

Bernard, René

Conference Paper

Mental Accounting and the Marginal Propensity to Consume

Beiträge zur Jahrestagung des Vereins für Socialpolitik 2022: Big Data in Economics

Provided in Cooperation with:

Verein für Socialpolitik / German Economic Association

Suggested Citation: Bernard, René (2022) : Mental Accounting and the Marginal Propensity to Consume, Beiträge zur Jahrestagung des Vereins für Socialpolitik 2022: Big Data in Economics, ZBW - Leibniz Information Centre for Economics, Kiel, Hamburg

This Version is available at:

<https://hdl.handle.net/10419/264186>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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

Mental Accounting and the Marginal Propensity to Consume*

René Bernard

Deutsche Bundesbank and Goethe University

This Draft: August 29, 2022

Abstract

This paper studies how consumers respond to unexpected, transitory income shocks and why. In a randomized control trial, I elicit marginal propensities to consume (MPC) out of different hypothetical income shock scenarios, varying the payment mode, the shock size, and the source of income. The results show respondents exhibit a higher MPC when exposed to a windfall paid out in cash or without any specification of the payment mode, respectively, compared to a windfall deposited in an instant-access savings account, suggesting consumers violate fungibility. Further, the MPC falls with the shock size, whereas it does not vary with the source of income. Using causal machine learning methods to explore treatment heterogeneity, I find that low liquidity, self-control problems, and a lack of cognitive sophistication contribute to MPC heterogeneity. The results are broadly in line with mental accounting theory.

Keywords: Randomized control trial, marginal propensity to consume, fiscal policy, mental accounting, causal forest.

JEL classification: C90, D12, D14, D15, D91.

*Contact address: René Bernard, Wilhelm-Epstein-Strasse 14, 60431 Frankfurt am Main, Germany. Phone: +49 69 9566-8656. EMail: rene.bernard@bundesbank.de. I thank Annika Bacher (discussant), Peter Egger, Hauke Feil, Olga Goldfayn-Frank, Győző Gyöngyösi, Andreas Hackethal, Tobin Hanspal, Steve Heinke, Susanne Helmschrott, Michael Knaus, Christine Laudenbach, Viktor Marinkov (discussant), Filip Matějka, Fabian Nemecek, Gabriel Okasa, Michaela Pagel, Thomas Pauls, Jonathan Parker, Luigi Pistaferri, Tobias Schmidt, Michael Weber, Nils Wehrhöfer, Joachim Winter, Basit Zafar, the audiences of the SAFE Household Finance Workshop 2021, the Workshop on Theoretical and Experimental Macroeconomics 2021 at TU Berlin & Pompeu Fabra University, the ESWM 2021, the RGS Doctoral Conference 2022, the SGF Conference 2022, the 2022 Annual Conference of the Society for Experimental Finance, the Behavioral Macroeconomics Workshop 2022, the SEM's 2022 annual conference as well as seminar participants at the Bundesbank, Goethe University, and the ECB for valuable comments. Special thanks goes to Konstantin Bräuer, Thomas Pauls, Seyyedshahabuldin Mirhosseini, and Marcel Möbus for fielding the online survey of retail customers at a large German bank. This work represents the author's personal opinion and does not necessarily reflect the views of the Deutsche Bundesbank or the Eurosystem. This paper uses data from the Bundesbank Online Panel – Households (BOP-HH). The results published and the related observations and analysis may not correspond to results or analysis of the data producers.

1 Introduction

The COVID-19 pandemic has brought renewed attention to government payments to households aimed at boosting aggregate demand. For an optimal design of such stimulus payments, it is crucial to understand how and why consumers respond to unexpected, transitory income shocks.

Textbook models of consumption ([Deaton, 1991](#); [Carroll, 1997](#)) predict that the marginal propensity to consume (MPC) out of unexpected, transitory income shocks declines with the level of liquidity and the size of the income shock, as some households face occasionally binding liquidity constraints. Yet, there is only mixed empirical evidence for a negative relation between liquidity and the MPC.¹ Moreover, only some papers have considered how MPCs vary with the shock size and arrive at contradictory conclusions.²

In contrast, the behavioral life-cycle hypothesis ([Shefrin and Thaler, 1988](#)) states that labeling affects how income shocks are spent. According to the theory, consumers label income shocks as belonging to different non-fungible mental accounts, depending on the income shocks' salient attributes, e.g. the size or the payment mode. Some mental accounts are more tempting to invade than others. Consequently, the MPC should vary with the salient attributes of income shocks.

In this paper, I study consumption responses to different types of windfalls in order to test competing theories of consumption. I use survey questions to elicit MPCs out of hypothetical income shocks. The survey-based approach allows me to estimate causal effects of different types of income shocks on the MPC that are difficult to identify using data on revealed behavior. In a randomized control trial, I generate exogenous variation in the size of the income shock by asking half of the respondents about their consumption response to a one month's income shock and the other half

¹[Baker, Farrokhnia, Meyer, Pagel, and Yannelis \(2020\)](#), [Coibion, Gorodnichenko, and Weber \(2020\)](#), [Fagereng, Holm, and Natvik \(2021\)](#), as well as [Jappelli and Pistaferri \(2014, 2020\)](#) find a negative relation between liquid wealth and the MPC, whereas [Bunn, Le Roux, Reinold, and Surico \(2017\)](#), [Christelis, Georgarakos, Jappelli, Pistaferri, and van Rooij \(2019\)](#), and [Fuster, Kaplan, and Zafar \(2021\)](#) find no significant relationship between liquid wealth and MPCs out of positive income shocks.

²[Andreolli and Surico \(2021\)](#), [Christelis et al. \(2019\)](#), and [Fagereng et al. \(2021\)](#) find that MPCs fall with the shock size, whereas [Fuster et al. \(2021\)](#) find the opposite.

about their response to a three months' income shock. Further, I exogenously vary the payment mode. One group of respondents receives an income shock paid out in cash, a second group receives it as a payment deposited in an instant-access savings account, and a third group receives it without any specification of the payment mode. In a further experiment, I test whether the source of income affects the MPC, exposing one group of consumers to a one-off payment from the government and another group to a lottery win.

I then study heterogeneity in consumption responses to understand the underlying mechanisms. The survey-based approach is well-suited for this task as it yields a direct MPC estimate for each individual and hence reduces problems of statistical power.³ Further, it enables me to elicit a rich set of behavioral characteristics which is not normally available in studies using transaction or administrative data.

To explore heterogeneity one typically interacts a treatment indicator with a covariate of interest. One problem with this method is that an interaction can show up as significant by pure chance if one conducts hypothesis tests for multiple covariates. Further, interaction models may miss important (unexpected) heterogeneity if the true treatment heterogeneity is more complex. To overcome these problems, I use causal machine learning methods that allow me to estimate treatment heterogeneity in a completely flexible, non-parametric, data-driven way. Specifically, I use a causal forest (Athey, Tibshirani, and Wager, 2019), predicting treatment effects for each individual by searching for the most responsive subgroup over all combinations of observable characteristics. Notably, the causal forest avoids capturing spurious heterogeneity by using a *honest* sample-splitting technique.

The main findings can be summarized as follows. First, the MPC varies with the payment mode of the income shock. Respondents exposed to an income shock paid out in cash and without any specification of the payment mode, respectively, exhibit a significantly higher MPC than those exposed to an income shock deposited in an instant-access savings account. The payment mode effect is significant in economic terms as it corresponds to a relative increase in the MPC of about 14% and 16%,

³Statistical power can be a limiting factor for heterogeneity analysis in studies that use revealed preference data to estimate MPCs in natural experiments (Parker and Souleles, 2019).

respectively. A notable subgroup of respondents does not spend anything out of the windfall deposited in an instant-access savings account, indicating that the extensive margin of spending adjustment plays an important role for the payment mode effect.

Second, the causal forest identifies considerable heterogeneity in the payment mode effect across respondents. Individuals who strongly adjust their MPC in response to changes in the payment mode are more impatient and impulsive than those who hardly react. Simulating how, on average, the payment mode effect varies with a single covariate while holding confounders constant, I find that the effect on the MPC grows with self-control problems and a lack of cognitive sophistication.

Third, I find evidence of a negative shock size effect. Tripling the shock size from one month to three months of household income, the MPC falls by about 10%. The shock size effect mainly originates from a decline in the MPC conditional on spending any money, which is indicative of the intensive margin of adjustment.

