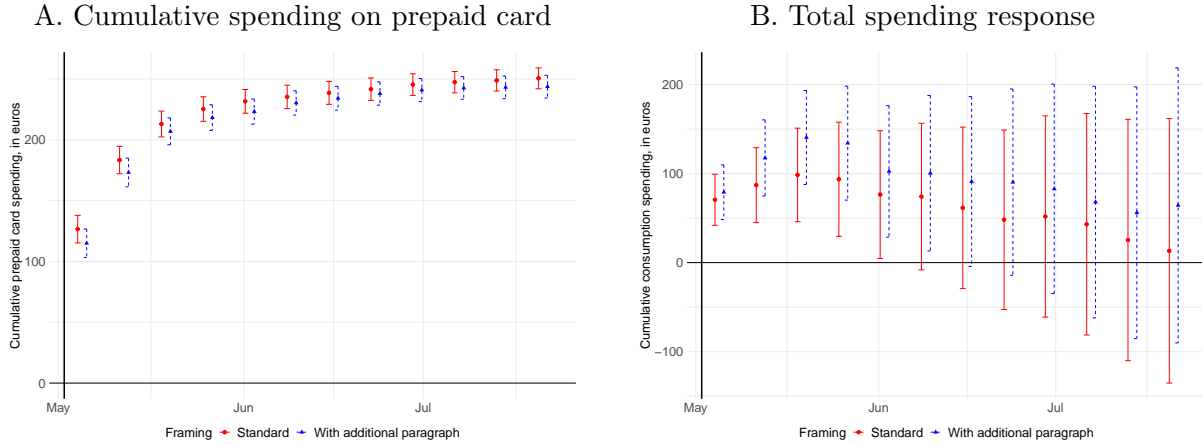
A potential concern is that the higher overall spending response with Cards 2 and 3 might come at the expense of the “quality” of spending. For example, a recent study by [Jaroszewicz et al. \(2022\)](#) finds that unconditional cash transfer taking place during Covid-19 sometimes had detrimental effects 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, better, and lottery products) differs across treatment arm. Panel C of Figure 7 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”. Panel D of Figure 7 rejects this hypothesis: participants in Groups 2 and 3 spend more on nondurables in the short run, and experience no fall in the longer-run. Finally, we use the survey to elicit the subjective impact of the transfer. In response to the question: “Has the transfer of the 300 euro card increased your happiness?”, only eight out of 391 respondents (or 2%) report that the transfer has not at all increased their happiness. 92% 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.

3.4 MPC Estimates By Framing Group

Finally, we evaluate whether households have different average MPCs depending on the framing of the intervention. Figure 8 reports the results. We find that households that received the additional framing

Figure 8 MPC Estimates by Framing Group



Notes: This figure reports MPC estimates depending on the framing of the intervention. We compare the consumption behavior of participants in who received the standard letter to those of participants who received a letter with an additional paragraph encouraging them to spend the money quickly and on local goods or services. Panel A reports spending patterns on the prepaid card, while panel B report the overall MPC. In panel B, 95% confidence intervals are reported, clustering the data at the household level.

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 < 10pp, and not statistically significant) as households that did not receive the framing treatment. 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 (Appendix Figure A17). Households in the framing group are however more likely to respond in the post-experiment survey that they had used the transfer to buy a “treat” for themselves (Appendix Figure A18), which is consistent with experimenter demand effects.

Thus, we conclude that implementation design choices are powerful tools to increase the MPC, while written framing treatments are unlikely to have large effects.

4 MPC Heterogeneity across Households: Facts #4 & #5

We now turn to the analysis of MPC heterogeneity across households. Estimating MPC heterogeneity is key both for policy – as policymakers may want to target certain 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 potential falsify certain models. We first document MPC heterogeneity by observable household characteristics (Subsection 4.1), establishing our fourth key fact about MPCs. Finally, we present estimates of the unconditional distribution of MPC across households with a deconvolution approach (Subsection 4.2), our fifth key fact about MPCs. We discuss the implications of our findings in Section 5.

4.1 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 marginal propensity to consume for households with different characteristics. Specifically, we estimate specifications of the form:

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

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 pre-treatment 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 5. 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 correspond to the sum of household-level current account and liquid saving deposits net of short-term debt (for instance consumption debt) at the bank. 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 pre-treatment consumption is measured as the average monthly consumption expenditure in the year prior to treatment, at the household level. Lastly, we define income at the household level as the sum of all incoming transfers.¹⁴ Except for net liquid wealth, which is measured at the beginning of the experiment, and average pre-period 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.¹⁵

Figure 9 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 the experiment. Panel B turns to illiquid wealth, depicting a negative relationship between MPCs and illiquid wealth quartiles. Despite these negative relationships, Appendix Figure A19 shows that the MPC remains high even for households who have substantial liquid wealth, a fact we will use later on when drawing

¹⁴Inflows above 15 000 euros are trimmed out at the household level.

¹⁵When 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 individual.

implications of our findings for consumption models. Figure A19 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. The MPC is particularly high for households in the lowest quartile, with a cumulative MPC of about 75% after three months. For the fourth quartile, the cumulative MPC is close to zero starting one month after treatment. Panel D reports that MPCs fall with household income. The MPC is again very high for the bottom quartile, around 75% after three months, while it is close to zero for the top quartile.

Finally, Panels E and F consider in turn age and gender. Panel E shows that households in the highest quartile of the age distribution (above 55) have a higher MPC, close to 60% after three month, while at this horizon the MPC is close to zero for households below the median age. Turning to gender, Panel F show that women have a much lower MPC than men. After three month, the cumulative MPC is close to zero for women while it is around 40% for men. Appendix Figure A20 shows that the higher MPC for women is observed regardless of the number of household members, and Appendix Figure A21 reports differences in spending composition by gender, showing that women spend more on food at home after receiving the prepaid card.

In sum, income and gender constitute the strongest sources of observable MPC heterogeneity among the predictors we consider. Appendix Figure A22 shows that the results for each of the six dimensions of heterogeneity we study remain similar when using a FGLS estimator to reduce noise.

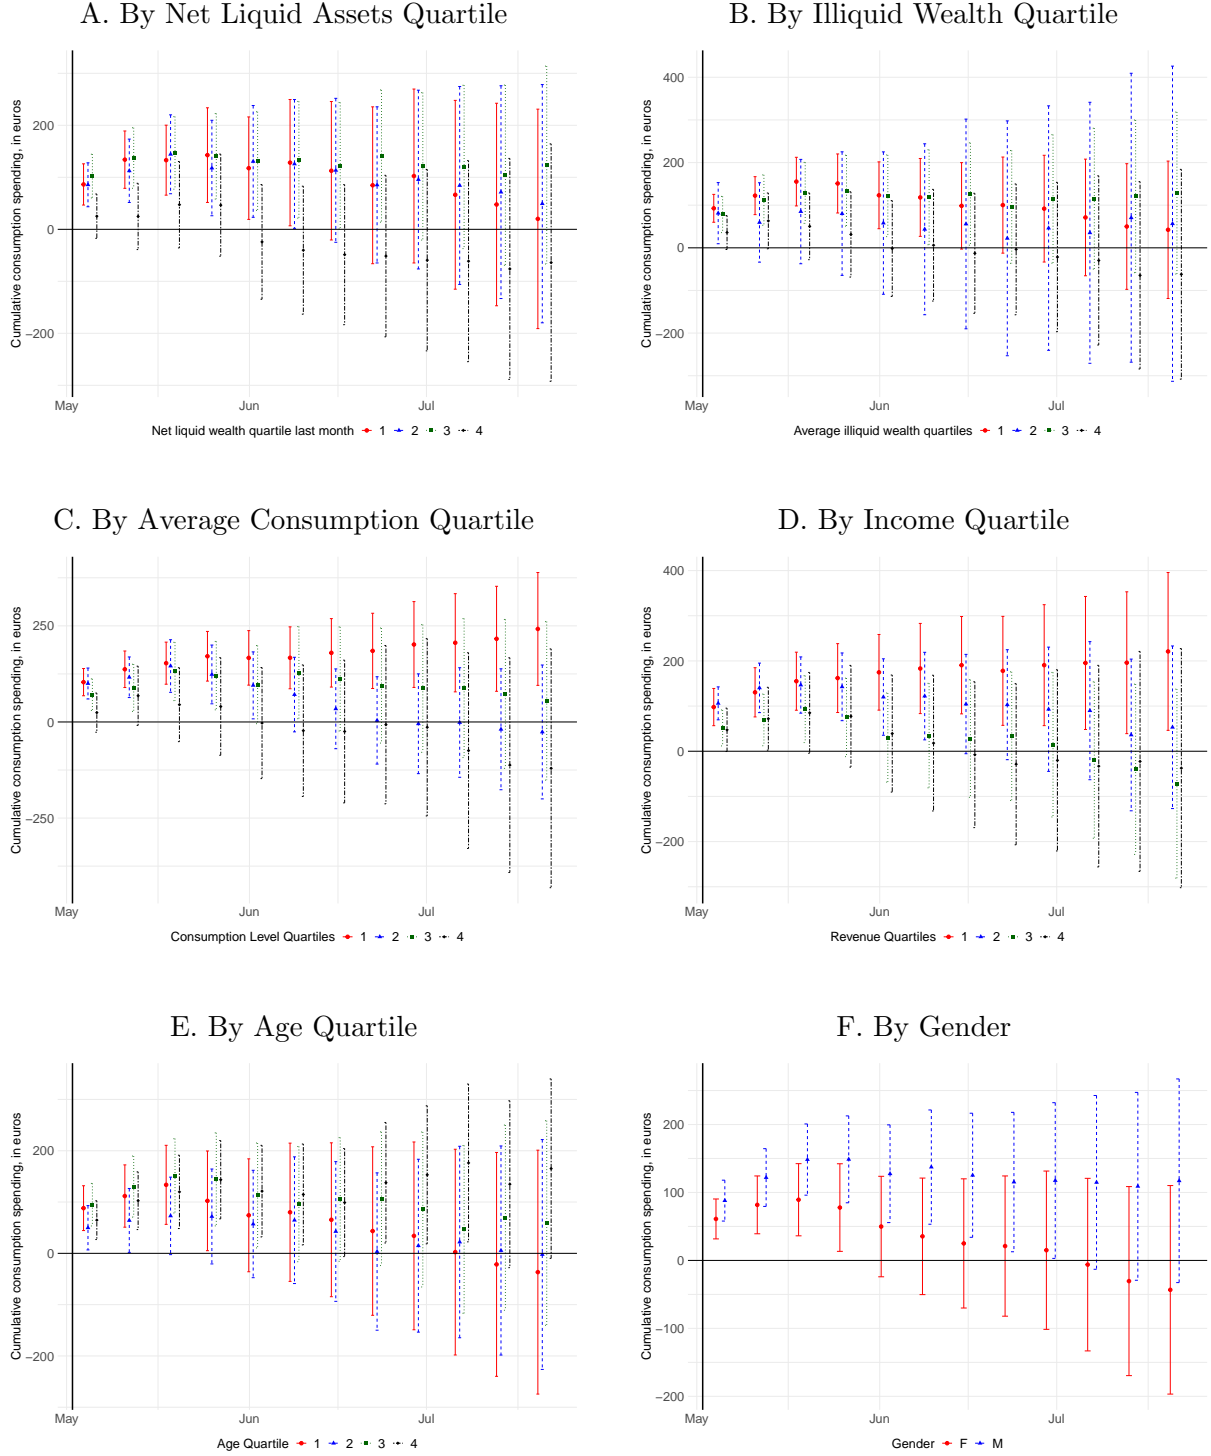
LASSO analysis. We now turn to a set of regressions that attempts to uncover which household characteristics are most relevant for explaining MPC heterogeneity. We implement specification (2) with all six variables (divided into quartiles when relevant) included jointly, as well as some additional variables (unemployment, local area characteristics, household size). In order to avoid over-fitting, 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 10 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 vs. 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, others 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.¹⁶ 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.

While leading macroeconomic model highlight the role of liquid and illiquid wealth as key predictors of treatment effect heterogeneity, our LASSO analysis show that other predictors are more powerful. We


¹⁶Results are similar when including only treatment group 1 and control group, see Appendix Figure A23.

Figure 9 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 and marital status. 95% confidence intervals, with standard errors clustered at the household level, are reported in all panels.

Figure 10 LASSO Estimates of Treatment Effect 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). The dashed vertical line shows the regularization parameter chosen by 5-fold cross validation.

further discuss the implications of these results for household targeting in Section 5.2.

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

Fact 4: MPCs strongly vary across observed characteristics, in particular by gender, age, and proxies for permanent income.

4.2 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

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

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

correspond 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 marginal propensity to consume is 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 non-experimental setting. To avoid confounding treatment effects with the potentially different tails of households of different size, 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 et al. \(2023\)](#).¹⁷ In the second step, we construct cumulatives of de-meaned weekly consumption expenditure:

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

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 Appendix Figure [A24](#), 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.

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 non-negative on its support and to integrate to one, restrictions that we also impose. Since deconvolution estimates often suffer from oscillating densities in the tails, [Yang et al. \(2020\)](#) 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.¹⁸ We obtain standard errors for the estimated quantiles of the treatment effect distribution

¹⁷We first estimate household effects α_i from all pre-treatment 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 $\hat{\alpha}_i$, i.e. we ensure that we have the same number of observations to estimate households fixed effects in treatment and control groups.

¹⁸An alternative, which we pursue in Appendix Figure [A25](#), is to estimate a model that is linear in log consumption. A

through a bootstrap of the entire two-step procedure.

Results. Figure 11 shows estimates for the distribution of 4-week MPCs by treatment group. The median MPCs are close to but slightly higher than the ATE estimates we obtained in Section 3.2 for each card type, at 28%, 81%, and 53% for groups 1, 2, and 3, respectively.¹⁹

The estimates show a substantial heterogeneity in the propensity to consume out of the transfer, with the bottom quartile having a 4-week MPC of less than 13% (card 1), 52% (card 2), and 30% (card 3), while the top quartile has MPCs above 48% (card 1), 103% (card 2), and 77% (card 3). The distribution of estimated treatment effects for group 2 first-order stochastically dominates the distribution of group 1. Appendix Figure A26 shows results for specifications where we drop non-negativity constraints and regularization, yielding similar findings. Appendix Figure A27 reports the results of the deconvolution by pooling together treatment cards 2 and 3, obtaining more precise estimates that confirm that the cards with negative rates yield MPCs that are higher than those associated with treatment card 1. Consistent with the large MPC heterogeneity uncovered by our deconvolution approach, Appendix Figure A28 plots the quantiles of spending on the treatment card and shows that there is a lot of heterogeneity in the speed at which households spend these funds.

These results establish our fifth key fact:

Fact 5: There is substantial heterogeneity in the unconditional marginal propensity to consume out of a windfall transfer, and a large fraction of households has a high MPC.


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 can also not 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

disadvantage of such a model is that the treatment effect distribution estimates cannot be directly interpreted as MPCs, but instead as cumulative percentage deviations from the household’s average level of consumption. Both approaches lead to similar shapes of treatment effect distributions.