Fourth, individuals who strongly adjust their MPC in response to an increase in the shock size are characterized by impatience, impulsiveness, and a lack of cognitive sophistication, and are more likely to face liquidity constraints. Importantly, liquid households also respond to changes in the shock size, i.e. they are able to smooth consumption but choose to do so only if the shock size is large. Further, cognitive sophistication and self-control problems explain a larger portion of the variance in the shock size effect than liquidity constraints. Taken together, the results suggest that behavioral characteristics are relatively more important than liquidity constraints in determining how consumers adjust their MPC in reaction to a rise in the shock size.

Lastly, consumption responses do not vary with the source of income. The MPC of those receiving a one-off payment from the government does not differ from those receiving a lottery win, implying that Ricardian equivalence seems to have no impact on consumption responses to stimulus payments.

The results are broadly in line with mental accounting theory. The difference in MPCs across payment modes suggests that consumers violate fungibility, i.e. they treat money differently depending on its labeling. The negative shock size effect matches the prediction that large windfalls are labeled as *current assets* and saved, while smaller

windfalls are labeled as *current income* and spent more easily (Thaler, 1999). Lastly, the prominent role played by self-control problems and cognitive sophistication in MPC heterogeneity is in line with the view that consumers use mental accounting to simplify decision-making and to deal with self-control problems. In contrast, liquidity constraints – as proposed by textbook models of consumption – can account for some MPC heterogeneity related to variation in the size of the income shock, while they cannot explain the differences in MPCs across payment modes.

The paper adds to the understanding of how consumers respond to unexpected, transitory income shocks and contributes to several streams of literature. First, the paper speaks to a broad literature aimed at estimating MPCs. One strand of literature follows a *revealed-preference* approach that draws inference on the MPC from spending data by exploiting quasi-experimental settings involving income shocks from lottery winnings (Fagereng et al., 2021), state-owned funds (Kueng, 2018), stimulus payments (Baker et al., 2020; Broda and Parker, 2014; Parker, Souleles, Johnson, and McClelland, 2013; Parker, 2017), tax refunds (Baugh, Ben-David, Park, and Parker, 2021), and dividend payments (Bräuer, Hackethal, and Hanspal, 2022). One constraint inherent in this approach is that it typically allows only one type of income shock to be studied at a time and yields MPCs only for those eligible for stimulus payments and exposed to a shock in natural experiments, respectively.

A complementary strand of literature uses a *reported-preference* approach to elicit MPCs, relying on survey questions that ask respondents to report their consumption responses either to stimulus payments already disbursed (Coibion et al., 2020; Shapiro and Slemrod, 2003; Sahm, Shapiro, and Slemrod, 2010) or to hypothetical income shocks (Bunn et al., 2017; Jappelli and Pistaferri, 2014, 2020).⁴ A key assumption under this approach is that survey responses are informative of actual consumption choices.⁵ This paper fits well into this strand of literature. Just like Christelis et al.

⁴Jappelli and Pistaferri (2020) provide an excellent review of the literature using survey questions to elicit MPCs. For a detailed discussion on the pros and cons of the *revealed-preference* and *reported-preference* approach, see Parker and Souleles (2019).

⁵Parker and Souleles (2019) compare MPCs elicited with the *revealed-preference* and *reported-preference* approach and conclude that both methods yield comparable spending propensities. Moreover, Parker, Schild, Erhard, and Johnson (2022) use a *revealed-preference* method to estimate MPCs out of the 2020 US stimulus payments and compare them with answers to a qualitative survey question on

(2019) and [Fuster et al. \(2021\)](#), I distinguish between small and large income shocks. I depart from them by additionally examining the effect of different payment modes and income sources on the MPC.

Second, the paper contributes to the literature on how mental accounting influences consumer behavior.⁶ Relying on laboratory experiments with students, early work has shown that acquired income is spent differently depending on its categorization ([Shefrin and Thaler, 1988](#); [Karlsson, Garling, and Selart, 1997](#); [Chambers and Spencer, 2008](#)). More recent work documents a violation of fungibility in the context of gasoline price shocks ([Hastings and Shapiro, 2013](#)), food stamps ([Hastings and Shapiro, 2018](#)), child allowances ([Kooreman, 2000](#)), restaurant vouchers ([Abeler and Marklein, 2017](#)), or coupon programs ([Milkman and Beshears, 2009](#); [Liu, Shen, Li, and Chen, 2021](#)). I add to this literature by providing evidence for consumer behavior consistent with mental accounting in a highly policy-relevant setting using a population-wide sample.

Lastly, the paper adds to an emerging body of empirical literature employing causal machine learning methods to study heterogeneous treatment effects on household behavior ([Davis and Heller, 2020](#); [Knittel and Stolper, 2019](#); [Farbmacher, Kögel, and Spindler, 2021](#); [Medina and Pagel, 2021](#)). To the best of my knowledge, this is the first paper that uses a causal forest to explore sources of MPC heterogeneity.

2 Data and Survey Design

This section describes the data sources and the survey design used to elicit MPCs out of different income shock scenarios and provides descriptive statistics of MPCs.

the use of the stimulus payments. They find that only those households exhibit a positive MPC estimate who also *report* that they mostly spent their stimulus payments.

⁶[Zhang and Sussman \(2018\)](#) and [Antonides and Ranyard \(2018\)](#) provide an excellent review of the literature.

2.1 Bundesbank Online Panel Households

The Bundesbank Online Panel Households (BOP-HH) constitutes the primary source of data. Similar to the New York Fed Survey of Consumer Expectations, the BOP-HH is a monthly internet-based survey focused on eliciting households' perceptions and expectations about various economic outcomes.⁷ A leading European survey company, Forsa, administers the survey and selects the gross sample using random sampling from the forsa.omninet database, with quotas for age, gender, and level of educational attainment. The sampling frame of the forsa.onminet database is individuals aged 14 years or older with internet access living in Germany.⁸ Each survey typically takes around 18 minutes to complete. Participation in the survey is incentivized to improve the quality of responses.⁹

I added a special module on the MPC to the December 2020 and June 2021 waves, which had a sample size of about 4,000 and 2,500 individuals, respectively. In total, I collected 6,373 consumption responses to hypothetical income shock scenarios.¹⁰ [Table 1](#) reports descriptive statistics of the survey.

2.2 Online Survey of Retail Customers at a Large German Bank

As a supplementary data source I use an online panel survey of retail customers at a large German bank with a branch network all over Germany.¹¹ The goal of the survey is to study retail customers' preferences, beliefs, as well as consumption and investment behavior. The survey was administered by the Goethe University in close cooperation with the bank. It was launched in May 2018 and continued on a roughly semi-annual

⁷For a detailed description of the survey, see [Beckmann and Schmidt \(2020\)](#).

⁸The forsa.omninet database consists of 75,000 individuals that were recruited by telephone. This offline recruitment method allows less internet-savvy respondents to be included in the sample and thus reduces a potential online selection bias.

⁹Participants in the BOP-HH receive bonus points as an incentive to participate in a survey wave. These bonus points can be redeemed for various small items.

¹⁰There were 187 non-responses to the MPC module, corresponding to an item non-response rate of less than 2.9%.

¹¹[Bräuer \(2021\)](#) provides a detailed description of the survey.

basis.¹² Respondents were recruited via email from a sampling frame comprising about 150,000 retail customers of the bank. About 4,000 respondents participated in the first wave and stayed in the panel until they dropped out voluntarily.¹³ A refresher sample of about 1,200 respondents was added in March 2020. Upon entry into the panel, all respondents had to complete a baseline questionnaire eliciting sociodemographic characteristics, while the follow-up waves asked them project-specific questions. Participation in each wave took about 10 minutes and was incentivized.¹⁴ I added a survey module on the MPC in the October 2020 wave, which comprised about 2,200 respondents. [Table D.10](#) reports summary statistics for this survey wave.

In the following, I describe the survey module administered in the BOP-HH. [Section 4.3](#) explains the survey experiment administered in the retail bank survey.

2.3 Eliciting the MPC

Using a between subject-design, I randomly split the sample into six equally-sized treatment groups and expose each group to a different hypothetical income shock scenario. I ask the baseline treatment group (BASE-1) about their consumption response to a hypothetical one month's income shock:¹⁵

BASE-1: one month's income shock

Imagine that you unexpectedly receive a one-off payment from the government, with no repayment obligation, in the amount of your monthly net household income. What part of this would you spend in the next twelve months and what part would you save or use to repay loans? Please state the percentage that you would spend.

The survey question has several noteworthy features:¹⁶ first, it yields a direct

¹²The follow-up survey waves were fielded in September 2018, February 2019, September 2019, March 2020, and October 2020.

¹³The invitation response rate was 2.7% ([Bräuer, 2021](#)). For an in-depth discussion of attrition across survey waves, see [Bräuer \(2021\)](#).

¹⁴Participants in each wave received online shopping vouchers worth €5.

¹⁵[Table D.9](#) shows that the different treatment groups are comparable along observable characteristics.

¹⁶The survey question is similar to [Jappelli and Pistaferri \(2014, 2020\)](#) and the MPC questions asked

quantitative estimate of the MPC as opposed to survey questions in the spirit of [Shapiro and Slemrod \(1995, 2003\)](#) that elicit qualitative spending responses ("*mostly spend*" / "*mostly save*"). Second, the question is explicit about the time horizon for spending ("*twelve months*"). It thus avoids artificial heterogeneity stemming from differences in the timing of spending and also allows predictions of the permanent income hypothesis to be tested. Third, the size of the income shock is proportional to respondents' household net income. This avoids unintended size effects – an absolute shock, rather than a proportional one, could be perceived as negligibly small by income-rich households.¹⁷ Lastly, the question explicitly reminds respondents of alternatives to spending, i.e. saving or repaying debt, in order to avoid priming participants towards a non-zero response.

2.4 Treatments

To study how consumers adjust their MPC in response to changes in the shock size and payment mode, I design the following five additional income shock scenarios:

CASH-1: one month's income shock paid in cash

SAVINGS-ACCOUNT-1: one month's income shock deposited in an instant-access savings account that pays interest

BASE-3: three months' income shock

CASH-3: three months' income shock paid in cash

SAVINGS-ACCOUNT-3: three months' income shock deposited in an instant-access savings account that pays interest

The treatments are tailored to test predictions of mental accounting theory. Ac-

in the 2017 wave of the Bundesbank's Panel on Household Finances (PHF). For a description of the PHF, see [Altmann, Bernard, Le Blanc, Gabor-Toth, Hebbat, Kothmayr, Schmidt, Tzamourani, Werner, and Zhu \(2020\)](#).

¹⁷For a discussion on the role of relative and absolute income shocks for MPC estimates, see also [Christelis et al. \(2019\)](#).

cording to the behavioral life-cycle hypothesis (Shefrin and Thaler, 1988), consumers categorize their resources into a system of broad mental accounts to simplify decision-making and to deal with self-control problems. Broad mental accounts differ in terms of how tempting it is to spend out of them, implying that the MPC is account-specific: first, the *current income* mental account, comprising cash and checking accounts, is the easiest to invade as funds in this account are designated for spending ($MPC=1$). Second, the *future income* account, comprising retirement savings, is considered off-limits for everyday consumption ($MPC=0$). Lastly, the *current assets* mental account encompasses several types of liquid assets such as discretionary savings, mutual funds, and stocks. Funds assigned to the *current assets* mental account are typically designated for saving ($MPC<1$). Consequently, the principle of fungibility of money is violated, i.e. income shocks are spent more readily or more reluctantly depending on how they are categorized or labeled.

Similarity judgments can influence how funds are labeled (Heath and Soll, 1996; Evers, Imas, and Kang, 2021), i.e. consumers associate a windfall with a particular mental account if it shares a salient attribute with a reference object in that mental account. For example, the size of a windfall relative to regular income receipts represents one salient attribute: the larger a windfall, the more wealth-like it is perceived to be, and hence the greater the likelihood of it being labeled as belonging to the *current assets* mental account rather than to the *current income* mental account. Accordingly, the MPC should decrease in the windfall's size (Shefrin and Thaler, 1988). To test this, I expose half of the sample to a one month's income shock (BASE-1, CASH-1, SAVINGS-ACCOUNT-1) and the other half to a three months' income shock (BASE-3, CASH-3, SAVINGS-ACCOUNT-3).

The payment mode of windfalls is another attribute that may influence how windfalls are labeled. Suppose consumers park funds associated with the *current assets* mental account in a particular bank account and they receive a windfall deposited into that bank account. In that case, they should be more likely to assign the windfall to the *current assets* mental account and mostly save it. Further, if consumers hold their *current income* in cash and receive a windfall paid out in cash, they should be more

likely to assign it to the *current income* account and mostly spend it. To test these predictions, I provide treatments on the payment mode of the income shock: one-third of respondents receive a "*one-off payment paid out in cash*" (CASH-1, CASH-3), which I refer to as cash payment scenario. Another third receives a "*one-off payment deposited in an instant-access savings account that pays interest*" (SAVINGS-ACCOUNT-1, SAVINGS-ACCOUNT-3), which I refer to as the savings account scenario.¹⁸ Note that instant-access savings accounts (*Tagesgeldkonto*) are highly liquid. Hence, transaction costs, as in the model of [Kaplan and Violante \(2014\)](#), are unlikely to drive differences in responses across payment mode scenarios.¹⁹ The remainder of the sample receives the income shock without any specification of the payment mode (BASE-1, BASE-3), referred to as the baseline scenario.

All respondents reporting that they would spend a positive amount are subsequently asked which categories of consumption they would use the money for, i.e. how much they would spend on durables (e.g. "*major purchases*", "*repairs*", "*education*"), nondurables (e.g. "*leisure spending*", "*gifts*", "*general living expenses*"), and other goods.²⁰ The follow-up question makes it possible to examine where the spending response to changes in the shock size and payment mode comes from.

¹⁸I used a slightly different wording in the June 2021 wave, referring to a classic savings account (*Sparkonto*) that is accessible at any time rather than to an instant-access savings account (*Tagesgeldkonto*) as in the December 2020 wave. In [Appendix C](#), I consider each of these savings account scenarios in isolation to study how consumption responses vary with the type of savings account. [Appendix C](#) shows that the savings account scenario in both the December 2020 and June 2021 waves induces consumers to adjust their MPC in the same direction. For that reason, I pool the different savings account scenarios in the main analysis to increase statistical power.

¹⁹Instant-access savings accounts (*Tagesgeldkonto*) allow savers to withdraw funds at any time, without facing any fees or being obligated to keep money in the account for a set period of time. I further discuss the role of transaction costs in [Section 5](#).

²⁰The survey question, along with the exact wording of the consumption categories, is shown in [Appendix A.1](#). The definition of consumption categories follows [Fuster et al. \(2021\)](#). To alleviate concerns about survey fatigue influencing results, observations are not considered in the analysis if they do not answer the follow-up question on the distribution of spending, despite reporting an intention to spend a positive amount in the first step of the MPC module.

2.5 Additional Measures

To explore possible heterogeneity in the consumption responses, I ask additional questions related to cognitive sophistication, self-control problems, and liquidity constraints, using an item battery listing statements to which respondents can disagree or agree with on a 7-point Likert scale.²¹

Cognitive sophistication Following Ameriks, Caplin, and Leahy (2003), I use respondents' self-assessed mathematical skills as a proxy for cognitive sophistication, asking respondents to what extent they agree with the statement *"I have a lot of confidence in my mathematical skills."* I invert the scale of the item such that higher values indicate lower mathematical skills. Further, a lack of financial planning can be a symptom of low cognitive sophistication. I elicit respondents' propensity to plan with the following item: *"I plan major spending and investment decisions more than one year ahead."*²²

Self-control problems I rely on various items to measure different dimensions of self-control problems. First, I use respondents' ability to commit to pre-set goals as a proxy for willpower, asking them to what extent they *"actively follow through with the plans [they] make."*²³ Second, I elicit respondents' degree of impatience with the item *"I am generally a very patient person."* I invert the scale of the variable such that higher values indicate being more impatient. Lastly, I elicit different aspects of trait impulsivity: to capture motor impulsiveness I use the item *"I rarely do anything without thinking about it thoroughly"* and invert its scale such that higher values indicate a stronger tendency to act without thinking. Further, I elicit myopic decision-making as part of non-planning impulsiveness with the item *"I live in the here and now and don't really think about the future."*²⁴

²¹Other questions of the core survey were asked between the MPC module and the item battery to ensure that responses to them are not affected by the income shock scenarios and vice versa.

²²The item is adapted from the PHF.

²³The survey item is adapted from the Health and Retirement Survey (HRS).

²⁴Both survey items are adapted from the HRS.