¹⁹Prior work has also attempted to estimate the distribution of MPCs. Misra and Surico (2014) compare spending distributions of US households around the 2001 and 2008 tax rebates using 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 vs untreated households are very similar (Appendix Figure A24) is difficult to reconcile with negative MPCs in our data. Furthermore, Lewis et al. (2019) use clustering-based Gaussian Mixture linear models to estimate MPC heterogeneity following the 2008 tax rebates. In their model MPCs vary across groups; group memberships and average MPCs for each group are identified through parametric assumption on the error terms. In contrast to their approach, our treatment effect distribution is identified and estimated entirely nonparametrically.

Figure 11 Household-level Quantiles of the 4-week MPC Distribution



Notes: This figure reports the quantiles of the distribution of 4-week treatment effects by treatment group. Shaded regions are delineated by the 10th and 90th percent quantile of the bootstrapped simulated distribution of the corresponding moment.

distribution, shown in Appendix Figure A29, 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 et al. (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 pre-treatment period. We perform the deconvolution exercise separately by treatment group on each of those samples. The estimates, shown in Appendix Figure A30, 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.

5 Implications

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

5.1 Implications for Models

Our experiment is not designed to test any particular model of consumption, but to instead 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.

5.1.1 Benchmark Rational Models

We first compare our findings with the predictions of canonical “rational” models. In the Heterogeneous-Agent New Keynesian (HANK) model of Kaplan et al. (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 et al. (2018)): the estimated MPC is about 17% over a quarter (for a \$300 transfer), about 25% over two quarters, and 32% over one year; furthermore, the high MPCs are entirely 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.²⁰

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 4).²¹ In contrast, as previously mentioned the MPC response is much more long-lived in HANK and in canonical buffer-stock saving models.²² While spending on durables could in principle explain a short-run spending burst in a standard model (Laibson et al. (2022)), we find that the response is also concentrated in the short run for non-durables. We compare our MPC estimates to the standard calibration of the HANK model more formally in Appendix Figure A31, 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 9). In Appendix Figure A19, 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 et al. (2021) among lottery winners in Norway, and Baugh et al. (2021) in response to expected tax refunds in the United States.²³

Third, our estimates of the unconditional distributions of MPCs reveal that MPCs are high for a large majority of the population (Figure 11). For Group 1 participants, our estimates indicate that half of the population has a one-month MPC above 28%, and three quarters has an MPC above 13%. In contrast, in the benchmark HANK models high MPCs are concentrated among a smaller fraction of the population, namely agents hitting their borrowing constraint.

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 non-homothetic preferences, allowing for MPCs that vary with permanent income.

Finally, it is instructive to compare our estimates to those used in standard calibrations of HANK

²⁰See Achdou et al. (2022) for a characterization of how cumulative MPCs vary with the time horizon in the Aiyagari–Bewley–Huggett model.

²¹Consistent with our experimental finding, Borusyak et al. (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 average, Fagereng et al. (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.

²²Auclert et al. (2023b) and Angeletos et al. (2023) show how, for certain policy experiments in certain classes of structural macro models, average (intertemporal) MPCs are key statistics for the behavior of macroeconomic aggregates.

²³In contrast to these studies, which isolate transitory income shocks, Ganong et al. (2020) study the consumption response to typical month-to-month fluctuations in labor income and find an MPC close to zero for households with high liquid wealth.

Table III First-quarter MPX and MPC Estimates for Calibration of Macroeconomic Models

	Parker et al. (2013)	Broda and Parker (2014)	Borusyak et al. (2023)	Orchard et al. (2023b)	This paper, treatment group 1
	(1)	(2)	(3)	(4)	(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) through (4)) and for treatment Card 1 participants in our experiment (column (5)). The first row reports the marginal propensity to spend on all goods and services, while the second row focuses on nondurables alone. The third row follows the methodology of Laibson et al. (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) correspond to the lowest and highest values among the three rescaling methods used by Broda and Parker (2014) and Borusyak et al. (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 et al. (2023b). We compare our estimates to a larger set of papers in Appendix Figure A32.

models. In Table III, we summarize estimates of the consumption response to the 2008 tax rebates in the United States, and we compare it to our estimates. Following Laibson et al. (2022), we now draw a distinction between the observed marginal propensity to spend (denoted MPX) and the model-consistent, or “notional,” marginal propensity to consume that should be used as a target for macroeconomic models.²⁴ 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 et al. (2018) and Auclert et al. (2023a) 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 et al. (2023) and Orchard et al. (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 et al. (2023) and Orchard et al. (2023b). The last row of Table III reports the notional MPC that should be used as a target for macroeconomic models, following the methodology of Laibson et al. (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.²⁵ A potential avenue is to augment standard consumption models with certain behavioral frictions. For example, in recent work Boutros (2023) and Lian (2021) develop structural behavioral models in which high-liquidity households have large

²⁴The 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 et al. (2022) showed that the notional MPC is the relevant target, state-of-the-art macroeconomic models targeted nondurable MPX estimates.

²⁵See Wolf (2023) for a characterisation of the shape that intertemporal MPCs in HANK models can take, and the extent to which they can be well approximation by simple models with occasionally binding borrowing constraints (as in, e.g., Farhi and Werning, 2019).

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.

5.1.2 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 euros (which can be made with the treatment card), we find that 88% of households in Group 2 spent at least 300 euros 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 household to substitute current account spending for prepaid card spending. In other words, under the rational benchmark, we expect that the 3-week expiry date for most households in treatment group 2 should not be a binding constraint, i.e., their MPC should be similar to households in treatment group 1, in contrast with our findings.

Figure 12 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 non-treatment cards) but for some reason did not use the card. A non-negligible share of households in groups 2 and 3 have a high enough remaining balance on the treatment card to cover the day’s expenditures but choose 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 not started to use the card. But even after more than a week into the experiment, the ratio remains above 15%. Appendix Figure A33 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 multi-person households of whom only one has access to the treatment card.


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 and Roussanov (2023)) and in sociology (e.g. Zelizer, 1989) that emphasizes the non-fungibility of money.²⁶

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 et al., 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 non-binding 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.²⁷

²⁶In contrast to work that has found that the labelling of cash transfers has an impact on spending patterns (Beatty et al., 2014, Benhassine et al., 2015), in Section 3.4 we detected no significant effect of framing on the magnitude and composition of expenditures.

²⁷Loss aversion is another bias that has been widely studied (e.g., Tversky and Kahneman (1992)). 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.

Figure 12 A Simple Test of the Fungibility of Money



Notes: This figure shows the fraction of households that should have used the treatment card but did not, by day and card type. Specifically, the figure shows the fraction of households that satisfy the following conditions (i) at the start of the day, they have a higher remaining balance on the treatment card than the realized consumption expenditure on other cards during the day; (ii) they do not use the treatment card during that day; (iii) they have a nontrivial amount of money left on the treatment card (more than 20 euros); (iv) they have positive consumption expenditures on that day; (v) they use the treatment card at some point during the experiment. The results are reported separately for Card 2, which expires after three weeks, and Card 3, which implements a negative interest rate of approximately 10% on the remaining balance on the treatment card every Monday at 11:59pm.

Second, another class of models that has been used for macro policy analysis is models that feature two sets of agents, “savers” and “spenders”, who have low and, respectively, high MPCs (Campbell and Mankiw (1989)). While implementations of such models can be made to feature consumption responses that are concentrated very early on, they would also imply strongly bi-modal distributions of MPCs, which we do not find (Figure 11). 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 et al., 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 Appendix C, 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, 2021)). 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 faster, 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).²⁸

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

²⁸Analyzing 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.

anyway, implying a lower marginal propensity to consume 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.

5.2 Implications for Macroeconomic Stabilization Policies

Our results have two immediate implications for policy. First, the large difference in MPCs across treatment groups show that the design of stimulus transfers is important to maximize MPCs. Treatment cards with negative rates — in the form of an expiry date or a weekly negative interest rate — deliver much larger MPCs than a standard cash-like stimulus transfer. Note that because some money ends up being returned in treatment designs 2 and 3, the average consumption stimulus per euro spent by the transferer is actually larger than the MPC estimate reported above. In Appendix Figure A34 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 groups. The resulting effective stimulus at the 4-week horizon is about 75 cent per euro of net transfer for group 2, and about 40 cent per euro for group 3.

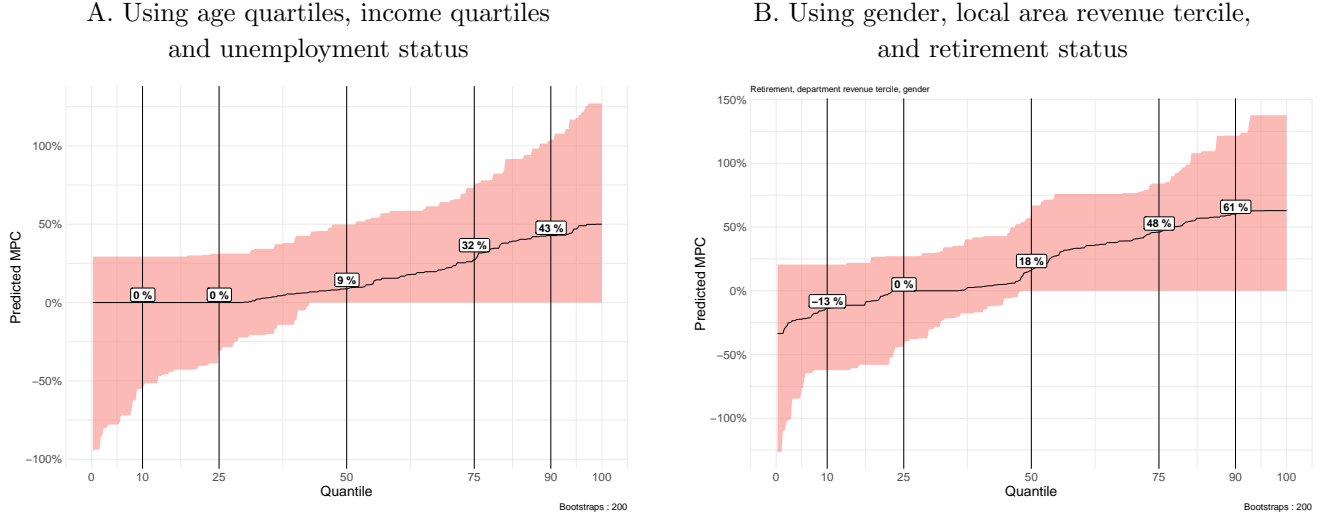
The external validity of our experimental estimates, and its broad applicability to high-income countries, appears plausible given that (i) we used a representative sample of the French population, and (ii) our estimates for group 1 are very similar to those obtained when studying the 2008 tax rebate response in the United States with robust estimators (Borusyak et al. (2023), Orchard et al. (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 et al. (2023) and Orchard et al. (2023b) for the 2008 stimulus payments. Using prepaid card 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. It is also worth noting that short-term interest rates were close to zero at the time when our experiment was implemented (see Appendix Figure A35), indicating the possible potency of particular types of stimulus policies even in a liquidity trap.²⁹

Second, our estimates of MPC heterogeneity have implications for the targeting of transfers by observable household characteristics. We documented in Section 4.1 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 4.1 with two sets of characteristics that policymakers might be able to observe as regressors, and estimate the distribution of

²⁹If 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.

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 13 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% confidence intervals are obtained by bootstrap and shown as shaded regions.

Figure 13 plots the distribution of predicted MPCs. In the first exercise, reported in panel A, we run LASSO with age quartiles, income quartiles, and unemployment status. This panel shows that by using these observables it is possible to identify households with substantially above-average MPCs. For example, 10% of households are predicted to have an MPC above 43%. In the second exercise, shown in panel B, we predict the MPC distribution using gender, local area revenue tercile, and retirement status. The top ten percent of households have an MPC of 61%. 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%. Thus, our estimates highlight that implementation design choices are a more powerful tool, compared to targeting, to increase the recipients' average MPC. In addition, targeting may raise political economy or fairness considerations that are avoided by providing a treatment card with an expiry date to all.

6 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% with a standard treatment card, without negative interest rates. Second, the design of the transfer matters: the one-month MPC is higher when treatment cards feature a negative interest rate: 61% when the remaining balance is reduced to zero after three weeks, and 35% when the remaining balance is reduced by approximately 10 percent every week. Third, the spending responses are concentrated early on, in the first one 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 model 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.

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.

References

- Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll, “Income and wealth distribution in macroeconomics: A continuous-time approach,” *The review of economic studies*, 2022, 89 (1), 45–86. 20
- Aguiar, Mark A, Mark Bils, and Corina Boar, “Who are the Hand-to-Mouth?,” Technical Report 2023. 6
- Altonji, Joseph G, Todd E Elder, and Christopher R Taber, “Selection on observed and unobserved variables: Assessing the effectiveness of Catholic schools,” *Journal of political economy*, 2005, 113 (1), 151–184. 4.2
- Angeletos, George-Marios, Chen Lian, and Christian K Wolf, “Can Deficits Finance Themselves?,” *Working paper*, 2023. 22
- Auclert, Adrien, Bence Bardóczy, and Matthew Rognlie, “MPCs, MPEs, and multipliers: A trilemma for New Keynesian models,” *Review of Economics and Statistics*, 2023, 105 (3), 700–712. 5.1.1
- , Matthew Rognlie, and Ludwig Straub, “The intertemporal keynesian cross,” *NBER Working Paper 25020*, 2023. 1, 22
- Baugh, Brian, Itzhak Ben-David, Hoonsuk Park, and Jonathan A Parker, “Asymmetric consumption smoothing,” *American Economic Review*, 2021, 111 (1), 192–230. 1, 5.1.1, 21, 5.1.2, 28, 29, C
- Beatty, Timothy KM, Laura Blow, Thomas F Crossley, and Cormac O’Dea, “Cash by any other name? Evidence on labeling from the UK Winter Fuel Payment,” *Journal of Public Economics*, 2014, 118, 86–96. 26
- Benhassine, Najy, Florencia Devoto, Esther Duflo, Pascaline Dupas, and Victor Pouliquen, “Turning a shove into a nudge? A “labeled cash transfer” for education,” *American Economic Journal: Economic Policy*, 2015, 7 (3), 86–125. 26
- Blundell, Richard, Luigi Pistaferri, and Ian Preston, “Consumption inequality and partial insurance,” *American Economic Review*, 2008, 98 (5), 1887–1921. 1
- Bonnet, Odran, Etienne Fize, Tristan Loisel, and Lionel Wilner, “How does fuel demand respond to anticipated price changes? Quasi-experimental evidence based on high-frequency data,” *Working Paper*, 2023. 2.1, A.1
- , Tom Olivia, and Théo Roudil-Valentin, “En 2020, la chute de la consommation a alimenté l’âépargne, faisant progresser notamment les hauts patrimoines financiers: quelques résultats de l’âexploitation de données bancaires,” *Note de conjoncture*, 2021. A.1
- Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, “Salience theory of choice under risk,” *The Quarterly journal of economics*, 2012, 127 (3), 1243–1285. 1, 5.1.2

- , —, and —, “Salience and consumer choice,” *Journal of Political Economy*, 2013, 121 (5), 803–843. 1, 5.1.2
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess**, “Revisiting event study designs: Robust and efficient estimation,” *arXiv preprint arXiv:2108.12419*, 2023. 1, 3, 5, 3.1.2, 4.2, 21, 5.1.1, 5.1.1, 5.2
- Bounie, David, Youssouf Camara, Etienne Fize, John Galbraith, Camille Landais, Chloe Lavest, Tatiana Pazem, Baptiste Savatier et al.**, “Consumption dynamics in the covid crisis: real time insights from French transaction bank data,” *Covid Economics*, 2020, 59, 1–39. 2.1, A.1
- Boutros, Michael**, “Windfall Income Shocks with Finite Planning Horizons,” *Working Paper*, 2023. 21, 5.1.1, 5.1.2
- Broda, Christian and Jonathan A Parker**, “The economic stimulus payments of 2008 and the aggregate demand for consumption,” *Journal of Monetary Economics*, 2014, 68, S20–S36. 1, 5.1.1, 5.1.1
- Bunn, Philip, Jeanne Le Roux, Kate Reinold, and Paolo Surico**, “The consumption response to positive and negative income shocks,” *Journal of Monetary Economics*, 2018, 96, 1–15. 1
- Campbell, John Y and N Gregory Mankiw**, “Consumption, income, and interest rates: Reinterpreting the time series evidence,” *NBER macroeconomics annual*, 1989, 4, 185–216. 1, 5.1.2
- Chaisemartin, Clement De and Xavier Haultfoeuille**, “Two-way fixed effects estimators with heterogeneous treatment effects,” *American Economic Review*, 2020. 3.1.2
- Commault, Jeanne**, “Does consumption respond to transitory shocks? Reconciling natural experiments and semistructural methods,” *American Economic Journal: Macroeconomics*, 2022, 14 (2), 96–122. 1
- , “How Do Persistent Earnings Affect the Response of Consumption to Transitory Shocks?,” Technical Report, HAL 2022. 1
- Ding, Jing, Lei Jiang, Lucy Msall, and Matthew J Notowidigdo**, “Consumer-Financed Fiscal Stimulus: Evidence from Digital Coupons in China,” Technical Report, National Bureau of Economic Research 2023. 1
- Ellison, Glenn and Alexander Wolitzky**, “A search cost model of obfuscation,” *The RAND Journal of Economics*, 2012, 43 (3), 417–441. C
- Evans, George W and Garey Ramey**, “Expectation calculation and macroeconomic dynamics,” *The American Economic Review*, 1992, pp. 207–224. C
- Fagereng, Andreas, Martin B Holm, and Gisle J Natvik**, “MPC heterogeneity and household balance sheets,” *American Economic Journal: Macroeconomics*, 2021, 13 (4), 1–54. 1, 5.1.1, 21
- Farhi, Emmanuel and Iván Werning**, “Monetary policy, bounded rationality, and incomplete markets,” *American Economic Review*, 2019, 109 (11), 3887–3928. 25

- Fuster, Andreas, Greg Kaplan, and Basit Zafar**, “What would you do with \$500? Spending responses to gains, losses, news, and loans,” *The Review of Economic Studies*, 2021, 88 (4), 1760–1795. [1](#)
- Ganong, Peter and Pascal Noel**, “Consumer spending during unemployment: Positive and normative implications,” *American economic review*, 2019, 109 (7), 2383–2424. [1](#), [3.3](#), [A.3](#)
- , **Damon Jones, Pascal J Noel, Fiona E Greig, Diana Farrell, and Chris Wheat**, “Wealth, race, and consumption smoothing of typical income shocks,” Technical Report, National Bureau of Economic Research 2020. [1](#), [23](#)
- Garber, Gabriel, Atif R Mian, Jacopo Ponticelli, and Amir Sufi**, “Consumption Smoothing or Consumption Binging? The effects of government-led consumer credit expansion in Brazil,” *Working Paper*, 2022. [3.3](#)
- Gelman, Michael**, “What drives heterogeneity in the marginal propensity to consume? Temporary shocks vs persistent characteristics,” *Journal of Monetary Economics*, 2021, 117, 521–542. [6](#)
- , “The Self-Constrained Hand-to-Mouth,” *Review of Economics and Statistics*, 2022, 104 (5), 1096–1109. [5.1.2](#)
- and **Nikolai Roussanov**, “Managing mental accounts: Payment cards and consumption expenditures,” *Working Paper*, 2023. [1](#), [5.1.2](#)
- Geng, Hao, Ce Shi, and Michael Zheng Song**, “Evaluating Hong Kong Consumption Voucher Scheme,” 2022. [1](#)
- Golosov, Mikhail, Michael Graber, Magne Mogstad, and David Novgorodsky**, “How Americans respond to idiosyncratic and exogenous changes in household wealth and unearned income,” Technical Report, National Bureau of Economic Research 2021. [1](#)
- Hastings, Justine and Jesse M Shapiro**, “How are SNAP benefits spent? Evidence from a retail panel,” *American Economic Review*, 2018, 108 (12), 3493–3540. [1](#), [5.1.2](#), [29](#)
- Hastings, Justine S and Jesse M Shapiro**, “Fungibility and consumer choice: Evidence from commodity price shocks,” *The quarterly journal of economics*, 2013, 128 (4), 1449–1498. [1](#), [5.1.2](#), [29](#)
- Hsieh, Chang-Tai, Satoshi Shimizutani, and Masahiro Hori**, “Did Japan’s shopping coupon program increase spending?,” *Journal of Public Economics*, 2010, 94 (7-8), 523–529. [1](#)
- Ilut, Cosmin and Rosen Valchev**, “Economic agents as imperfect problem solvers,” *The Quarterly Journal of Economics*, 2023, 138 (1), 313–362. [1](#), [5.1.2](#)
- Jappelli, Tullio and Luigi Pistaferri**, “Fiscal policy and MPC heterogeneity,” *American Economic Journal: Macroeconomics*, 2014, 6 (4), 107–136. [1](#)
- Jaroszewicz, Ania, Jon Jachimowicz, Oliver Hauser, and Julian Jamison**, “How effective is (more) money? Randomizing unconditional cash transfer amounts in the US,” Technical Report 2022. [3.3](#)

- Johnson, David S, Jonathan A Parker, and Nicholas S Souleles**, “Household expenditure and the income tax rebates of 2001,” *American Economic Review*, 2006, 96 (5), 1589–1610. [1](#)
- Kan, Kamhon, Shin-Kun Peng, and Ping Wang**, “Understanding consumption behavior: Evidence from consumers’ reaction to shopping vouchers,” *American Economic Journal: Economic Policy*, 2017, 9 (1), 137–153. [1](#)
- Kaplan, Greg and Giovanni L Violante**, “A model of the consumption response to fiscal stimulus payments,” *Econometrica*, 2014, 82 (4), 1199–1239. [1](#), [4.1](#), [5.1.1](#)
- , **Benjamin Moll, and Giovanni L Violante**, “Monetary policy according to HANK,” *American Economic Review*, 2018, 108 (3), 697–743. [1](#), [5.1.1](#), [D](#)
- Kueng, Lorenz**, “Excess sensitivity of high-income consumers,” *The Quarterly Journal of Economics*, 2018, 133 (4), 1693–1751. [1](#), [5.1.1](#)
- Laibson, David**, “Golden eggs and hyperbolic discounting,” *The Quarterly Journal of Economics*, 1997, 112 (2), 443–478. [1](#), [5.1.2](#)
- , **Peter Maxted, and Benjamin Moll**, “Present bias amplifies the household balance-sheet channels of macroeconomic policy,” Technical Report, National Bureau of Economic Research 2021. [1](#), [5.1.2](#)
- , —, and —, “A simple mapping from MPCs to MPXs,” Technical Report, National Bureau of Economic Research 2022. [1](#), [2.1](#), [5.1.1](#), [5.1.1](#), [24](#)
- Lewis, Daniel J, Davide Melcangi, and Laura Pilossoph**, “Latent heterogeneity in the marginal propensity to consume,” *FRB of New York Staff Report*, 2019, (902). [1](#), [19](#)
- Lian, Chen**, “Mistakes in future consumption, high MPCs now,” Technical Report, National Bureau of Economic Research 2021. [5.1.1](#)
- Maxted, Peter**, “Present bias in consumption-saving models: A tractable continuous-time approach,” Technical Report, Mimeo 2020. [1](#), [5.1.2](#)
- McDowall, Robert A**, “Consumption behavior across the distribution of liquid assets,” *Working paper*, 2020. [5.1.2](#)
- Misra, Kanishka and Paolo Surico**, “Consumption, income changes, and heterogeneity: Evidence from two fiscal stimulus programs,” *American Economic Journal: Macroeconomics*, 2014, 6 (4), 84–106. [1](#), [19](#)
- Nakamura, Emi and Jón Steinsson**, “Five facts about prices: A reevaluation of menu cost models,” *The Quarterly Journal of Economics*, 2008, 123 (4), 1415–1464. [6](#)
- and —, “Identification in macroeconomics,” *Journal of Economic Perspectives*, 2018, 32 (3), 59–86. [1](#)
- Olafsson, Arna and Michaela Pagel**, “The liquid hand-to-mouth: Evidence from personal finance management software,” *The Review of Financial Studies*, 2018, 31 (11), 4398–4446. [5.1.1](#)

- Orchard, Jacob, Valerie A Ramey, and Johannes F Wieland**, “Using Macro Counterfactuals to Assess Plausibility: An Illustration using the 2001 Rebate MPCs,” *Working Paper*, 2023. [3](#)
- , —, and **Johannes Wieland**, “Micro MPCs and macro counterfactuals: the case of the 2008 rebates,” Technical Report, Working Paper 2023. [1](#), [3](#), [5.1.1](#), [5.1.1](#), [5.2](#), [C](#)
- Oster, Emily**, “Unobservable selection and coefficient stability: Theory and evidence,” *Journal of Business & Economic Statistics*, 2019, *37* (2), 187–204. [4.2](#)
- Parker, Jonathan A and Nicholas S Souleles**, “Reported effects versus revealed-preference estimates: Evidence from the propensity to spend tax rebates,” *American Economic Review: Insights*, 2019, *1* (3), 273–90. [1](#)
- , **Jake Schild, Laura Erhard, and David Johnson**, “Economic Impact Payments and Household Spending During the Pandemic,” *Brookings Papers on Economic Activity*, 2022. [3](#), [5.2](#)
- , —, —, and —, “Household spending responses to the economic impact payments of 2020: Evidence from the consumer expenditure survey,” *NBER Working Paper*, 2022. [3](#), [5.2](#)
- , **Nicholas S Souleles, David S Johnson, and Robert McClelland**, “Consumer spending and the economic stimulus payments of 2008,” *American Economic Review*, 2013, *103* (6), 2530–2553. [1](#), [5.1.1](#), [5.1.1](#)
- Schennach, Susanne M**, “Recent advances in the measurement error literature,” *Annual Review of Economics*, 2016, *8*, 341–377. [4.2](#)
- Shapiro, Matthew D and Joel Slemrod**, “Consumer response to tax rebates,” *American Economic Review*, 2003, *93* (1), 381–396. [1](#)
- and —, “Did the 2008 tax rebates stimulate spending?,” *American Economic Review*, 2009, *99* (2), 374–379. [1](#)
- Shefrin, Hersh M and Richard H Thaler**, “The behavioral life-cycle hypothesis,” *Economic inquiry*, 1988, *26* (4), 609–643. [5.1.2](#), [28](#), [C](#)
- Straub, Ludwig**, “Consumption, savings, and the distribution of permanent income,” *Working paper*, 2019. [4.1](#), [5.1.1](#)
- Thaler, Richard H**, “Anomalies: Saving, fungibility, and mental accounts,” *Journal of economic perspectives*, 1990, *4* (1), 193–205. [5.1.2](#), [C](#)
- Tversky, Amos and Daniel Kahneman**, “Advances in prospect theory: Cumulative representation of uncertainty,” *Journal of Risk and uncertainty*, 1992, *5*, 297–323. [27](#)
- Wolf, Christian K**, “Interest rate cuts vs. stimulus payments: An equivalence result,” *Working Paper*, 2023. [25](#)

Xing, Jianwei, Eric Zou, Zhentao Yin, Yong Wang, and Zhenhua Li, “”Quick Response” Economic Stimulus: The Effect of Small-Value Digital Coupons on Spending,” *American Economic Journal: Macroeconomics*, 2023. [1](#)

Yang, Ran, Daniel W Apley, Jeremy Staum, and David Ruppert, “Density deconvolution with additive measurement errors using quadratic programming,” *Journal of Computational and Graphical Statistics*, 2020, *29* (3), 580–591. [4.2](#), [D](#)

Zelizer, Viviana A, “The social meaning of money: special monies,” *American journal of sociology*, 1989, *95* (2), 342–377. [5.1.2](#), [5.1.2](#)

– , *The social meaning of money: Pin money, paychecks, poor relief, and other currencies*, Princeton University Press, 2021. [5.1.2](#)

Appendix to “Five Facts about MPCs”

Johannes Boehm, *Sciences Po & CEPR*

Etienne Fize, *PSE & IPP*

Xavier Jaravel, *LSE & CEPR*

November 2023

A Data Appendix

In this appendix, we discuss the representativeness of the data, the data structure, and the exact definition of our variables for replication purposes.

A.1 Sampling and Representativeness

We build on a sample of households that the bank drew in June 2020, using the following steps. First, in order to be eligible for inclusion in the sample, the bank had to be the main bank used by the households (i.e., households could be using multiple banks but must have located their main assets, credits and income at the bank). Second, households had to be client of the bank in January 2019. Third, French overseas territory and employees of the bank were excluded of the sampling process. Finally, the sampling procedure drew clients from cells at the regional (“*département*”) \times age bin level. Specifically, 94 different *départements* and six age bins were used: 18-25, 26-35, 36-45, 46-55, 56-65 and 66+ years. For the largest 31 *départements* 1,000 households per cell were selected, then 500 for the next 26 *départements* and finally 100 for the least populated *département*. The initial sample size was around 300,000 households. The sample was never renewed and, because of attrition, the sample size decreased over time. We received remote access to anonymized versions of the data that start in January 2019.

This dataset is by design representative of the population of clients at the bank. Bounie et al. (2020) and Bonnet et al. (2023) find that the sample is also broadly representative of the French population along several characteristics, with some slight differences. Specifically, compared to the French population, the bank sample is younger, with fewer retired people, features a higher share of individuals out of the labor force,¹ and a higher share of single households. The distribution of spending (and the ratio of spending over income) by income deciles in the bank sample are in line with the French consumption survey (“*Budget des Familles*”). The trends in card spending and liquid bank account balances also match macro aggregates from the French national accounts (see Bounie et al., 2020 and Bonnet et al., 2021). Bonnet et al. (2023) also show that the customer base of the bank is representative of the French population in terms of financial wealth and disposable income.

Our paper focuses on a sub-sample of this panel. In order for an individual to be eligible in our experimental draw, a number of conditions have to be satisfied:

1. Age above 26 and below 75 years at the time of draw, and be in legal capacity, i.e. not having a legal guardian (“*majeur non protégé*”)

¹The “inactive” category includes students, the unemployed, and any other person with no professional activity.