Liquidity constraints To measure liquidity constraints, I ask respondents whether *"[they] have put aside money for a possible emergency so that [they] can cover expenses for at least three months with no income."* I invert the scale of the item such that higher values indicate a higher probability of being liquidity-constrained. The question is similar to [Parker \(2017\)](#) and inspired by the classic buffer-stock model ([Carroll, 1997](#)) in which consumers try to reach a target ratio between cash-on-hand (wealth plus income) and the permanent component of income.²⁵

2.6 Descriptive Statistics

Average MPCs Table 2 reports summary statistics for the consumption response to different income shock scenarios. As the size of the income shock is tied to household income, the one month's income shock amounts to about €3,400 for the average respondent.²⁶ The average MPC out of the one month's income shock comes to 52% in the baseline scenario (BASE-1). The MPC estimate sits at the upper end of estimates in the literature, and is remarkably close to other reported MPC studies ([Drescher, Fessler, and Lindner, 2020](#); [Jappelli and Pistaferri, 2014](#)).²⁷

The MPC falls to 47% if the shock size increases to three months of household income (BASE-3). Further, the MPC varies with the payment mode. Respondents exposed to a one month's income shock deposited in an instant-access savings account (SAVINGS-ACCOUNT-1) exhibit an MPC of 45% compared to an MPC of 54% for those exposed to an equally-sized shock paid out in cash (CASH-1).

MPC distributions The large standard deviations indicate that the MPC is very heterogeneous among respondents. To further illustrate this point, [Figure 1a](#) plots the

²⁵Around 30% of respondents in the sample lack a buffer-stock of savings of at least three months of income, i.e. scoring 4 or lower on the inverted scale. The number matches quite well the fraction of poor and wealthy hand-to-mouth households of about 30% in Germany, as estimated by [Kaplan, Violante, and Weidner \(2014\)](#).

²⁶Using the midpoint of each income category, the average household net income is at around €3,400 in the sample.

²⁷Using data from a euro area-wide survey (HFCS), [Drescher et al. \(2020\)](#) document an average MPC of about 51% for Germany. [Jappelli and Pistaferri \(2014\)](#) use a survey representative for Italian households (SHIW) and find an average MPC of about 48%.

distribution of MPCs out of one month's income shocks (blue bars) and three months' income shocks (red bars).²⁸ Both distributions exhibit heaping in the middle and at the tails, as is often found in reported MPC studies (Jappelli and Pistaferri, 2014, 2020; Drescher et al., 2020): around 15% of respondents spend nothing out of the small income shock (MPC = 0), around 20% spend the entire shock (MPC=100), and around 27% spend half of the shock (MPC=50).²⁹ Increasing the size of the income shock reduces heaping, in particular, the share of respondents consuming everything (MPC=100). At the same time, it raises the share of respondents spending less than half of the income shock.

Figure 1b shows MPC distributions for the different payment mode scenarios.³⁰ The number of individuals not spending any money (MPC=0) is much larger in the savings account scenario compared to the baseline and cash payment scenario, suggesting that the difference in MPCs across payment modes originates from the extensive margin of spending adjustment.

Average MPC and sociodemographic characteristics To assess the external validity of the MPC estimate, I study how the MPC out of the one month's income shock (BASE-1) relates to sociodemographics. The upper left panel in Figure 2 shows that the average MPC increases up to the age of 54 and drops strongly after retirement, corroborating the findings of Jappelli and Pistaferri (2014) and Andreolli and Surico (2021). Further, low-income households exhibit a higher average MPC than middle-income and high-income households (upper right panel), consistent with Baker et al. (2020), Christelis et al. (2019), and Coibion et al. (2020). Income losses due to the COVID-19 pandemic are associated with a higher MPC (lower left panel), in line with Baker et al. (2020), who document a strong consumption response to stimulus payments among US households with income losses induced by the COVID-19 pandemic. Lastly, respondents lacking a liquidity buffer of at least three months of income have a slightly higher MPC than others, consistent with Parker (2017) and Coibion et al. (2020). In

²⁸The different payment mode treatment groups are pooled for each shock size in Figure 1a.

²⁹Studying subcategories of spending, Figure D.1 shows the heterogeneity in spending is present for almost all subcategories of consumption.

³⁰The different shock size treatments are pooled for each payment mode.

sum, the survey results on the MPC estimate are broadly in line with those of other studies.

3 Econometric Framework

This section lays out the econometric framework that I use to estimate how consumers adjust their MPC in response to changes in the size and payment mode of income shocks.

3.1 Estimation of Average Treatment Effects

I estimate average treatment effects (ATEs) with the following equation using OLS:³¹

$$MPC_i = \alpha + \gamma \times LargeShock + \sum_{s=1}^{S=2} \beta_s \times PaymentMode_s + \delta \mathbf{X}_i + \epsilon_i, \quad (1)$$

where MPC_i is the share (scaled to be between 0 and 100) of individual i 's income shock that is used for spending. *LargeShock* is a randomized binary treatment variable that equals 1 if respondent i received a three months' income shock (BASE-3, CASH-3, SAVINGS-ACCOUNT-3) and 0 if respondent i received a one month's income shock (BASE-1, CASH-1, SAVINGS-ACCOUNT-1). The γ coefficient identifies the ATE of increasing the shock size from one month to three months of income, which I refer to as the shock size effect. *PaymentMode_s* is a randomized binary treatment variable that equals 1 if respondent i receives an income shock paid out in cash (CASH-1, CASH-3) and without any specification of the payment mode (BASE-1, BASE-3), respectively, and 0 if respondent i receives an income shock deposited in an instant-access savings account (SAVINGS-ACCOUNT-1, SAVINGS-ACCOUNT-3). The β_s coefficients identify the ATEs of the different payment modes. \mathbf{X}_i is a vector of household and individual-specific control variables.³²

³¹The ATE is defined as $\tau = \mathbb{E}[MPC(1) - MPC(0)]$, where $MPC(1)$ and $MPC(0)$ refer to the potential outcomes that one would observe given that individual i received treatment ($W = 1$) or control ($W = 0$).

³²The control variables have no material effect on the estimates because the assignment to different income shock scenarios is randomized. Nevertheless, I incorporate them to increase the precision of

3.2 Estimation of Heterogeneous Treatment Effects

To explore why consumers respond to the different income shock scenarios, I concentrate on estimating conditional average treatment effects (CATEs), which are defined as $\tau(x) = \mathbb{E}[MPC(1) - MPC(0)|X = x]$.

The standard approach to estimating CATEs proposes to either include an interaction term of the treatment variable and the heterogeneity variable of interest in the regression model (e.g. interact a treatment dummy with age) and test for its significance, or estimate the treatment effect for subgroups of interest and study how the treatment effect varies across these subgroups (e.g. young vs. old). One problem with this approach is that one might capture spurious heterogeneity, as coefficients can turn significant by pure chance if one tests for multiple heterogeneity variables of interest (multiple hypothesis testing problem). Further, specifying a correct interaction model a priori is challenging if treatment heterogeneity is characterized by a non-linear function of a single variable or by an interaction of more than one variable at a time. Consequently, the standard approach may miss important (unexpected) treatment heterogeneity.

To overcome these problems, I apply the causal forest algorithm developed by [Athey et al. \(2019\)](#), which predicts a treatment effect for each individual based on high-dimensional, non-linear functions of observable characteristics. The algorithm proceeds in two steps: first, the data is recursively split into subsamples subject to a splitting criterion. Second, a treatment effect is estimated within each subsample and used as a prediction for the treatment effect of each observation that falls into that particular subsample. Importantly, the splitting criterion is tailored to address treatment effect heterogeneity. It chooses splits that maximize the variance of treatment effects across subsamples by searching across all combinations of covariate values.³³ The algorithm thus learns in a flexible, non-parametric, data-driven way

the estimates. The vector of household and individual-specific control variables comprises age, age squared, female, unemployment indicator, homeownership status, education (indicator variable for each category), household income (indicator variable for each category), household size (indicator variable for each category), living in eastern Germany and city size (indicator variable for each category).

³³The splitting rule of regular random forests optimizes for heterogeneity in predicted values \hat{Y} , to maximize predictive accuracy. In contrast, the splitting rule of the causal forest seeks to maximize

which combinations of covariates are relevant for treatment heterogeneity, including non-linearities and interactions of covariates, that would be difficult to model with a standard parametric approach to treatment heterogeneity. Further, the causal forest avoids capturing spurious treatment heterogeneity by relying on an *honest* splitting technique that uses different partitions of the data for splitting and for treatment effect estimation. Lastly, to reduce the uncertainty of predictions the algorithm iterates the sub-sampling procedure many times and takes a weighted average over all predictions for each individual.