2. Resident at the address registered with the bank
3. Account is a personal account, not a professional account. According to the records of the bank, the customer is either entirely or mostly banking with them.
4. The checking account is active and the account holder is in good standing with the bank ("*compte courant ouvert et sain*"), and there were movements on the checking account within the last 10 days.
5. At least one transfer received every month over the course of the last six months; at least 10 transfers received in 2020, and at least 10 transfers received in 2019.
6. At least one payment made every month in the last six months, in 2020, and in 2019.
7. Has an active bank card that it has used at least once within the prior 20 days.
8. Is using internet/distance banking ("*détenteur d'un contrat banque à distance*") with at least one connection within the prior 20 days. Has an email address registered on file.

Constraints 1, 2, and 8 are imposed for operational and ethical reasons. Constraints 3, 4, 5, 6, and 7 are imposed to exclude people whose consumption expenditure patterns are poorly captured through the data available to us.

A.2 Data Structure

The data is divided in six different tables:

- The first table is at the individual \times month level, it contains socio-demographic information for all individuals in the household at month t .
- The second table is at the household \times month level. It contains information on the balance of all different bank accounts of the household (current account, liquid savings account, life insurance and illiquid savings). The table also provides information on household debt (total debt, and by subcategories such as mortgage debt or consumption debt), and on the sum of incoming and outgoing banking movement for some categories of banking operations (checks, cash withdrawal, card purchases). Finally, it includes information on payment or other financial difficulties faced by the household, such as overdraft.
- The third table is at the household \times operation level. This dataset provides information (time, amount) on all banking operations, i.e. all inflows and outflows. These flows cover a vast range of transactions, including card purchases, wire transfers, checks, and direct debit. The bank also provides information on incoming wire transfers. The bank classifies the incoming wire transfers into distinct categories: pensions, unemployment insurance, government subsidies, and salaries.
- The fourth table is at the household \times operation level. This dataset gives information on all card transactions. Compared to the previous table, this table gives more information for the card transactions (e.g., the Merchant Category Code (MCC) for the purchase). Moreover, while the previous table records the date at which there is a banking movement, this table records the date at which the transaction occurs (i.e., when the card is actually used). The two dates may differ for several

reasons. For instance, some household choose to have a deferred debit, where the banking movements comes at the end of each month for all card transactions. The difference can also comes from delays from either the bank (in case the purchase is made on a bank holiday, or on a Sunday) or from the merchant (for instance, for fuel and gas purchases).

- The fifth table is at the household \times operation level and provide provides a classification of all direct debit operations (phone bill, water bill...).
- The sixth table is at the household \times period level. This table is a snapshot of all real estate wealth owned by the household, according to the bank’s records. The information was collected twice, in September of 2020 and in November of 2021.

All of these tables can be joined thanks to an anonymized household identifier.

A.3 Variable Definitions

Our main variables are defined as follows:

- Consumption expenditure, per week: sum of card purchases and cash withdrawals of the household within the week (according to the third table described in Appendix A.2). We construct our winsorized weekly consumption spending by winsorizing regular client transactions at the 99th percentile (1940 euros in a week) and add up treatment card expenses in a week at the household level.
- Consumption expenditure on non-durable / durable / semi-durable goods or services: sum of card purchases and cash withdrawals of the household within the week linked to Merchant Category Codes classification to include only expenditure on specific categories of consumption expenditure (following the classification of Ganong and Noel (2019)).
- Treatment card expenditure: sum of household treatment card purchases within a week. Treated household who have effectively used the treatment card will have at least one week with positive value during the treatment period. Control group households have a value of zero for this variable.
- Regular card expenditure: sum of card purchases and cash withdrawals of the household within the week (according to the third table described in Appendix A.2) that are not identified as treatment card expenditure for treated households.
- Withdrawals: sum of cash withdrawals of the household within the week (according to the third table described in Appendix A.2).
- Weekly consumption expenditure, broad measure: sum of card purchases and cash withdrawals of the household within the week (according the third table described in Appendix A.2) plus all other outflows (direct debits, wire transfers, etc).
- Take-up dummy: Time-invariant dummy variable, equal to one for treated households who have used the treatment card at some point during the treatment period.

- Number of eligible individuals in the household: number of individuals in the household that would have been eligible to receive the treatment, at the time of randomization.
- Unemployed: time-invariant dummy, equal to one for households that receive at least one transfer from the unemployment benefits agency (“Pôle emploi”) within the 6 months prior to treatment.
- Aggregation of individual characteristics to the household level: We aggregate individual characteristics to the household level by using the characteristic of the eligible household member. For control group households composed of two eligible people, we randomly choose one person’s characteristic to represent the entire household. For treatment group households with two eligible members, we use the characteristic of the individual that has been chosen (at random) as treated. The relevant individual characteristics are as follows:
 - Age: time invariant variable that corresponds to the age of the individual.
 - Location: capture the département where the household lives.
 - Location type: this variable measures whether the household lives in rural, periphery or urban areas.
 - Occupation: this variable measures whether the individual works in one of the following occupations: farmers, artisans, executives, intermediate professions, employee, worker, retired, unemployed/students.
- Number of household members: this variable is used to correct time invariant characteristics like income and wealth. We account for the presence of children to compute a unit of consumption (UC) for each household. Following the OECD scale, we attribute 1 UC to a first adult in a household, 0.5 UC to the following one and 0.3 UC for every child below 14 years old.
- Variables for time-invariant heterogeneity analysis: all of these variables are divided by the sum of the unit of consumption in the household (see above):
 - Household monthly mean expenditure: average of monthly card expenditures in a week for 1 year before treatment.
 - Household monthly income: average monthly inflows to the household’s bank account within the six months prior to treatment. Individual transactions value above 15,000 euros are trimmed.
 - Household identified income: average monthly identified transfers to the household’s bank account within the six months prior to treatment. Identified transfers correspond to inflows that the bank could classify as salary, social transfers, pensions, or unemployment benefits.
 - Household wealth: we build two versions of the variables listed below over different time horizons, taking averages over either one month or six months prior to treatment.
 - * Household bank current account: current account balance prior to treatment. This variable captures the bank account funds that the household can use at any point in time.

- * Household liquid saving accounts: liquid saving balance prior to treatment. This variable capture the funds available on liquid, tax-free savings accounts with instant access: *Livret A*, *Livret d'épargne populaire*, *Livret Jeune*, and philanthropic savings accounts, etc.
- * Household life insurance accounts: life insurance balance prior to treatment.
- * Household illiquid savings: illiquid saving accounts balance prior to treatment, including the “share savings plan” (*Plan d'épargne en action*).
- * Net liquid wealth: sum of the balances of the household’s current account and liquid saving accounts at the end of the month, net of consumer debt.
- * Net illiquid wealth: sum of the balances of the household’s illiquid saving accounts, share dealing accounts, and brokerage accounts at the end of the month, net of mortgage debt.
- * Household real estate wealth: real estate wealth reported by the household during a survey conducted in November 2021.

B Experimental Design Appendix

In this appendix, we describe the letter sent to the participants, as well as the survey administered in June 2022.

B.1 Letter Sent to the Participants

The letter sent to participants is printed on the bank’s letterhead and is personally addressed to the selected participant:

Vous avez été sélectionné pour participer à une étude et ainsi bénéficier d’une enveloppe d’un montant de 300 EUR, qui vous est offerte.*

En effet, afin de contribuer au débat économique, le CIC participe à une étude scientifique menée par le Conseil d’Analyse Economique (CAE) et financée par l’Agence Nationale de la Recherche (ANR). L’objectif de cette initiative est d’étudier, dans le cadre d’une politique destinée à favoriser la relance économique, les comportements de dépenses des personnes lorsqu’une somme d’argent leur est distribuée gratuitement.

Le CIC veille à la protection des données de ses clients. Toutes les analyses réalisées dans le cadre de cette étude seront effectuées sur des données strictement anonymisées sur les seuls systèmes d’information sécurisés du CIC. Il s’agit des mouvements bancaires, de la situation financière et de données socio-économiques.

Ce montant de 300 EUR sera utilisable au moyen d’une carte de paiement spécifique. Cette carte vous sera adressée gratuitement par courrier postal dans les prochains jours.

- *Le code confidentiel de cette carte est identique à celui de la carte que vous possédez déjà. Vous pouvez le retrouver dans votre espace personnel en ligne, sur l’application mobile ou le site internet www.cic.fr.*
- *Cette carte peut être utilisée auprès des établissements affichant les logos CB ou Mastercard, ainsi que pour des achats en ligne, dans la limite du solde disponible.*
- *Il n’est pas possible de retirer des espèces, ni d’effectuer des dépôts.*
- *Le suivi des opérations et le solde disponible sur cette carte sont consultables dans votre espace personnel en ligne, sur l’application mobile ou sur le site internet www.cic.fr.*
- *Les conditions générales d’utilisation qui régissent votre carte actuelle, s’appliquent également à cette carte (CG.03.20).*

Translation:

You have been selected to participate in a study and, as a result, benefit from an amount of 300 EUR, which is offered to you.*

Indeed, in order to contribute to the economic debate, CIC is participating in a scientific study conducted by the Council of Economic Analysis (CAE) and funded by the National Research Agency (ANR). 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.

CIC ensures the protection of its clients' data. All analyses carried out as part of this study will be performed on strictly anonymized data on CIC's secure information systems. This includes banking transactions, financial situation, and socio-economic data.

This amount of 300 EUR will be available for use through a specific payment card. This card will be sent to you free of charge by postal mail in the coming days.

- *The confidential code for this card is the same as the one for the card you already possess. You can find it in your personal online space, on the mobile application, or on the website www.cic.fr.*
- *This card can be used at establishments displaying the CB or Mastercard logos, as well as for online purchases, up to the available balance.*
- *It is not possible to withdraw cash or make deposits. The operations and available balance on this card can be checked in your personal online space, on the mobile application, or on the website www.cic.fr.*
- *The general terms of use that govern your current card also apply to this card (CG.03.20).*

The next paragraph contains information that is specific to the treatment group.

For treatment group 1:

La carte fonctionne jusqu'au 03/10/2022. Si vous ne dépensez pas l'intégralité du montant de 300 EUR avant cette date, le solde restant sera automatiquement transféré sur votre compte courant habituel du CIC.

Transl.: *The card is valid until 10/03/2022. If you do not spend the entire amount of 300 EUR before this date, the remaining balance will be automatically transferred to your regular current account at CIC.*

For treatment group 2:

L'objectif de cette expérience est d'encourager une hausse de la consommation à court terme, dans le cadre d'une politique économique de relance. Pour cette raison, la carte fonctionne jusqu'au 23/05/2022 à 23 heures 59. Il ne sera plus possible d'utiliser les fonds après cette date limite; les fonds inutilisés seront perdus.

Transl.: *The objective of this experiment is to encourage an increase in short-term consumption, as part of an economic policy for recovery. For this reason, the card is valid until 05/23/2022 at 11:59 PM. It will no longer be possible to use the funds after this deadline, and any unused funds will be lost.*

For treatment group 3:

L'objectif de cette expérience est d'encourager une hausse de la consommation à court terme, dans le cadre d'une politique économique de relance. Pour cette raison, le montant disponible de la carte est débité automatiquement d'un certain montant chaque lundi à 23 heures 59 (à partir du lundi 09/05/2022). Le montant débité dépend du solde restant à ce moment, avec un montant débité plus élevé lorsque le solde restant est plus élevé afin d'encourager une consommation rapide. Ainsi, le solde disponible sera diminué :

- *de 30 EUR si le solde restant est supérieur à 200 EUR ;*
- *de 20 EUR si le solde restant est entre 100 EUR et 200 EUR ;*
- *de 10 EUR si le solde est inférieur à 100 EUR (le débit correspond au solde restant si celui-ci est inférieur à 10 EUR).*

Par exemple, si vous dépensez le montant de 300 EUR avant le lundi 09/05/2022 à 23 heures 59, le solde restant est nul et aucun montant ne sera débité. Si vous dépensez seulement 50 EUR avant le lundi 09/05/2022 à 23 heures 59, le solde disponible sera diminué de 30 EUR et le solde disponible le mardi 10/05/2022 à 00h00 sera de 220 EUR (= 300 – 50 – 30).

Transl.: *The goal of this experiment is to promote an increase in short-term consumption as part of an economic policy for recovery. For this reason, the available amount on the card is automatically debited by a certain amount every Monday at 11:59 PM (starting from Monday, 05/09/2022). The debited amount depends on the remaining balance at that moment, with a higher amount debited when the remaining balance is higher, to encourage rapid consumption. As a result, the available balance will be reduced as follows:*

- *by 30 EUR if the remaining balance is above 200 EUR;*
- *by 20 EUR if the remaining balance is between 100 EUR and 200 EUR;*
- *by 10 EUR if the remaining balance is below 100 EUR (the debit amount will be equal to the remaining balance if it is below 10 EUR). For example, if you spend the full amount of 300 EUR before Monday, 05/09/2022, at 11:59 PM, the remaining balance will be zero, and no amount will be debited. If you only spend 50 EUR before Monday, 05/09/2022, at 11:59 PM, the available balance will be reduced by 30 EUR, and the available balance on Tuesday, 05/10/2022, at 12:00 AM will be 220 EUR (= 300 – 50 – 30).*

Next, a paragraph that depends on whether the participant is part of a framing group. Participants that are not in the framing group receive the following message:

Vous êtes totalement libre d'utiliser le montant de 300 EUR comme vous le souhaitez.

Transl.: *You are completely free to use the amount of 300 EUR as you wish.*

Participants that are in the framing group receive instead:

Bien que vous soyez libre d'utiliser le montant de 300 euros comme vous le souhaitez, nous vous invitons à:

- *dépenser l'argent aussi rapidement que possible;*

- *acheter des produits fabriqués en France et des services qui soutiennent l'emploi local, car l'objectif de ce transfert est la relance de l'économie française, en encourageant la consommation de produits made in France;*
- *acheter des produits ou services que vous n'achèteriez pas habituellement (autres que vos dépenses courantes) afin d'augmenter vos dépenses totales, et ainsi de contribuer à la relance économique, plutôt que de couvrir des dépenses déjà prévues.*

Transl.: *Although you are free to use the amount of 300 euros 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 "made in France" products;*
- *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.*

All groups conclude with the following:

L'utilisation de cette carte n'entraîne aucun frais pour vous. Si vous ne souhaitez pas participer à cette étude, n'utilisez pas la carte et détruisez la. En utilisant la carte, vous acceptez de participer à l'étude. En vous remerciant pour votre confiance, votre conseiller CIC se tient à disposition pour répondre à toutes vos questions.