I estimate two causal forests: one for the shock size effect, and one for the payment mode effect. For the latter, I pool the treatment groups from the cash payment scenario and the baseline scenario to increase statistical power, i.e. I use a treatment dummy that equals 1 if respondent i receives a windfall either paid out in cash (CASH-1, CASH-3) or without any specification of the payment mode (BASE-1, BASE-3), and 0 when the respondent receives an income shock deposited in an instant-access savings account (SAVINGS-ACCOUNT-1, SAVINGS-ACCOUNT-3).³⁴ For each causal forest, I use 18 variables as potential covariates to split, including measures for cognitive sophistication, self-control problems, and liquidity constraints, and a vector of sociodemographics. As a result, I obtain a prediction for each individual of how much she adjusts her MPC in response to changes in the shock size and the payment mode, respectively. [Appendix B](#) explains the intuition and implementation of the causal forest algorithm in detail.

4 Results

In the following, I examine how consumers respond to different types of windfalls. I first focus on how a windfall's size and payment mode affects the MPC on average. I then turn to MPC heterogeneity. Lastly, I consider how the MPC varies with the

treatment heterogeneity across sample splits. That's what makes it *causal*.

³⁴In [Appendix C](#), I use a standard sample-splitting approach to analyze treatment heterogeneity. [Table D.8](#) shows that consumption responses to the cash payment and the baseline scenario are characterized by similar heterogeneity patterns. It therefore appears reasonable to pool the cash payment and the baseline scenario for the heterogeneity analysis.

income source.

4.1 Average Effects of Payment Mode and Shock Sizes Scenarios

Table 3 reports regression results for the specification in Equation 1 with and without control variables. Tripling the shock size from one month to three months of household income, respondents reduce their MPC by 4.9 percentage points. This negative size effect is statistically and economically significant as it corresponds to a 10% decrease in the MPC.³⁵ The result is consistent with mental accounting theory and models of consumption with liquidity constraints and corroborates the findings of Andreolli and Surico (2021), Christelis et al. (2019), and Fagereng et al. (2021) that MPCs decline with the shock size.³⁶

The MPC varies with the payment mode of the income shock. Respondents who receive an income shock paid out in cash and without any specification of the payment mode, respectively, exhibit a 5.9 and 6.9 percentage points higher MPC than those receiving an income shock deposited in an instant-access savings account. The effect is statistically significant and large in economic terms as it corresponds to a relative increase in the MPC of about 14% and 16%. While the point estimate is larger for the cash payment scenario than for the baseline scenario, I cannot reject the null hypothesis that the two coefficients are equal (p-value=0.33).³⁷ Overall, the payment mode effect is in line with the hypothesis that consumers violate the principle of fungibility as they categorize windfalls as belonging to different mental accounts depending on windfalls' salient attributes and spend them differently once they are categorized.

³⁵I further analyze how increasing the shock size affects spending on different categories of durable and non-durable goods and services. Table D.2 shows that respondents reduce spending on almost all categories of consumption, except for large durables, likely because a sufficiently large shock allows them to buy more expensive durables.

³⁶Fuster et al. (2021) find average MPCs increase with the shock size. However, their study considers consumption responses to absolute gains instead of relative gains, i.e. proportional to household income, as in this paper. The results might thus be not directly comparable.

³⁷Table D.2 shows that the adjustment in spending in response to changes in the payment mode comes from both durable and nondurable goods and services.

4.2 Heterogeneity in Consumption Responses

Intensive and extensive margin The difference in MPCs across income shock scenarios might originate either from differences in the number of individuals spending any money (extensive margin) or from changes in the MPC conditional on spending any money (intensive margin). To study the role of the extensive margin, I use a dummy as outcome variable that equals 1 if respondent i spends any money ($MPC > 0$), and 0 otherwise. Columns (3) and (4) of [Table 4](#) report marginal effects from logistic regressions. If the income shock is tripled, the fraction of respondents spending any money increases slightly by 2 percentage points (column 2), suggesting the extensive margin of adjustment is of minor importance for the shock size effect.

To study the intensive margin, I regress the MPC on the treatments conditional on a positive MPC. Columns (5) and (6) illustrate that the average MPC conditional on spending falls by more than 7 percentage points if the shock size increases. Taken together, the results suggest that most of the adjustment in spending in response to an increase in the shock size comes from the intensive margin.

Turning to the payment mode effect, I find a significant role for the extensive margin. The fraction of respondents spending any money is about 6 and 7 percentage points higher among those receiving a one-off payment without any specification of the payment mode and paid out in cash, respectively, compared to those receiving a one-off payment deposited in an instant-access savings account (column 3). The results suggest a subgroup of individuals does not consume out of a windfall if it is deposited in an account that is typically associated with the *current assets* mental account. The intensive margin, on the other hand, is statistically significant, though the estimated coefficients are only about half as large as those for the total effect (column 5).

Conditional Average Treatment Effects Next, I estimate conditional average treatment effects using the causal forest. Based on observable characteristics I predict how each individual adjusts her MPC in response to changes in the size and the payment mode of income shocks, respectively.

The causal forest identifies substantial heterogeneity in treatment effects. [Figure 3a](#) plots the distribution of the predicted effect of increasing the shock size on the MPC. The 25th and 75th percentile of the distribution is at -6.4 and -3.3 percentage points, respectively, indicating that some individuals (in the left tail) strongly reduce their MPC when the shock size triples from one month to three months of household income, while others (in the right tail) hardly react. Similarly, the causal forest predicts heterogeneity in consumption responses to changes in the payment mode. The 25th and 75th percentile the distribution is at 5.1 and 8.0 percentage points, respectively, as shown in [Figure 3b](#).

[Appendix B.2](#) examines the validity of the causal forest's predictions. All the tests performed yield the same conclusion that the individualized predictions for the shock size and payment mode effects capture actual treatment heterogeneity in the data.

Who reacts the most? To understand where the heterogeneity comes from, I study how those adjusting their MPC most strongly in response to changes in the shock size and payment mode, respectively, differ along observables from those who hardly react. To this end, I plot differences in means between respondents located in the bottom and the top quartile of each predicted treatment effect distribution.

[Figure 4a](#) shows that respondents who strongly reduce their MPC in response to an increase in the shock size are more likely to be impatient, act without thinking, and take myopic decisions, and less likely to commit to pre-set goals, than those who hardly react. They also have a lower propensity to plan their finances and exhibit lower mathematical skills. The pattern is consistent with mental accounting being used to simplify decision-making and as a self-commitment device in particular by those consumers who are less cognitively sophisticated and prone to self-control problems.

Further, the share of liquidity-constrained individuals is higher among those who react strongly, in line with the textbook model prediction that *temporary* liquidity-constrained consumers are more likely to overcome constraints and hence smooth consumption when the shock size is large. In line with this thinking, one would expect that many of those reacting strongly recently lost income. However, quite the opposite

is the case – the share of individuals who lost income during the pandemic is lower. Taken together, these results might suggest that the higher prevalence of liquidity constraints does not result from *temporary* shocks, but rather reflect a personal trait of holding *permanently* low liquidity.³⁸

Turning to sociodemographic characteristics, individuals predicted to reduce their MPC most strongly when the size of the income shock increases are older, earn a higher income, and are more likely to own their home.

Figure 4b reveals that those who adjust their MPC most strongly in response to changes in the payment mode are significantly more impatient, and more likely to act without thinking, and less likely to plan their finances or face liquidity constraints, than those who adjust their MPC the least. These results suggest that those individuals who violate the principle of fungibility do have to cope with self-control problems and are less cognitively sophisticated, as implied by mental accounting theory. Concerning sociodemographics, the higher share of men, high earners, and homeowners among those responding strongly stands out.

The role of behavioral characteristics and liquidity constraints The previous analysis has shown that behavioral characteristics and liquidity constraints seem to matter for consumption responses to transitory income shocks. Next, I examine partial dependence plots (PDs) to study how on average the treatment effect varies with one of these characteristics, while holding all other characteristics constant.³⁹

Figure 5 reveals that the shock size effect becomes more negative with self-control problems: the higher the respondents' degree of impatience (upper left panel), holding

³⁸The observed pattern is consistent with a growing body of evidence that a lack of consumption smoothing is a persistent behavioral characteristic associated with holding low liquidity, rather than being primarily driven by transitory income shocks or temporary low liquidity (Parker, 2017; Baugh et al., 2021; Aguiar, Bils, and Boar, 2020).