Transl.: *The use of this card does not incur any fees for you. If you do not wish to participate in this study, do not use the card and destroy it. By using the card, you agree to participate in the study. Thank you for your trust; your CIC advisor is available to answer any questions you may have.*

The footnote is as follows:

** L'étude est menée et a été définie par une équipe scientifique du CAE et financée par l'Agence Nationale de la Recherche. Les critères de sélection des participants, l'utilisation des cartes, les données étudiées et la durée de l'étude qui s'étend du 01/10/2021 au 03/10/2022 ont été définis par le CAE. Les 1 000 participants qui bénéficient de la somme de 300 EUR ont été tirés au sort sous contrôle d'huissier.*

Transl.: *The study is conducted and has been defined by a scientific team from the CAE and funded by the National Research Agency. The criteria for selecting participants, the use of the cards, the data studied, and the duration of the study, which extends from 10/01/2021 to 10/03/2022, have been determined by the CAE. The 1,000 participants who are receiving the sum of 300 EUR have been randomly selected under the supervision of a bailiff.*

B.2 Survey Questions

Participants were contacted by email with the following message:

Bonjour,

Vous avez récemment fait appel au service Etudes, Satisfaction et Qualité pour vous accompagner dans le cadre du projet : Enquête de satisfaction CAE / CARTE DE PAIEMENT 300 euros. Afin d'améliorer la qualité de nos prestations, nous sollicitons votre retour d'expérience. Nous vous proposons donc une courte enquête composée de quelques questions. Cela vous prendra moins de 5 minutes pour y répondre.

[Hyperlink: Répondre à l'enquête]

Nous vous remercions par avance.

Notre équipe reste bien évidemment à votre disposition.

Bonne journée.

Le service Etudes, Satisfaction et Qualité

Translation:

Hello,

You recently used the Studies, Satisfaction, and Quality service to assist you in the context of the project: Satisfaction Survey CAE / 300 Euro Payment Card. In order to improve the quality of our services, we would appreciate your feedback. We invite you to participate in a short survey consisting of a few questions. It will take you less than 5 minutes to complete.

[Hyperlink: Respond to the survey]

Thank you in advance.

Our team remains at your disposal.

Have a great day.

The Studies, Satisfaction, and Quality service

The full text of the questionnaire is available from the authors upon request.

C A Stylized Model

In this appendix, we present a simple model to make predictions that qualitatively match our main empirical results.

Overview. The model relies on three key ingredients: (i) mental accounts; (ii) search costs; (iii) memory (i.e., certain agents can make purchases without incurring search costs). In our model, spending the prepaid card on “windfall consumption” involves a key tradeoff: (1) it delivers a utility boost λ because of mental accounts; but (2) it requires incurring search costs to find suitable windfall purchases, except for some agents who remember “windfall purchase opportunities”.

We summarize below the three key results we obtain in the model, thanks to the three key ingredients:

- For Group 1 participants, the spending response is concentrated in the short run.
 - Key channel: for Group 1 participants who remember suitable “windfall purchase opportunities”, it is optimal to purchase immediately.
- For Group 2 participants, the spending response is larger than for Group 1.
 - Key channel: while Group 1 participants smooth search costs across a large number of periods, Group 2 participants search for and buy more windfall consumption goods and services in period 0 using the prepaid card, in order to spend it down before it expires in period 1.
- For Group 3 participants, the spending response lies in between Group 1 and Group 2.
 - Key channel: the search costs are higher for Group 3 participants, leading them to prefer to spend relatively more on regular consumption (compared to Group 2) rather than incurring very high search costs for windfall consumption goods and services in period 0.

Note that the model below produces these results with a common “mental account” parameter for all three groups – rather than assuming different types of mental accounts for each card, which would be mechanical. Although the simple model below makes predictions that qualitatively match the main patterns in our data, it is not meant to provide a quantitative match of the estimated marginal propensities to consume.

Setting. Agents in the model receive a prepaid card and optimize consumption at an infinite horizon. There are three treatment groups, motivated by our experiment. Group 1 participants have access to the remaining balance on the prepaid card for $T + 1$ periods. In contrast, Groups 2 and 3 both lose access to the remaining balance on the card after the initial period.²

Preferences. The agent optimizes consumption over an infinite horizon with two goods, general consumption c_t and “windfall consumption” g_t . The utility function is:

$$U = \sum_{t=0}^{\infty} \beta^t [\lambda_t \cdot v(g_t, s_t) + u(c_t) - \psi(s_t)],$$

²In our experiment, Group 3 has a high negative interest rate, which we could model as well. However, for simplicity we can model Group 3 by varying the search cost parameter, as discussed below.

where s_t denotes search costs that help increase the marginal utility of windfall consumption, while λ_t is a marginal utility shifter for windfall consumption. s_t captures the idea that agents must incur calculation costs to find windfall consumption goods and services that suit their tastes (in the spirit of [Evans and Ramey \(1992\)](#) and [Orchard et al. \(2023b\)](#)); the convexity of costs is akin to [Ellison and Wolitzky \(2012\)](#).

The parameter λ_t captures the idea that utility for specific windfall goods and services may shift because of mental accounting (in the spirit of [Shefrin and Thaler \(1988\)](#), [Thaler \(1990\)](#), and [Baugh et al. \(2021\)](#)). Specifically, we assume that the households who receive a prepaid card in our experiment perceive it as a windfall, akin to a gift, and that they incur a utility boost if they spend this windfall on unplanned “windfall consumption” (e.g., going to a fancy restaurant, going out more frequently than usual, purchasing a treat, etc.) rather than on regular consumption.³ To capture the idea that the marginal utility of spending on windfall consumption goods and services is larger when spending from the prepaid card, we use a simple functional form:

$$\lambda_0 = \lambda \cdot 1_{\{p_{gg_0}=G_0-G_1>0\}},$$

with $\lambda > 0$, i.e. marginal utility is positive when the agent buys a positive amount of treats using the prepaid card, while it is null otherwise. G_t denotes the amount available on the prepaid card at time t , equal to 300 euros in our experiment. Note that in the functional form for λ_0 , the budget constraint is intertwined with the utility function. This approach is standard in models of mental accounts, going back to [Shefrin and Thaler \(1988\)](#). In this way, the utility derived from a purchase differs depending on the income source used to make the purchase, which is the very idea of a “mental account.” Our chosen functional form simply means that the marginal utility of purchasing windfall consumption goods and services is positive only when spending down the prepaid card to purchase these products.

For subsequent periods, the functional form is the same for Group 1, i.e. $\lambda_t = \lambda 1_{\{p_{gg_t}=G_t-G_{t+1}>0\}}$. We set $\lambda_t = 0$ for $t \geq 1$ for Groups 2 and 3, because these participants lose the remaining balance on the prepaid card after the initial period.

We make additional simple parametric assumptions to obtain closed-form solutions:

$$\begin{aligned}\psi(s_t) &= \frac{\kappa}{\eta} s_t^\eta, \eta > 1, \\ u(c) &= \log(c), \\ v(g_t, s_t) &= \min(g_t, s_t + e_{it}),\end{aligned}$$

where e_{it} denotes an individual-specific “endowment” of ideas about which windfall consumption goods and services to purchase. The functional form for $v(g_t, s_t)$ captures the idea that to enjoy windfall consumption the agent needs to purchase g_t units of windfall consumption but also to incur search costs s_t , or leverage their search endowment e_{it} . We set $e_{i0} = e_0 > 0$ for a fraction of agents, i.e. these agents know which windfall consumption goods and services to purchase – as if they remembered past

³This assumption is a line with the economics and sociology literature on the non-fungibility of money. To illustrate our assumption, consider a different context: our assumption means that a households receiving money for Christmas or a birthday will disproportionately spend them on windfall consumption goods and services (rather than regular consumption, e.g. laundry supplies).

opportunities to consume such goods and services. These agents can purchase up to $e_0 < G$ units of windfall consumption without the need to incur search costs. The endowment is set to zero for other agents.⁴

Furthermore, we assume that search costs are larger for Group 3, which we will study below with comparative statics on κ . This is motivated by the fact that, in our experiment, Group 3 participants faced a large negative interest rate after a week only: search costs can be seen as particularly costly for this group given the limited time available.

Thus, spending the prepaid card on windfall consumption involves a key tradeoff in the model: (i) it delivers a utility boost λ ; but (ii) it requires incurring search costs. To obtain simple closed-form solutions, we study the case of quadratic search costs, i.e. $\eta = 2$.

Budget constraint. The household faces a stream of per-period income z growing at rate g . The amount available on the prepaid card is denoted G . The budget constraint is:

$$\sum_{t=0}^{\infty} \left(\frac{1+g}{1+r} \right)^t z + G = \sum_{t=0}^{\infty} \frac{1}{(1+r)^t} c_t + \sum_{t=0}^{\infty} p_g \cdot g_t,$$

using the price of general consumption as the numeraire and denoting the interest rate by r .⁵ Note that the interest rate does not apply to future period windfall consumption, because in equilibrium the agent purchases treats with the prepaid card, where interests do not accrue.

We make the standard assumptions $\beta = \frac{1}{1+r}$ and $g < r$ so that the equilibrium is well-behaved.

Equilibrium. To solve the consumption problem, we first consider the standard problem without a prepaid card, setting $G = 0$. In this case, utility maximization yields the standard result that it is optimal to equate consumption in each period:

$$c_t^* = r \cdot \frac{z}{r-g} \quad \forall t,$$

i.e. the agent consumes the annuity value of their total income stream in each period.

Group 1 participants. We now consider participants with a prepaid card expiring after $T + 1$ periods. We first discuss some parameter restrictions so that we can focus on an equilibrium in which the agent finds it optimal to spend the entire balance available on the prepaid card, G , on windfall consumption across the $T + 1$ periods, and nothing on regular consumption. This equilibrium is sustained if, in each period until T , the marginal utility of spending on windfall consumption – net of search costs and scaled by the price of windfall consumption – is above the marginal utility of regular consumption in that period, equal to $u'(c_t^*) = 1/c_t^*$. Algebra yields that this is satisfied if

$$\frac{\lambda - \kappa s_0^*}{p_G} > \frac{r-g}{r \cdot z}, \quad (\text{A1})$$

⁴In equilibrium, the endowment is depleted in the initial period as we discuss below, i.e. $e_{it} = 0$ for all i and $t \geq 1$.

⁵Note that, in principle, agents in Groups 2 and 3 could lose some of the prepaid card balance G due to the expiration date. However, in our model, by utility maximization agents never lose money and always spend it either on windfall consumption or regular consumption.

where s_0^* is defined below in terms of exogenous parameters. We assume that this condition holds, which is guaranteed when λ is large enough.

Next, we consider an interior solution for search costs, i.e. the agent will decide to spend the entire prepaid card balance on windfall consumption but will not do it at once in order to smooth the search costs across all $T + 1$ periods. For Group 1 participants endowed with $e_0 > 0$, it is optimal to buy at least g_0 units of windfall consumption at no search cost: it would be suboptimal to wait until later periods to spend the endowment, since later periods are discounted at rate β and the prepaid card yields no interest rate. In addition, the agent exerts some search effort to buy additional windfall consumption. Considering an interior solution for search effort and windfall consumption spending in all periods up to T , the first-order conditions yield:

$$\beta(\lambda - \kappa s_{t+1}) = \lambda - \kappa s_t \quad \forall t < T - 1.$$

Since the agents purchase windfall consumption with the prepaid card only, we have the budget constraint:

$$\frac{G}{p_G} - e_0 = \sum_{t=0}^T s_t$$

From this we obtain:

$$\begin{aligned} s_0^* &= \left(\frac{G}{p_G} - e_0 \right) \frac{1 - 1/\beta}{1 - 1/\beta^{T+1}} + \nu \\ s_{t+1}^* &= \frac{1}{\beta} s_t^* - \frac{\lambda(1-\beta)}{\kappa\beta} \quad \forall t \in [1, T] \end{aligned}$$

with $\nu = \frac{\lambda(1-\beta)}{\kappa\beta} T \frac{1-1/\beta^T}{1-1/\beta^{T+1}} + \beta \frac{(1-1/\beta)(T/\beta^{T+1}) - 1/\beta^2(1-1/\beta^T)}{(1-1/\beta^{T+1})(1-1/\beta)}$.

This yields the optimal allocations:

$$\begin{aligned} g_0^* &= e_{i0} + s_0^*, \\ g_t^* &= s_t^* \quad \forall t \in [1, T], \end{aligned}$$

Note that $g_0^* > g_t^* \quad \forall t \in [1, T]$, especially for households endowed with $e_0 > 0$. This establishes our first key result: for Group 1 participant, the extra spending is concentrated in the short run. Intuitively, households who remember “windfall purchase opportunities” buy them immediately, at no search cost. They then smooth the search costs over time.⁶

Group 2 participants. For Group 2 participants, the problem is the same as above except that $\lambda_t = 0 \quad \forall t > 0$. The agent now exerts optimal search effort s_0^* in period 0 to take advantage of the fact that the marginal utility of spending on windfall goods and services is larger, through λ , in this period alone. The agent thus buys $e_{i0} + s_0^*$ windfall consumption goods and services at price p_g and spends the remainder on regular consumption, with perfect consumption smoothing over time (i.e., consuming

⁶Note that in this tractable version of the model, with parameter restrictions such that the card is entirely spent on windfall consumption, the marginal propensity to consume out of the prepaid card is 100 % over T periods. To be in line with our empirical findings of a modest MPC concentrated in the short run for Group 1 participants, we can set $T \rightarrow \infty$ to obtain a small cumulative MPC over time, with a burst of spending in the initial period.

$\frac{r}{1+r} \cdot [G - p_g (e_{i0} + s_0^*)]$ every period). We assume that the optimum satisfies an interior solution, i.e. the agents exerts search effort up to the point where the marginal utility of getting more windfall consumption equates the marginal utility of spending on regular consumption in the initial period:

$$\begin{aligned} \frac{\lambda - \kappa s_0^*}{p_g} &= u'(c_0^*) \\ &= \frac{1}{r \cdot \frac{z}{r-g} + \frac{r}{1+r} \cdot [G - p_g (e_{i0} + s_0^*)]} \end{aligned} \quad (\text{A2})$$

This characterizes the optimal choice of search effort s_0 , and thus of windfall purchases g_0 ; optimal choices can be found by solving the quadratic formula: $A + Bs_0^* + Cs_0^{*2} = 0$.⁷

To compare the consumption response of Group 2 to Group 1, note that when G is small relative to lifetime income $\frac{z}{r-g}$, as in the data, the right-hand side of equation (A2) remains essentially unchanged regardless of the choice of s_0^* . Group 2 now equates the marginal utility of spending on windfall consumption (net of search costs) to the marginal utility of regular consumption, while in the case of Group 1 the marginal utility of spending on windfall goods and services remains larger, per equation (A1). Indeed, Group 1 agents are able to smooth the search costs over many periods. Instead, Group 2 agents search more and buy more windfall consumption in the initial period. We thus obtain our second key result: Group 2 participants have a larger increase in spending than Group 1 participants in the short run.

Group 3 participants. For Group 3 participants, the optimal allocation is also given by the equation (A2), but with the higher value of κ that characterizes Group 3. We can directly infer from equation (A2) that the equilibrium levels of search and spending on windfall goods and services fall with higher search costs κ (again noting that G is small relative to lifetime income $\frac{z}{r-g}$ on the right-hand side). Per the comparison of equation (A2) to equation (A1), the spending of Group 3 remains larger than the spending of Group 1 in the initial period. This establishes our third key result: the extra spending of Group 3 falls between that of Group 1 and Group 2.

Additional prediction. The model above highlights that Group 1 participants who spend early on after receiving the prepaid card should have a large MPC. In the model, these agents are endowed with $e_0 > 0$ and are able to buy windfall consumption immediately at no search cost, while other agents smooth search costs over time and experience no spending burst upon receiving the card. Taking this prediction to the data, we analyze the sub-sample of Group 1 participants who spent the prepaid card within the first three weeks. Consistent with the prediction, we estimate a large MPC in this sub-sample of Group 1 participants: their MPC is close to that of Group 2 participants, as reported in Appendix Figure A36. This finding provides additional evidence about the channel whereby the expiry date can act as a spur to make purchases for Group 2 participants, which in our model requires incurring higher search costs that Group 1 participants prefer to avoid.

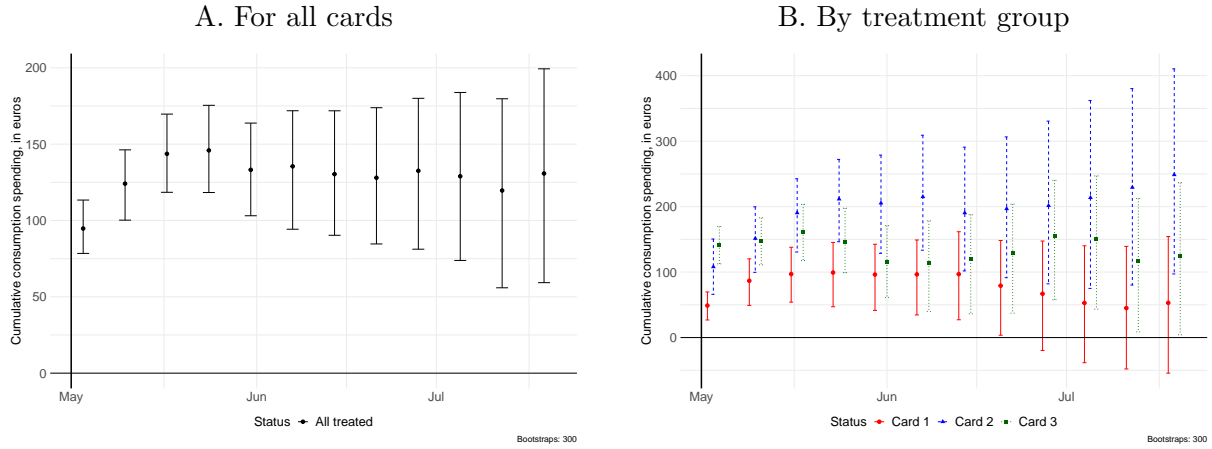
⁷The parameters are as follows:

$$\begin{aligned} A &= \frac{\lambda}{p_g} \left[r \cdot \frac{z}{r-g} + \frac{r}{1+r} (G - p_g e_{i,0}) \right] - 1, \\ B &= -\lambda \frac{r}{1+r} - \frac{\kappa}{p_g} \left[r \cdot \frac{z}{r-g} + \frac{r}{1+r} (G - p_g e_0) \right], \\ C &= \frac{r}{1+r} \kappa. \end{aligned}$$

Extension: time-varying salience. Note that the salience effects above are only tied to the prepaid card. An alternative modeling approach could assume that λ_t falls over time, i.e. the reference point for salience is not just the card but also the time of receipt. This assumption would also yield a spending response concentrated in the short run, without the need for the assumption that agents have a “search endowment”.


D Additional Figures and Tables

Figure A1 FGLS estimates of the MPC, by treatment group




Notes: This figure shows the estimated MPC using a feasible generalized least square (FGLS) procedure, where standard errors of each household's error term are parameterized to be able to vary with each bin of time-invariant characteristics calculated from pre-period data (10 age bins, 10 income bins, gender dummy, 10 liquid wealth bins, 10 average consumption expenditure bins, 95 local area dummies), i.e. in each iteration we calculate weights from $1/\hat{\sigma}_i^2$, where $\hat{\sigma}_i$ is the predicted standard error from a regression of the household-level standard error in the previous iteration on characteristic bin dummies. While Panel A considers all cards, Panel B presents the estimates by treatment group. The 95% confidence intervals are obtained by bootstrap.

Figure A2 Long-term MPC estimates



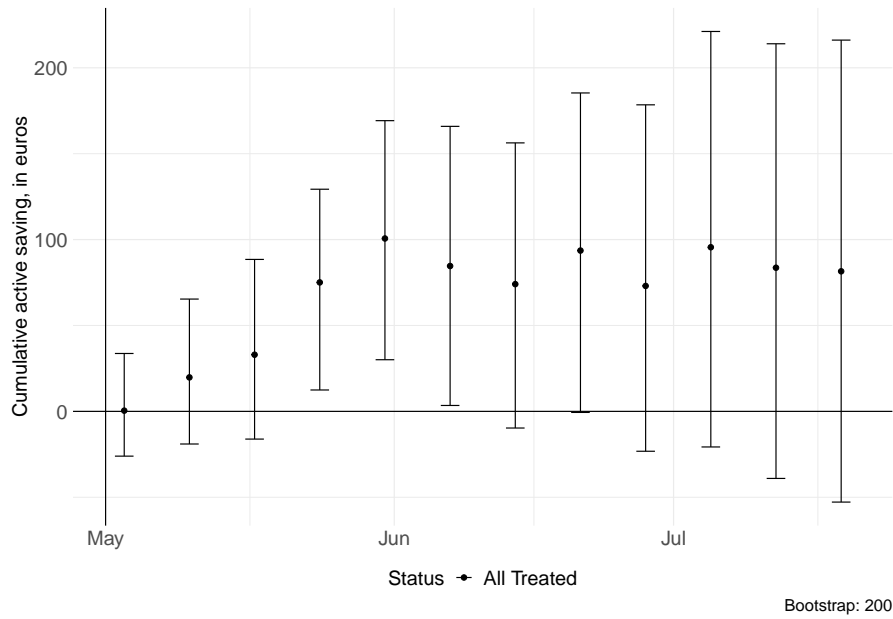
Notes: In this figure, we run a specification analogous to (1), except that to reduce noise we use a seventh-order polynomial to model the weekly MPC response after treatment: $Y_{it} = \sum_{k=1}^8 \beta_{\tau}^{k-1} \cdot \tau_{it}^{k-1} + \alpha_i + \alpha_{tE} + \varepsilon_{it}$. We still view this specification as non-parametric estimation of the MPC, given the flexibility of the 7th-order polynomial. To reduce noise further, we use the same feasible generalized least square (FGLS) procedure as in Figure A1. The figure reports the cumulative MPC and both the 95% and 68% confidence intervals, clustered at the household level.

Figure A3 The Response of Consumption and Savings




Notes: In this figure, we run a specification analogous to (1), except that the outcome is the sum of consumption flows and savings into the checking account and liquid or illiquid savings accounts at the bank. Furthermore, to reduce noise we use a seventh-order polynomial to model the weekly outcome response after treatment: $Y_{it} = \sum_{k=1}^8 \beta_{\tau}^{k-1} \cdot \tau_{it}^{k-1} + \alpha_i + \alpha_{tE} + \varepsilon_{it}$, which we estimate with the same FGLS procedure as in Figure A1. The figure reports the cumulative change in the outcome and both the 95% and 68% confidence intervals, clustered at the household level.

Figure A4 Savings into liquid savings accounts




Notes: This figure analyzes the response of savings into liquid savings accounts at the bank (e.g., “Livret A”) for the treated participants. The figure reports the cumulative net flows of savings after the start of the experiment.

Figure A5 Total Spending Response, Weekly, with Treatment Leads




Notes: This figure shows the results of a regression estimating a specification analogous to equation 1, but including pre-treatment leads.

Figure A6 Observable Predictors of Non-Take-Up



Notes: This figure reports the predictors of non-take-up of the treatment card, using the full sample of treated households. We use a range of characteristics to predict whether the household never use the prepaid card during the six months following the experiment. We find that households who do not use the treatment card tend to be younger, with fewer children, higher liquid wealth, and lower illiquid wealth.


Figure A7 Cumulative MPC Estimates with Broader Consumption Measure




Notes: This figure reports the results from specification (1), with a broader measure of consumption including all direct debit transactions as well as wire transfers.

Figure A8 MPC Estimates for Reweighted Sample

A. Weekly MPC

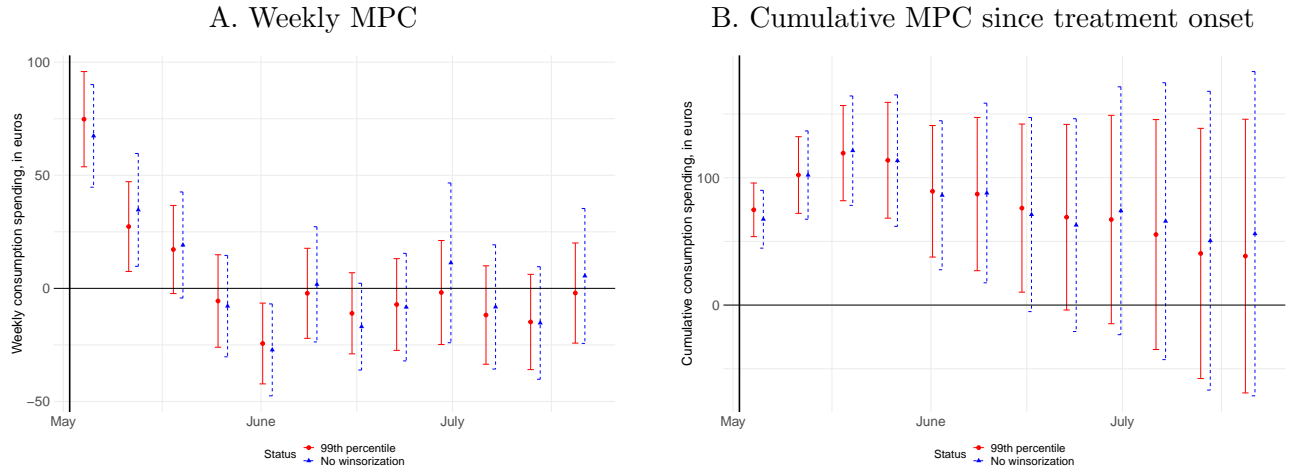


B. Cumulative MPC since treatment onset



Notes: This figure reports the results from specification (1) in a weighted OLS regression such that our sample is representative of the French population by local area (“”), income decile, age, and gender. The population weights are obtained from the French statistical institute INSEE.

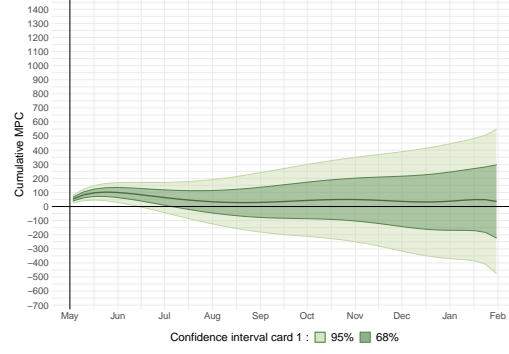
Figure A9 MPC Estimates without Winsorization




Notes: This figure reports the results from specification (1), comparing the results in our main sample (winsorizing consumption expenditures at the 99th percentile) and without winsorization.

Figure A10 Long-term MPC Estimates by Card Type

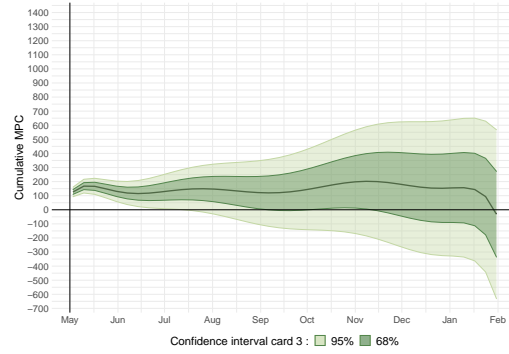
A. Group 1, no restrictions on treatment card



B. Group 2, expiration after three weeks




C. Group 3, negative rates every week



Notes: This figure plots the cumulative MPC by card type over a long horizon. To reduce noise we use a seventh-order polynomial to model the weekly outcome response after treatment: $Y_{it} = \sum_{k=1}^8 \beta_{\tau}^{k-1} \cdot \tau_{it}^{k-1} + \alpha_i + \alpha_{tE} + \varepsilon_{it}$, which we estimate with the same FGLS procedure as in Figure A1. The figure reports the cumulative change in the outcome and both the 95% and 68% confidence intervals, clustered at the household level.


Figure A11 Cumulative Spending on the Prepaid Card by Card Type




Notes: This figure reports average cumulative spending on the prepaid card for the three types of card.

Figure A12 Characteristics of Households with Low Take-Up by Card Type


A. Group 1, no restrictions on treatment card



B. Group 2, expiration after three weeks




C. Group 3, negative rates every week



Notes: This figures analyzes the socio-demographic characteristics of Group 1 participants who did not use the card by July 1st (panel A), and those of Group 2 and 3 participants who lost at least 50 Euros (panels B and C).


Figure A13 4-week MPC Differences between Card Types




Notes: This figure reports the 4-week MPC difference between card types, with 95% confidence intervals.

Figure A14 Long-term MPC Estimates for Durables and Nondurables

A. Non durable MPC




B. Durable MPC




Notes: This figure reports the cumulative MPC estimates for nondurables and durables, considering all cards together. To reduce noise we use a seventh-order polynomial to model the weekly outcome response after treatment: $Y_{it} = \sum_{k=1}^8 \beta_{\tau}^{k-1} \cdot \tau_{it}^{k-1} + \alpha_i + \alpha_{tE} + \varepsilon_{it}$, which we estimate with the same FGLS procedure as in Figure A1. The figure reports the cumulative change in the outcome and the 95% confidence intervals, clustered at the household level.