³⁹To construct PD plots introduced by Friedman (2001), I first force each observation of the sample to assume a particular value of the heterogeneity variable of interest, while holding all other observable characteristics fixed at their empirical value. Feeding the causal forest with these artificial data points, I then predict a treatment effect at the particular value for each observation. Lastly, I average over the sample distribution. Repeating these steps for all values of the heterogeneity variable, I obtain the marginal impact of the variable on the treatment effect. In Appendix C, I show that the results from the PD plots are robust to taking potential correlations of the variable of interest with other confounding variables into account.

all other covariates constant, the more they reduce their MPC if the shock size triples from one month to three months of household income. A similar, albeit weaker, pattern appears for measures of impulsiveness: the higher the respondents' tendency to act without thinking (upper middle panel) and engage in myopic decision-making (upper right panel), respectively, the more they adjust spending, consistent with individuals using mental accounting to cope with problems of self-control.

Further, the lower the respondents' confidence in their mathematical skills (lower left panel) and the lower their propensity to plan major spending and investment decisions (lower middle panel), respectively, the larger is the difference in MPCs between the one month's and the three months' income shock. The results are consistent with less sophisticated individuals following a mental accounting heuristic to reduce cognitive costs of optimization, i.e. consume small income shocks and save larger ones.

Liquidity constraints are associated with a stronger consumption response (lower right panel). In the event that households lack a buffer stock of savings (i.e. they score 4 or higher on the x-axis), they reduce their MPC with the shock size, consistent with consumption models with liquidity constraints ([Carroll, 1997](#)) or hand-to-mouth behavior ([Kaplan and Violante, 2014](#)). However, in the event that households are liquid (i.e. they score 3 or lower) they also adjust their spending in response to changes in the shock size, albeit less strongly, i.e. households that are able to smooth consumption choose to do so only if the shock size is large. The latter result is difficult to reconcile with textbook models of consumption and illustrates that the shock size effect is not driven purely by liquidity constraints.⁴⁰

Turning to the payment mode effect, the upper panel in [Figure 6](#) illustrates that the effect grows with self-control problems. In particular, the higher the degree of impatience, the larger the difference in MPCs between a windfall deposited in a savings account and a windfall paid out in cash or without any specification of the payment mode. A similar, albeit less pronounced pattern emerges for the tendency to

⁴⁰Models of consumption with liquidity constraints and precautionary savings predict that unconstrained households holding a sufficiently high buffer stock of savings should always smooth consumption and hence should not respond to changes in the shock size.

act without thinking (upper middle panel) and to engage in myopic decision-making (upper right panel), respectively. The results illustrate in another way that individuals facing self-control problems are more likely to spend windfalls differently depending on its salient attributes.

Further, low mathematical skills (lower left panel) and a lower propensity to plan one's own finances (lower middle panel) are associated with a stronger adjustment in the MPC in response to changes in the payment mode, consistent with mental accounting being used to ease the cognitive load of decision-making. Lastly, the more likely it is that respondents face liquidity constraints (lower right panel), the less they adjust their MPC, possibly because liquidity-constrained households exhibit a higher MPC in general, leaving less room for an upward adjustment of the MPC.

Importance of liquidity and behavioral characteristics Lastly, I focus on the relative importance of liquidity constraints and behavioral characteristics for the heterogeneity in consumption responses to changes in the shock size. To this end, I examine how much variation in the shock size effect is accounted for by each characteristic. Using the individualized predictions of the shock size effect as pseudo outcomes, I linearly project liquidity constraints and behavioral characteristics onto the pseudo outcome. I then calculate the semi-partial R^2 for each characteristic, indicating the unique variance explained by the respective characteristic. [Figure 7](#) reveals that financial planning and impatience exhibit a semi-partial R^2 of around 0.12 and 0.04, respectively. In contrast, liquidity constraints exhibit a semi-partial R^2 of below 0.01. The finding suggests that cognitive sophistication and self-control problems are relatively more important than liquidity constraints in determining how consumers adjust their MPC in reaction to a rise in the shock size.

4.3 Consumption Responses by Source of Income

In assigning funds to different mental accounts, consumers may consider not only the size and payment mode but also the origin of funds. Prior research shows that pure windfall gains such as lottery wins are spent more readily than money from other

sources (Arkes, Joyner, Pezzo, Nash, Siegel-Jacobs, and Stone, 1994; Henderson and Peterson, 1992). It is thus instructive to study how the MPC varies with the source of income shocks.

To this end, I run another experiment using an online survey of retail customers at a large German bank. I generate exogenous variation in the source of the income by randomly splitting the sample into four equally-sized groups and exposing each group to a different hypothetical income shock.⁴¹ Half of the respondents are asked about their consumption response to a one-off payment from the government, with no repayment obligation, in the amount of their monthly household net income, just as in the BASE-1 scenario in the BOP-HH. The other half of the sample is asked about their response to an equally-sized lottery win. Additionally, I vary the payment mode within these two groups, i.e. respondents either receive the windfall paid out in cash or without any specification of the payment mode:⁴²

GOV-BASE: one month's income shock from the government

GOV-CASH: one month's income shock from the government paid in cash

LOTTERY-BASE: one month's income shock from a lottery win

LOTTERY-CASH: one month's income shock from a lottery win paid in cash

The average MPC in both the GOV-BASE and LOTTERY-BASE scenarios is at around 42%, whereas the average MPC is at 44% and 46% in the GOV-CASH and LOTTERY-CASH scenarios, respectively, as shown in Table D.11.⁴³ The MPC estimates are lower than those from the BOP-HH, possibly due to differences in respondents' characteristics.⁴⁴

To estimate the causal effect of the source and payment mode of income shocks

⁴¹Table D.11 shows that the four treatment groups are comparable along observable characteristics.

⁴²Appendix A.2 reports the exact wording of the survey questions.

⁴³The relatively high MPC out of the lottery win is consistent with Fagereng et al. (2021) who, using Norwegian administrative data on sizable lottery prizes, document a within-year expenditure response of around one-half.

⁴⁴Table D.12 compares observable characteristics between the retail bank survey and the Bundesbank survey (BOP-HH). Respondents to the retail bank survey are younger, more likely to be male, more educated, and less likely to live in eastern Germany, compared to respondents to the BOP-HH.

on the MPC, I use the following OLS regression equation:

$$MPC_i = \alpha + \beta \times Gov_i + \gamma \times Cash_i + \delta X_i + \epsilon_i, \quad (2)$$

where MPC_i is the share of individual i 's income shock that went to consumption. Gov_i is a treatment dummy that equals 1 if respondent i received a one-off payment from the government and 0 if she received a lottery win. $Cash_i$ is a treatment dummy that equals 1 if the transfer was paid out in cash and 0 otherwise. X_i is a vector of sociodemographic controls.⁴⁵

Table 5 shows that the MPC does not vary with the source of income. The MPC of those receiving a one-off payment from the government is not statistically different from those receiving a lottery win. The result is consistent with consumers labeling one-off payments from the government as pure windfalls, which in turn could explain the widely documented lack of consumption smoothing in response to stimulus payments.⁴⁶ Aside from this, the finding that the MPC out of the government transfer is just as high as it is for the lottery win indicates that consumers disregard general equilibrium effects and suggests that Ricardian equivalence does not influence consumption responses to stimulus payments.⁴⁷

Further, Table 5 shows again that the payment mode matters for the MPC. Respondents exposed to an income shock paid out in cash have a 3 percentage point higher MPC than those exposed to an income shock without any specification of the payment mode. The direction of the effect is consistent with the view that consumers assign a windfall to the *current income* mental account and spend it if the windfall

⁴⁵Sociodemographic controls are gender, age, age squared, indicator for unemployment, college degree, household income (indicator variable for each category), and household size (indicator variable for each size).

⁴⁶For an overview of MPC estimates out of stimulus payments, see Coibion et al. (2020).

⁴⁷For example, one might think that consumers anticipate general equilibrium effects such as a change in prices or interest rates as a consequence of higher aggregate consumption induced by stimulus payments. Ricardian equivalence proposes that consumers are forward-looking and take the government's budget constraint into account in their consumption choices. Accordingly, consumers should save more and hence exhibit a lower MPC when receiving a government transfer compared to a lottery win, as they anticipate future tax hikes to finance the government transfer. However, this prediction is at odds with the results.

exhibits an attribute, i.e. hard cash, that it has in common with resources typically associated with that mental account.