Figure A15 MPC by Spending Category


A. MPC on durables




B. MPC on semi-durables



C. MPC on non-durables

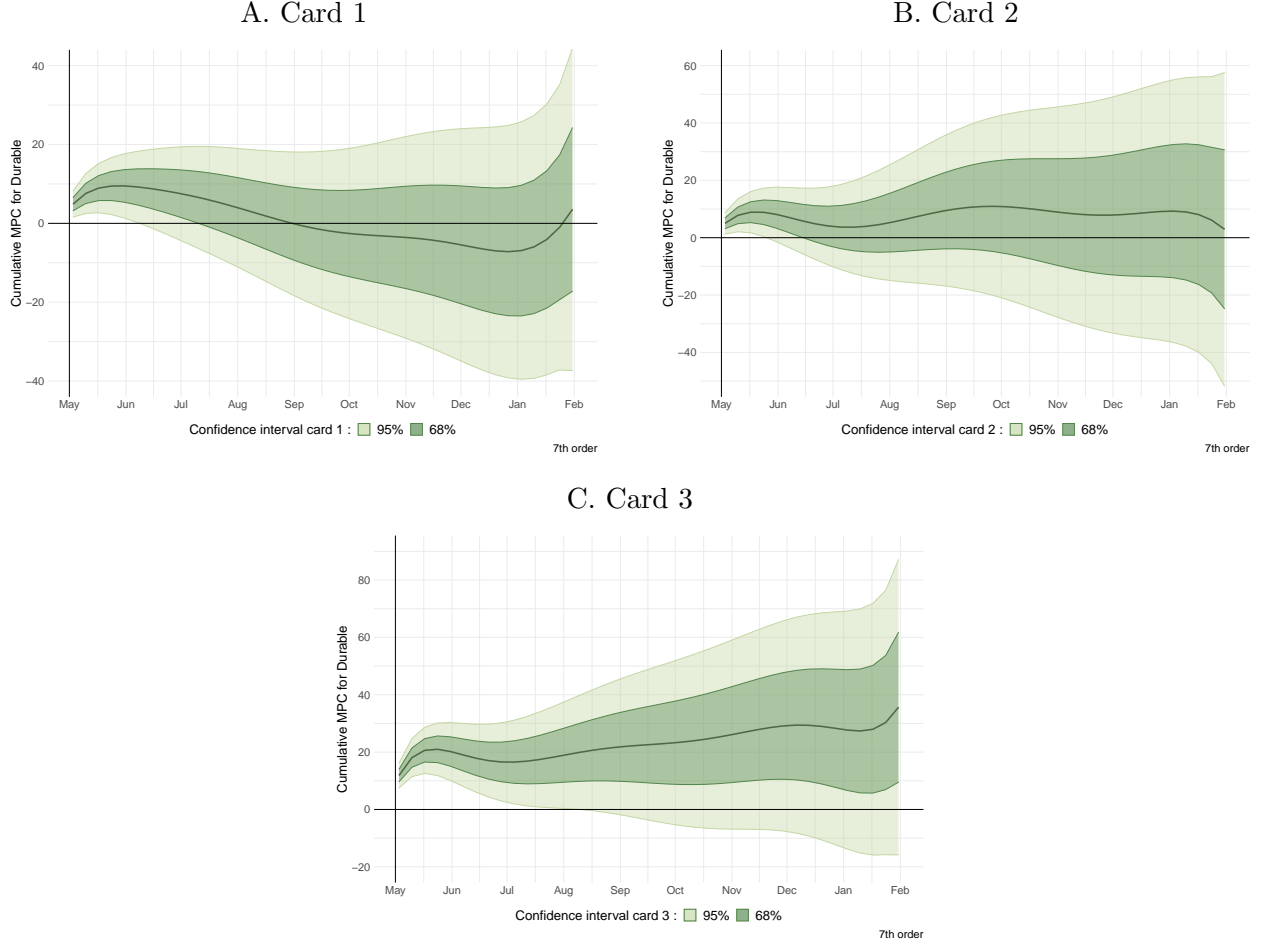


D. MPC on services




Notes: This figure reports MPCs by spending category and card type.

Figure A16 Long-term MPC Estimates for Durables by Card Type




Notes: This figure reports the cumulative MPC estimates for spending on durables, considering all cards together. To reduce noise we use a seventh-order polynomial to model the weekly outcome response after treatment: $Y_{it} = \sum_{k=1}^8 \beta_{\tau}^{k-1} \cdot \tau_{it}^{k-1} + \alpha_i + \alpha_{tE} + \varepsilon_{it}$, which we estimate with the same FGLS procedure as in Figure A1. We estimate the specification separately for each card type. The figure reports the cumulative change in the outcome and the 95% confidence intervals, clustered at the household level.

Figure A17 Spending on Imports, by Framing Group



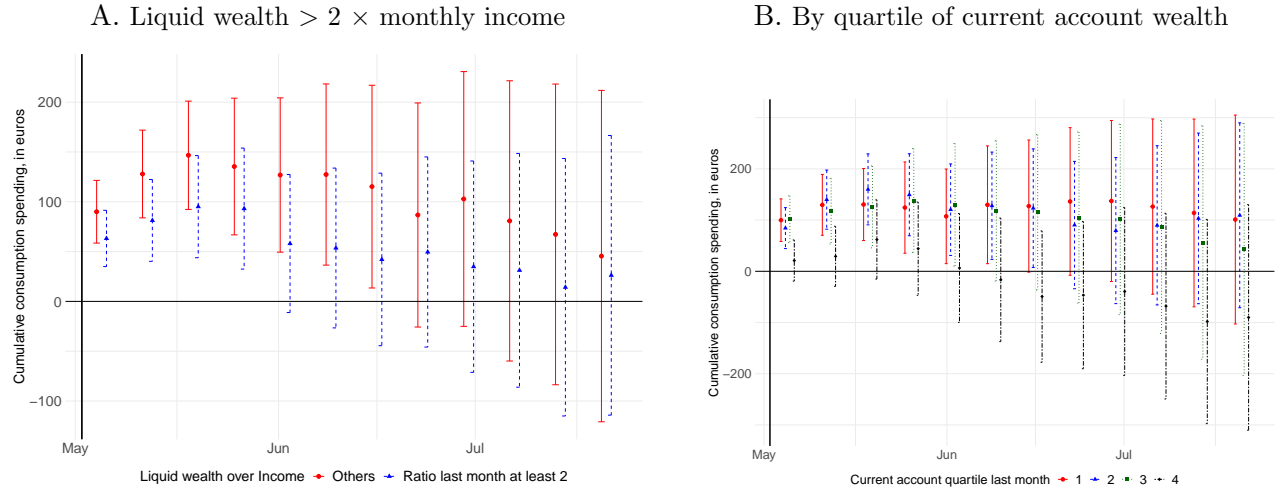
Notes: This figure shows the average expenditure share on imports for households in the two framing groups and the control group. Imports are calculated as the fraction of each product category that is directly imported from abroad, using the Input-Output table provided by the French statistical institute INSEE, and linked to MCC codes using our crosswalk.

Figure A18 Stated Propensity to Purchase “Treats” by Framing Group




Notes: This figure reports the fraction of respondents who answer that they used the funds on the prepaid cards primarily to purchase a “treat”, for the two framing groups. 95% confidence intervals are reported.

Figure A19 Total spending response, weekly, for households with high liquid wealth



Notes: The panels of this figure shows the results of estimating equation 1 in a subsample of households whose liquid wealth is larger than twice their monthly income, and by quartiles of current account wealth. The figure plots the estimates for the cumulative MPC at different time horizons.

Figure A20 Heterogeneity by Gender and Marital Status




Notes: This figure shows the cumulative MPC by gender and marital status.

Figure A21 Spending Composition by Gender

Panel A: Prepaid card alone, spending shares




Panel B: Bank account spending and prepaid card, differences relative to the control group



Notes: Panel A shows the spending shares by gender for the prepaid card alone. Panel B shows differences in spending shares by gender for the treated households compared to control households when considering both the prepaid card and the household's bank account


Figure A23 LASSO estimates of treatment effect heterogeneity coefficients, group 1 and control group




Notes: This figure shows estimates of specification (2) on the set of observations pertaining to treatment group 1 and control observations, for varying levels of the regularization parameter.

Figure A24 Quantile treatment effects: de-meaned cumulative consumption, treated vs. control


A. 4-week horizon



B. 8-week horizon

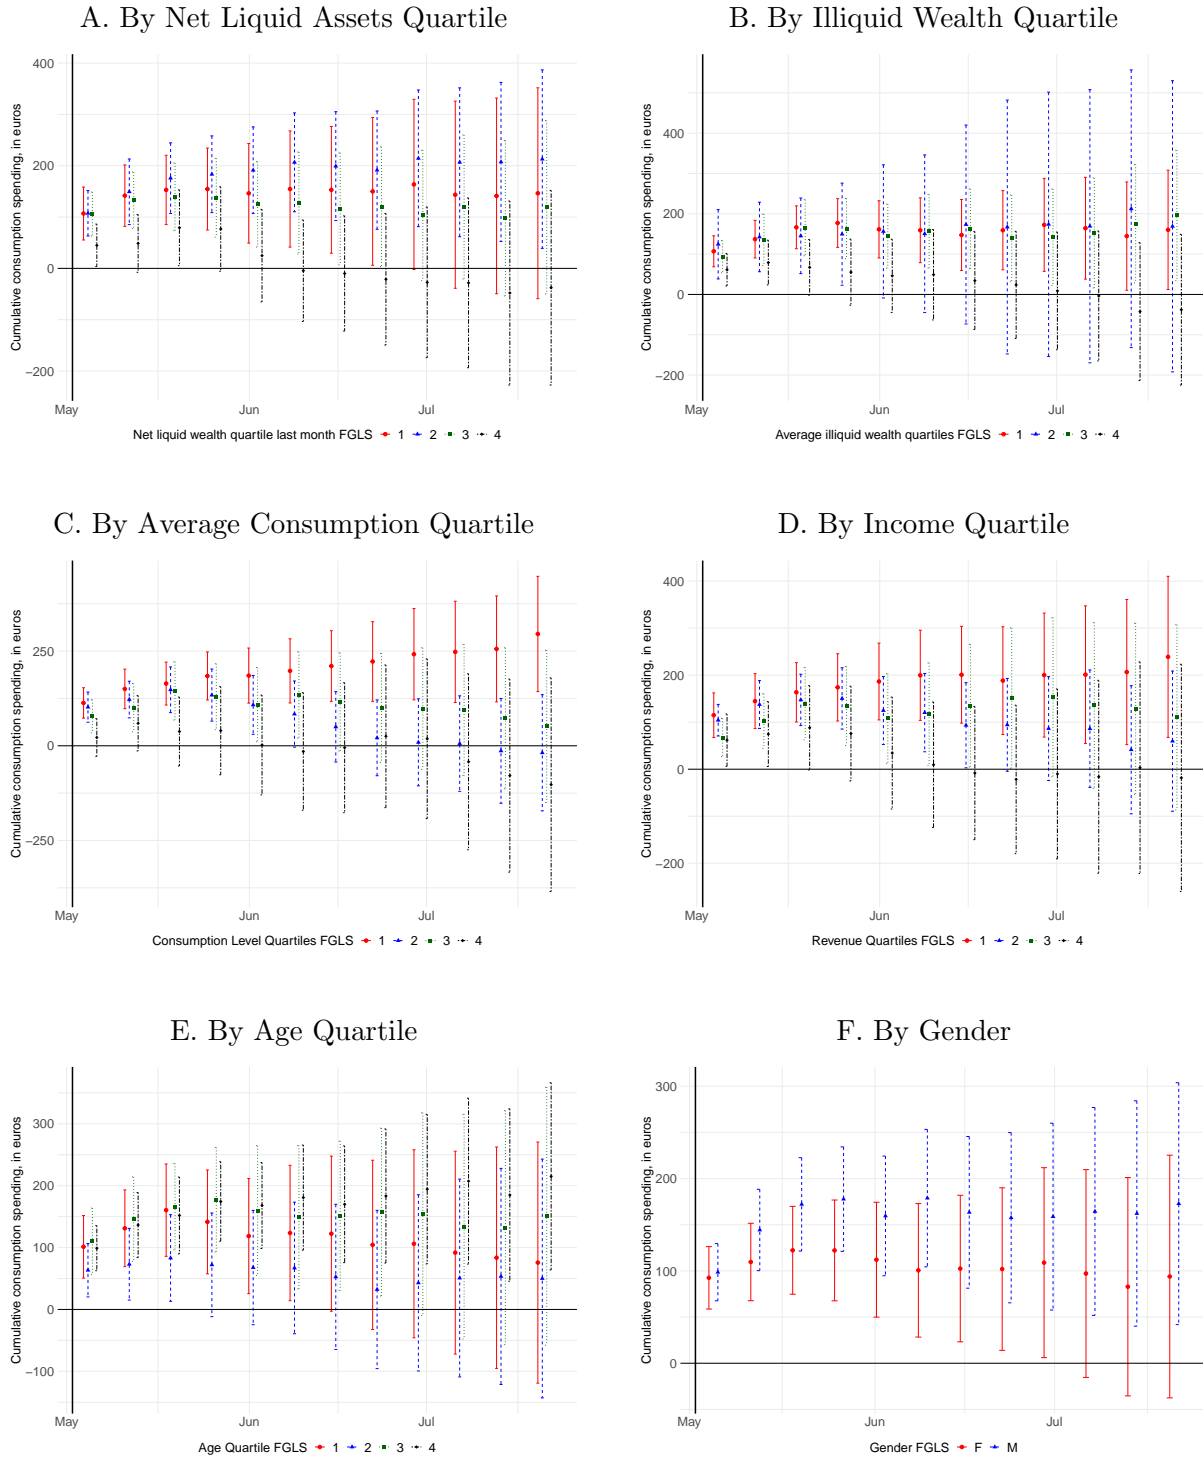


C. 12-week horizon




Notes: This figure shows quantile treatment effects—the difference between the quantiles of the distribution of treated and control groups—for cumulative de-meaned consumption expenditures. Standard errors are estimated using the bootstrap.

Figure A22 MPC Heterogeneity by Observable Household Characteristics



Notes: This figure reports MPC estimates depending on observable household characteristics using a feasible generalized least square (FGLS) procedure, where standard errors of each household's error term are parameterized to be able to vary with each bin of time-invariant characteristics calculated from pre-period data (10 age bins, 10 income bins, gender dummy, 10 liquid wealth bins, 10 average consumption expenditure bins, 95 *departement* dummies), i.e. in each iteration we calculate weights from $1/\hat{\sigma}_i^2$, where $\hat{\sigma}_i$ is the predicted standard error from a regression of the household-level standard error in the previous iteration on characteristic bin dummies. 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 and marital status. 95% confidence intervals, with standard errors clustered at the household level, are reported in all panels.

Figure A25 Quantiles of the 4-week cumulative percentage deviation from mean household consumption expenditure




Notes: This figure reports the quantiles of the distribution of 4-week treatment effects by treatment group. In contrast to the model from the baseline specification, the log of weekly consumption expenditure is here linear in treatment effects and fixed effects. Specifically, we estimate a specification in log weekly average consumption:

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

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


We estimate the distribution of treatment effects in this model using the same deconvolution approach as in the main text. Since the dependent variable, log consumption expenditure, is much less skewed than in the baseline model from Section 4.2, we winsorize it only at the 99th percentile. Note that here the treatment effect is not a marginal propensity to consume, but a cumulative percentage deviation from the household's mean consumption expenditure level (on average 417 euros). The estimates from this model show similar economic effects to the benchmark specification. Shaded regions are delineated by the 10th and 90th percent quantile of the bootstrapped simulated distribution of the corresponding moment.

Figure A26 Estimated distribution of MPCs without constraining it to have no mass on negative values




Notes: The figure shows the estimated distribution of 4-week MPCs using the flexible deconvolution procedure of [Yang et al. \(2020\)](#) when the support of the density of the distribution is not constrained to lie on the positive part of the real line.

Figure A27 MPC Distribution, Group 1 vs Groups 2 and 3 combined




Notes: The figure shows the quantiles of the estimated distribution of MPCs, when the estimation is performed separately for treatment group 1 and for treatment groups 2 and 3 (jointly). Standard errors are estimated using a bootstrap with 150 draws.