5 Robustness Checks

In [Appendix C](#), I assess the sensitivity of the findings to various robustness checks. In the following, I briefly summarize the main results.

Heaping One concern with the MPC estimate is that its distribution exhibits heaping at focal values of 0, 50, and 100%. Heaping could reflect respondents' uncertainty and hence might be related to specific characteristics such as low cognitive sophistication. However, [Table D.1](#) shows that neither respondents' mathematical skills, nor education predict their choice to report 0, 50, and 100% as a response.⁴⁸

Misunderstanding of the question A related problem might be that some respondents do not fully comprehend the hypothetical income shock scenarios. Excluding those more prone to misunderstand or misinterpret survey questions, the shock size and payment mode effect, respectively, on the MPC remain in the same ballpark, as shown in [Table D.3](#).

Censoring of the outcome variable Another potential concern is that the MPC variable is censored from below, at 0, and above, at 100, which can bias OLS estimates. Using a two-limit Tobit estimator instead of OLS, [Table D.4](#) shows that the estimated shock size and payment mode effects are almost identical across the two estimators.

Transaction costs One alternative explanation for the difference in MPC across payment modes is that those individuals who obtain the windfall on an instant-access savings account consider it too burdensome to transfer the money to another account to spend it. In June 2021 wave of the BOP-HH, participants were asked directly about

⁴⁸For a detailed discussion on the role of heaping in reported MPC measures, see also [Jappelli and Pistaferri \(2020\)](#).

their reasons for spending. [Figure D.2](#) shows that the share of respondents reporting motives related to transaction costs is negligible in the group predicted to react most strongly to changes in the payment mode. Importantly, within this group, a relatively high share states that they "*never spend*" or "*do not like to spend money from their savings account*", in line with consumers associating their savings account with the *current assets* mental account.

Macroeconomic environment Another concern might be that the shock size and payment mode effect on the MPC is contingent on the macroeconomic environment, as the survey module was fielded during the COVID-19 pandemic. To address this point, I consider results from a pilot study administered in the retail bank survey in September 2019. In a within-subject design, MPCs out of different shock sizes were elicited. Additionally, the payment mode was varied among respondents. [Table D.5](#) shows that the MPC out of a large shock is significantly lower than the MPC out of a small shock. Further, the MPC differs significantly across payment modes, suggesting that the external validity of the results is not limited to the period of the COVID-19 pandemic.

6 General Discussion

This section discusses the findings in light of the theories that aim to explain how consumers respond to unexpected, transitory income shocks.

Mental accounting The preferred interpretation of the results is that consumption responses to windfalls can best be understood through the lens of a model of mental accounting in which consumers categorize their resources into a system of broad mental accounts and use resources differently once they are categorized ([Shefrin and Thaler, 1988](#)).

The finding that the MPC falls with the size of the shock is consistent with the model's prediction that large windfalls are labeled as *current assets* and saved, while smaller windfalls are labeled as *current income* and spent more easily. The relatively

low MPC out of a windfall deposited in an instant-access savings account as compared to the higher MPC out of a windfall paid out in cash is consistent with consumers labeling a windfall as belonging to a specific mental account if it shares a salient attribute with a reference object of that mental account and using the money accordingly. Further, the evidence that impatience, impulsiveness, low mathematical skills, and a lack of financial planning account for heterogeneity in MPCs across different types of income shocks is in line with mental accounting being used as a tool to simplify decision-making and to solve self-control problems. Lastly, the finding that the MPC out of a one-off payment from the government equals the MPC out of a lottery win might suggest that consumers label the former as true windfall money.

Permanent income hypothesis A finite horizon model with quadratic utility and no borrowing constraints predicts that consumption is proportional to lifetime disposable resources and hence the MPC out of unexpected, transitory income gains is close to zero, independent of the shock size. This prediction is at odds with the finding that the MPC is large on average and falls with the size of the gain.

Liquidity constraints Workhorse models of consumption with temporary binding liquidity constraints suggest that the MPC falls with the shock size: constrained consumers spend small windfalls in full to close the gap between desired and actual consumption, while the constraint is less likely to bind if the shock size is sufficiently large, thereby allowing consumption to be smoothed ([Christelis et al., 2019](#)). This mechanism is consistent with the evidence of a higher MPC out of the small windfall as compared to the large windfall.⁴⁹ However, liquidity constraints cannot explain why those households who hold a buffer stock of savings adjust their MPC in response to an increase in the shock size. Further, the differences in MPCs across payment modes are difficult to reconcile with textbook models of consumption.

⁴⁹The size-asymmetry is also consistent with hand-to-mouth behavior as in the model of [Kaplan and Violante \(2014\)](#) featuring a low-return liquid asset account and a high-return illiquid asset account that can be accessed only by paying a transaction fee. In such a model, hand-to-mouth households spend small windfalls in full, while a sizeable windfall induces many of them to pay the transaction fee and deposit the windfall into the illiquid account to save it.

Consumption adjustment costs [Fuster et al. \(2021\)](#) consider a precautionary savings model with non-pecuniary costs of deviating from past consumption, which prevents consumers from adjusting consumption in response to small income shocks. Such a model generates a positive size effect on the extensive margin, consistent with the result that the share of respondents reporting that they spend any money increases with the windfall's size. However, the finding that the positive effect on the extensive margin is dominated by a negative size effect on the intensive margin – overall, the average MPC falls with the windfall's size – suggests that consumption adjustment costs play only a minor role for MPC heterogeneity.⁵⁰

Non-homothetic preferences for non-essential goods Among affluent households, liquidity constraints are less likely to bind, whereas consumption of luxury goods and services might be a stronger driver of the MPC. Following this idea, [Andreolli and Surico \(2021\)](#) consider a model with non-homothetic preferences for non-essential consumption which predicts that i) the MPC increases with cash-on-hand for large income gains, and ii) the MPC increases with the size of the windfall among affluent households. The evidence of a high MPC for unconstrained households is consistent with such a model. However, the finding that the MPC declines with the shock size among liquid households runs counter to the models' predictions.⁵¹

7 Summary and Concluding Remarks

This paper provides new causal evidence on how consumers respond to unexpected, transitory income shocks. Using a randomized control trial to generate exogenous variation in the type of income shocks, I elicit consumption responses in different hypothetical income shock scenarios. I find that MPCs decrease with the size of in-

⁵⁰The model of [Fuster et al. \(2021\)](#) is calibrated to their empirical findings so that, overall, the MPC increases with the size of the windfall.

⁵¹[Andreolli and Surico \(2021\)](#) compare MPCs out of shocks equal to one month and twelve month household income and find that among affluent households the larger windfall is associated with a higher MPC. Given that this is not the case when comparing MPCs out of one month's and three months' income shocks, non-homothetic preferences thus may become relevant only if the income shock is very large.

come shocks. Further, the consumption response to a windfall paid out in cash is significantly larger than the response to a windfall deposited in an instant-access savings account, suggesting that consumers violate fungibility. Causal machine learning methods uncover substantial heterogeneity in consumption responses to income shocks across consumers. Liquidity constraints contribute to the variance in responses. However, self-control problems and low cognitive sophistication are relatively more important for MPC heterogeneity. Overall, the results are consistent with a model of mental accounting in which consumers label windfalls as belonging to different mental accounts depending on the windfalls' attributes and use them differently once they are labeled.

The results have important implications for the design of stimulus programs. The findings suggest that small income shocks are likely to be spent, while larger ones are more likely to be saved. A stimulus program of a given size can thus be more effective in boosting consumption if the money is transferred to households in relatively small installments rather than as a lump sum. Further, consumers seem to spend windfalls differently depending on the payment mode. Policymakers should therefore take care of how stimulus payments are delivered. One-off payments in the form of cash can be successful in inducing consumers to label and use them for consumption purposes.

References

- Abeler, J. and F. Marklein (2017). Fungibility, labels, and consumption. *Journal of the European Economic Association* 15(1), 99–127.
- Aguiar, M., M. Bils, and C. Boar (2020). Who are the hand-to-mouth? NBER Working Paper 26643.
- Altmann, K., R. Bernard, J. Le Blanc, E. Gabor-Toth, M. Hebbat, L. Kothmayr, T. Schmidt, P. Tzamourani, D. Werner, and J. Zhu (2020). The Panel on household finances (PHF) – microdata on household wealth in Germany. *German Economic Review* 21(3), 373–400.