Figure A28 Distribution of cumulative spending on treatment card, across households



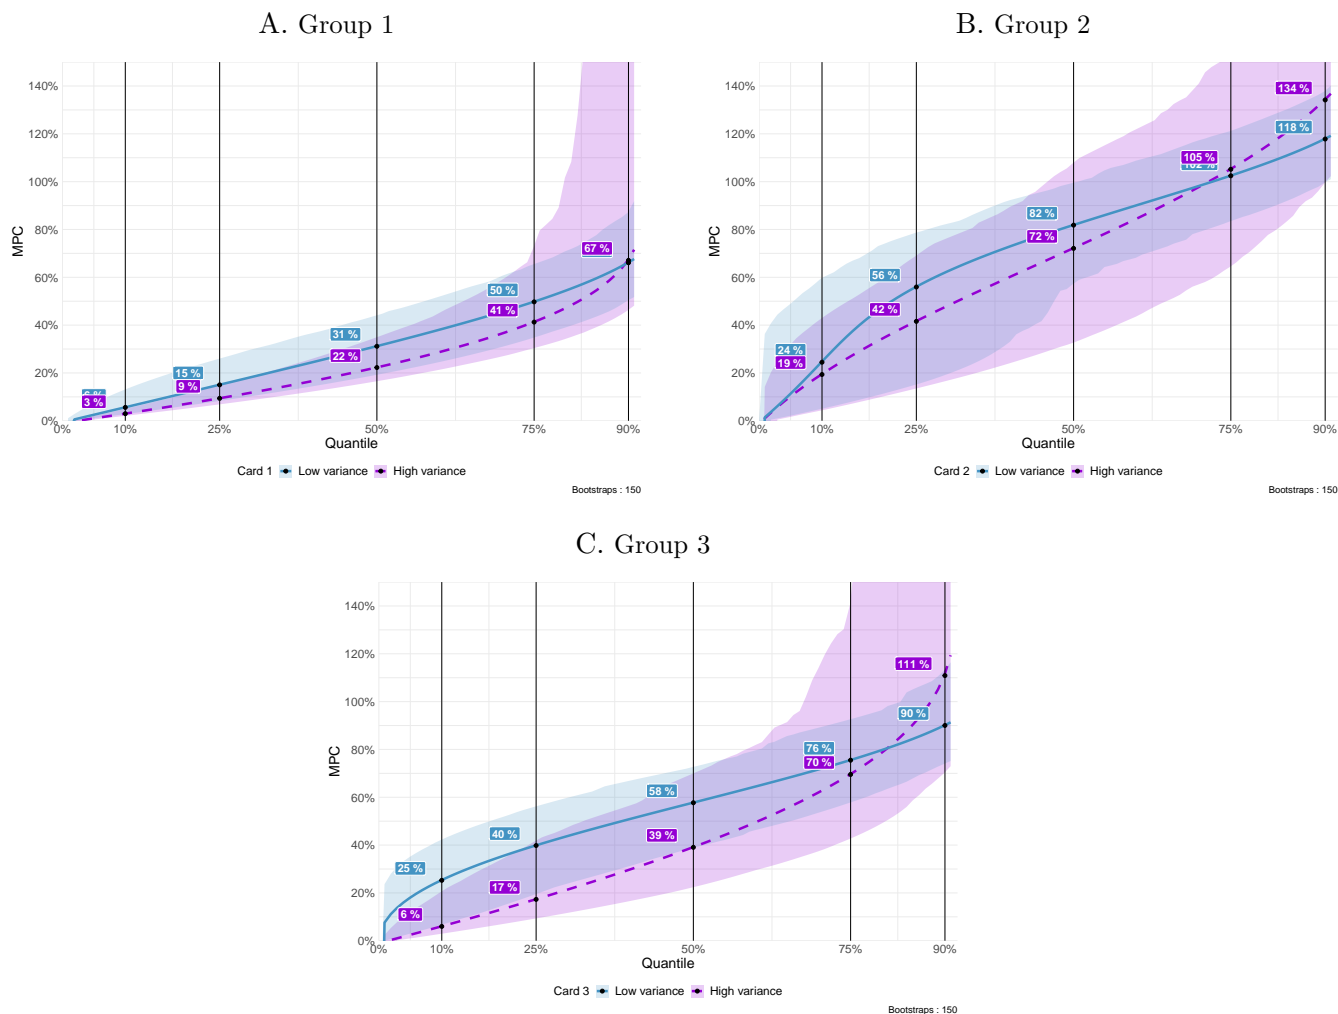
Notes: The figure shows moments of the distribution of cumulative expenditures on the treatment card, for each week.

Figure A29 Household-level Quantiles of the 4-week MPC Distribution: Robustness to inclusion of observed characteristics bins




Notes: This figure reports the quantiles of the distribution of 4-week treatment effects by treatment group in a two-stage procedure with controls. In the first stage of our estimation procedure we project weekly consumption expenditures on household 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. In the second state we deconvolve, as before, the outcome distribution for the treated by the empirical distribution of the error term of the control group. Constraints on the estimated distribution (mass on positive part of real line, and penalization) are as in the specification of the benchmark estimate. The figure shows the resulting treatment distribution estimates, which are similar to those reported using the baseline procedure. Shaded regions are delineated by the 10th and 90th percent quantile of the bootstrapped simulated distribution of the corresponding moment.

Figure A30 Unconditional MPC distribution estimates, high-variance vs low-variance households




Notes: The panels of this figure show estimates of the 4-week MPC distribution for each treatment group, where the sample is split into households that have a below-median variance of pre-period weekly consumption expenditures (green) and above-median variance of pre-period weekly consumption expenditures (purple). For each treatment card, the estimated distributions for both variance groups are very similar (indeed, no quantile is significantly different from each other), giving support to the view that treatment effect distributions are similar even for households that have different higher moments of consumption. Shaded regions are delineated by the 10th and 90th percent quantile of the bootstrapped simulated distribution of the corresponding moment.

Figure A31 Comparison of Dynamic MPC Estimates to Standard HANK Model




Notes: This figures compare the dynamics of the spending response in our experiment and in the calibrated HANK model of [Kaplan et al. \(2018\)](#). Panel A draws this comparison using Group 1 participants. The red line reports our baseline weekly MPC estimates for Card 1. The green line reports the estimates obtained by fitting an exponential model such that the MPC t weeks after treatment is given by $C_0 \cdot e^{-a \cdot t}$. Finally, the blue line reports the path of MPCs according to the calibrated HANK model of [Kaplan et al. \(2018\)](#); while their model is quarterly, we fit by nonlinear least squares an exponential model of weekly MPCs to match the calibrated quarterly MPCs reported in Figure 2 of [Kaplan et al. \(2018\)](#), which gives a cumulative MPC out of a \$300 transfer of 17%, 25%, and 36% at 13, 26 and 52 weeks respectively. Panel A shows that the estimated decay parameter a is one order of magnitude larger in our data than in HANK. We reject that the baseline weekly MPC model with Card 1 is the same as the MPC path from HANK with a p-value of 0.038. In contrast, when comparing our baseline weekly MPC estimates to our estimated exponential model, we cannot reject the exponential model (p-value = 0.63). Panel B repeats the analysis using all treatment groups, rather than Group 1 alone. We estimate an even larger decay parameter and reject the HANK model with a p-value close to zero (p-value = $5 \cdot 10^{-11}$). In contrast, we cannot reject our estimated exponential model (p-value = 0.38).

Figure A32 Summary of MPCs estimates



Notes: This figure reports the estimates of MPCs in the literature (typically from their baseline specifications), coded by time horizon and expenditure categories; 95% confidence intervals are also reported.

Figure A33 Fraction of households that should have used the treatment card but did not: narrow sample




Notes: This figure shows the same result as Figure 12 – the fraction of households that should have used the treatment card but did not – but only among the population of households that consist of a single adult and no children, and conditioning on making consumption expenditures on that day in physical stores only (i.e. not online or via phone). This is to rule out the possibility that this phenomenon is driven by multi-person households of whom only one has access to the treatment card, or by households that are not aware that they could use the card online.


Figure 1 is a line graph showing the consumption increase per euro of stimulus (Y-axis, ranging from -50% to 100%) over time (X-axis, May to Jul) for two groups: 2 (blue circles) and 3 (green triangles). The graph displays the mean consumption increase and its associated error bars for each group across several time points. Group 2 consistently shows a higher consumption increase than Group 3, and both groups show an increasing trend over time.

Time Point	Group 2 (Mean %)	Group 2 (Error %)	Group 3 (Mean %)	Group 3 (Error %)
May 1	30	±15	52	±12
May 2	48	±20	45	±15
May 3	68	±25	50	±20
May 4	72	±35	42	±25
May 5	65	±40	28	±30
May 6	68	±45	22	±35
May 7	60	±50	20	±40
May 8	62	±55	25	±45
May 9	60	±60	40	±50
May 10	60	±65	32	±55
May 11	60	±70	22	±60
May 12	58	±75	22	±65

Figure A35 Yield Curve for French bonds, May 2022



A41

Figure A36 Testing an Auxiliary Prediction of the Simple Model

Notes: Motivated by the auxiliary prediction of the simple model in Appendix C, this figure reports the cumulative MPC for three groups of participants: all participants in Group 2, participants in Group 1 who used the prepaid card to make purchases of a total value of more than 275 euros in the first three weeks of the experiment, and other Group 1 participants who did not spend down the prepaid card as quickly.

Table A1 Summary statistics, weekly consumption spending

	count	mean	std	10%	25%	50%	75%	90%
Weekly cons. expend. (cash and cards), total	2,571,000	417.66	435.02	67.30	163.25	315.95	542.63	848.64
Direct debits, debt payments, Subscriptions	2,571,000	143.83	1670.08	0.00	0	48.75	159.45	325.55
Wire transfers	2,571,000	111.50	845.94	0	0	0	0	100
Weekly cons. expend. (broad measure), excl. treatment cards	2,571,000	672.90	1947.09	120.28	246.67	421.70	704.66	1,000,670.55
Weekly cons. expend. (cash and cards), excl. treatment cards	2,571,000	417.57	434.99	67.25	163.19	315.88	542.53	848.50
Weekly cash withdrawals	2,571,000	23.74	83.71	0.00	0.00	0.00	0.00	70.00

Notes: The table shows summary statistics on different consumption categories by week. The sample consists of all household-weeks since January 2022. The broad measure of consumption includes the total of cash withdrawals, card spending, automatic debits, and wire transfers.

Table A2 Summary statistics, household characteristics

	count	mean	std	10%	25%	50%	75%	90%
Age of eligible household member	85,700.00	47.03	12.92	30.00	36.00	46.00	58.00	65.00
Number of eligible household members	85,700.00	1.15	0.36	1.00	1.00	1.00	1.00	2.00
Avg. monthly incoming transfers, 6m prior	85,685.00	2,654.04	1,439.56	1,317.69	1,796.14	2,381.15	3,159.71	4,217.33
Avg. monthly salaries, social allowance, pensions, benefits, 6m prior	80,034.00	2,109.55	4,968.86	493.06	1,049.91	1,667.95	2,348.71	3,259.71
Avg. monthly incoming salaries, 6m prior	80,034.00	1,630.59	5,003.78	0.00	95.65	1,171.51	2,077.47	3,053.53
Avg. monthly incoming pension payments, 6m prior	80,034.00	300.72	691.85	0.00	0.00	0.00	0.00	1,464.42
Avg. monthly incoming social allowances, 6m prior	80,034.00	98.06	199.81	0.00	0.00	0.00	110.932	298.32
Avg. monthly incoming unemployment benefits, 6m prior	80,034.00	80.18	284.12	0.00	0.00	0.00	0.00	242.28
Dummy: has received unemployment benefits within 6m prior	85,685.00	0.14	0.35	0.00	0.00	0.00	0.00	1.00
Avg. net liquid wealth, 1m prior	85,700.00	19,265.55	46,067.40	121.38	1,506.66	7,599.54	23,854.85	55,256.41
Avg. current account balance, 6m prior	85,698.00	4,448.41	19,976.04	63.50	424.17	1,006.25	2,487.53	7,563.32
Avg. liquid savings, 6m prior	85,698.00	16,896.51	34,466.19	17.30	617.74	5,462.02	19,253.96	44,896.89
Avg. value of life insurance assets, 6m prior	85,698.00	5,867.43	32,466.19	0.00	0.00	0.00	372.91	8,924.360
Avg. net illiquid wealth , 6m prior	85,700.00	64,746.97	185,159.99	-43,194.32	0	1,285.04	93,451.46	211,111.53
Avg. total debt, 6m prior	85,698.00	-33,300.8	55,007.44	-99,130.28	-52,931.94	-5,641.02	0.00	0.00
Avg. consumer debt, 6m prior	85,698.00	-2,388.27	5,194.01	-7,590.37	-2,979.42	0.00	0.00	0.00
Avg. mortgage debt, 6m prior	85,698.00	-30,872.90	54,288.18	-96,402.39	-50,413.24	0.00	0.00	0.00
Number of adult members in the household	85,698.00	1.53	0.50	1.00	1.00	2.00	2.00	2.00
Number of children in the household	85,698.00	0.61	0.96	0.00	0.00	0.00	1.00	2.00
Avg. monthly consumption expenditures (cash, card payments), 6m prior	85,690.00	1,205.63	658.22	545.96	794.64	1,102.89	1,480.69	1,940.18
Avg. monthly outgoing transfers, 6m prior	85,635.00	518.12	1,001.24	148.95	233.58	364.50	587.08	912.76
Avg. total monthly consumption (broad measure)	85,631.00	1,715.60	1,563.41	960.39	1,304.68	1,747.74	2,370.86	32,60.95

Notes: This table report the distributions of the household characteristics used in our analysis. The variable “Avg. monthly outgoing transfers” includes direct debits, debt payments, and subscriptions. The variable “Avg. total monthly consumption (broad measure)” includes the sum of cash withdrawals, card spending, automatic debits, and wire transfers.

Table A3 Examples of MCCs Classified across Product Categories

Description of MCC Product Category	Product Type
Veterinary Services	S
Agricultural Co-operatives	S
Horticultural Services, Landscaping Services	S
General Contractors-Residential and Commercial	S
Air Conditioning Contractors , Sales and Installation, etc.	S
Electrical Contractors	S
Insulation , Contractors, Masonry, Stonework Contractors, etc.	S
Carpentry Contractors	S
Roofing , Contractors, Sheet Metal Work, etc.	S
Motor vehicle supplies and new parts	D
Office and Commercial Furniture	D
Construction Materials, Not Elsewhere Classified	D
Office, Photographic, Photocopy, and Microfilm Equipment	D
Computers, Computer Peripheral Equipment, Software	D
Men's Women's and Children's Uniforms and Commercial Clothing	SD
Commercial Footwear	SD
Home Supply Warehouse Stores	SD
Variety Stores	SD
Misc. General Merchandise	SD
Grocery Stores, Supermarkets	ND
Meat Provisioners , Freezer and Locker	ND
Candy, Nut, and Confectionery Stores	ND
Dairy Products Stores	ND
Bakeries	ND
Misc. Food Stores , Convenience Stores and Specialty Markets	ND

Notes: This table illustrates the classification of product categories, defined by their Merchant Category Code (MCC), into four groups: services (S), durables (D), semi-durables (SD), and nondurables (ND). This table only focuses on a subset of products, out of the total of 933 MCC categories in our data.