- Ameriks, J., A. Caplin, and J. Leahy (2003). Wealth accumulation and the propensity to plan. *The Quarterly Journal of Economics* 118(3), 1007–1047.
- Andreolli, M. and P. Surico (2021). Less is more: Consumer spending and the size of economic stimulus payments. CEPR Discussion Paper 15918.
- Antonides, G. and R. Ranyard (2018). Mental accounting and economic behaviour. In R. Ranyard (Ed.), *Economic Psychology*, pp. 123–138. Hoboken: Wiley-Blackwell.
- Arkes, H., C. Joyner, M. Pezzo, J. Nash, K. Siegel-Jacobs, and E. Stone (1994). The psychology of windfall gains. *Organizational Behavior and Human Decision Processes* 59(3), 331–347.
- Athey, S., J. Tibshirani, and S. Wager (2019). Generalized random forests. *The Annals of Statistics* 47(2), 1148–1178.
- Athey, S. and S. Wager (2019). Estimating treatment effects with causal forests: An application. *Observational Studies* 5(2), 37–51.
- Baker, S., R. Farrokhnia, S. Meyer, M. Pagel, and C. Yannelis (2020). Income, liquidity, and the consumption response to the 2020 economic stimulus payments. NBER Working Paper 27097.
- Baugh, B., I. Ben-David, H. Park, and J. Parker (2021). Asymmetric consumption smoothing. *American Economic Review* 111(1), 192–230.
- Beckmann, E. and T. Schmidt (2020). Bundesbank online pilot survey on consumer expectations. Deutsche Bundesbank Technical Paper 01/2020.
- Bräuer, K. (2021). Inside the personality of retail investors during the COVID-19 pandemic. mimeo.
- Bräuer, K., A. Hackethal, and T. Hanspal (2022). Consuming dividends. *Review of Financial Studies*, in press.


- Broda, C. and J. Parker (2014). The economic stimulus payments of 2008 and the aggregate demand for consumption. *Journal of Monetary Economics* 68, 20–36.
- Bunn, P., J. Le Roux, K. Reinold, and P. Surico (2017). The consumption response to positive and negative income shocks. *Journal of Monetary Economics* (96), 1–15.
- Carroll, C. (1997). Buffer-stock saving and the life cycle/permanent income hypothesis. *The Quarterly Journal of Economics* 112(1), 1–55.
- Chambers, V. and M. Spencer (2008). Does changing the timing of a yearly individual tax refund change the amount spent vs. saved? *Journal of Economic Psychology* 29(6), 856–862.
- Christelis, D., D. Georgarakos, T. Jappelli, L. Pistaferri, and M. van Rooij (2019). Asymmetric consumption effects of transitory income shocks. *Economic Journal* 129(622), 2322–2341.
- Coibion, O., Y. Gorodnichenko, and M. Weber (2020). How did U.S. consumers use their stimulus payments? NBER Working Paper 27693.
- Davis, J. and S. Heller (2020). Rethinking the benefits of youth employment programs: The heterogeneous effects of summer jobs. *Review of Economics and Statistics* 102(4), 664–677.
- Deaton, A. (1991). Saving and liquidity constraints. *Econometrica* 59(5), 1221–1248.
- Drescher, K., P. Fessler, and P. Lindner (2020). Helicopter money in Europe: New evidence on the marginal propensity to consume across European households. *Economics Letters* 195, 109416.
- Evers, E., A. Imas, and C. Kang (2021). On the role of similarity in mental accounting and hedonic editing. *Psychological Review*, in press.
- Fagereng, A., M. Holm, and G. Natvik (2021). MPC heterogeneity and household balance sheets. *American Economic Journal: Macroeconomics* 13(4), 1–54.

- Farbmacher, H., H. Kögel, and M. Spindler (2021). Heterogeneous effects of poverty on attention. *Labour Economics* 71, 102028.
- Friedman, J. (2001). Greedy function approximation: A gradient boosting machine. *The Annals of Statistics* 29(5), 1189–1232.
- Fuster, A., G. Kaplan, and B. Zafar (2021). What would you do with \$500? Spending responses to gains, losses, news, and loans. *The Review of Economic Studies* 88(4), 1760–1795.
- Hastings, J. and J. Shapiro (2013). Fungibility and consumer choice: Evidence from commodity price shocks. *The Quarterly Journal of Economics* 128(4), 1449–1498.
- Hastings, J. and J. Shapiro (2018). How are SNAP benefits spent? Evidence from a retail panel. *American Economic Review* 108(12), 3493–3540.
- Heath, C. and J. Soll (1996). Mental budgeting and consumer decisions. *Journal of Consumer Research* 23(1), 40–52.
- Henderson, P. and R. Peterson (1992). Mental accounting and categorization. *Organizational Behavior and Human Decision Processes* 51(1), 92–117.
- Jappelli, T. and L. Pistaferri (2014). Fiscal policy and MPC heterogeneity. *American Economic Journal: Macroeconomics* 6(4), 107–136.
- Jappelli, T. and L. Pistaferri (2020). Reported MPC and unobserved heterogeneity. *American Economic Journal: Economic Policy* 12(4), 275–297.
- Kaplan, G. and G. Violante (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica* 82(4), 1199–1239.
- Kaplan, G., G. Violante, and J. Weidner (2014). The wealthy hand-to-mouth. *Brookings Papers on Economic Activity* (1), 77–138.
- Karlsson, N., T. Garling, and M. Selart (1997). Effects of mental accounting on intertemporal choice. *Psychological Research* 27(5).


- Knittel, C. and S. Stolper (2019). Using machine learning to target treatment: The case of household energy use. NBER Working Paper 26531.
- Kooreman, P. (2000). The labeling effect of a child benefit system. *The American Economic Review* 90(3), 571–583.
- Kueng, L. (2018). Excess sensitivity of high-income consumers. *The Quarterly Journal of Economics* 133(4), 1693–1751.
- Liu, Q., Q. Shen, Z. Li, and S. Chen (2021). Stimulating consumption at low budget: Evidence from a large-scale policy experiment amid the COVID-19 pandemic. *Management Science* 67(12), 7291–7307.
- Medina, P. and M. Pagel (2021). Does saving cause borrowing? NBER Working Paper 28956.
- Milkman, K. and J. Beshears (2009). Mental accounting and small windfalls: Evidence from an online grocer. *Journal of Economic Behavior & Organization* 71(2), 384–394.
- Parker, J. (2017). Why don't households smooth consumption? Evidence from a \$25 million experiment. *American Economic Journal: Macroeconomics* 9(4), 153–183.
- Parker, J., J. Schild, L. Erhard, and D. Johnson (2022). Household spending responses to the economic impact payments of 2020: Evidence from the Consumer Expenditure Survey. NBER Working Paper 29648.
- Parker, J. and N. Souleles (2019). Reported effects versus revealed-preference estimates: Evidence from the propensity to spend tax rebates. *American Economic Review: Insights* 1(3), 273–290.
- Parker, J., N. Souleles, D. Johnson, and R. McClelland (2013). Consumer spending and the economic stimulus payments of 2008. *The American Economic Review* 103(6), 2530–2553.

- Sahm, C., M. Shapiro, and J. Slemrod (2010). Household response to the 2008 tax rebate: Survey evidence and aggregate implications. *Tax Policy and the Economy* 24(1), 69–110.
- Shapiro, M. and J. Slemrod (1995). Consumer response to the timing of income: Evidence from a change in tax withholding. *The American Economic Review* 85(1), 274–283.
- Shapiro, M. and J. Slemrod (2003). Consumer response to tax rebates. *American Economic Review* 93(1), 381–396.
- Shefrin, H. and R. Thaler (1988). The behavioral life-cycle hypothesis. *Economic Inquiry* 26(4), 609–643.
- Thaler, R. (1999). Mental accounting matters. *Journal of Behavioral Decision Making* 12(3), 183–206.
- Zhang, C. and A. Sussman (2018). Perspectives on mental accounting: An exploration of budgeting and investing. *Financial Planning Review* 1(1-2), e1011.

Tables and Figures



(a) Shock size scenarios



(b) Payment mode scenarios

Figure 1: Distribution of MPCs for different income shock scenarios

Data source: BOP-HH, December 2020 and June 2021.

Notes: This figure plots the distribution of the MPC out of different income shock scenarios. Panel A shows the MPC distribution for the one month's income shock scenario and the three months' income shock scenario. Panel B shows the MPC distribution for the baseline scenario, the cash payment scenario, and the savings account scenario, respectively.

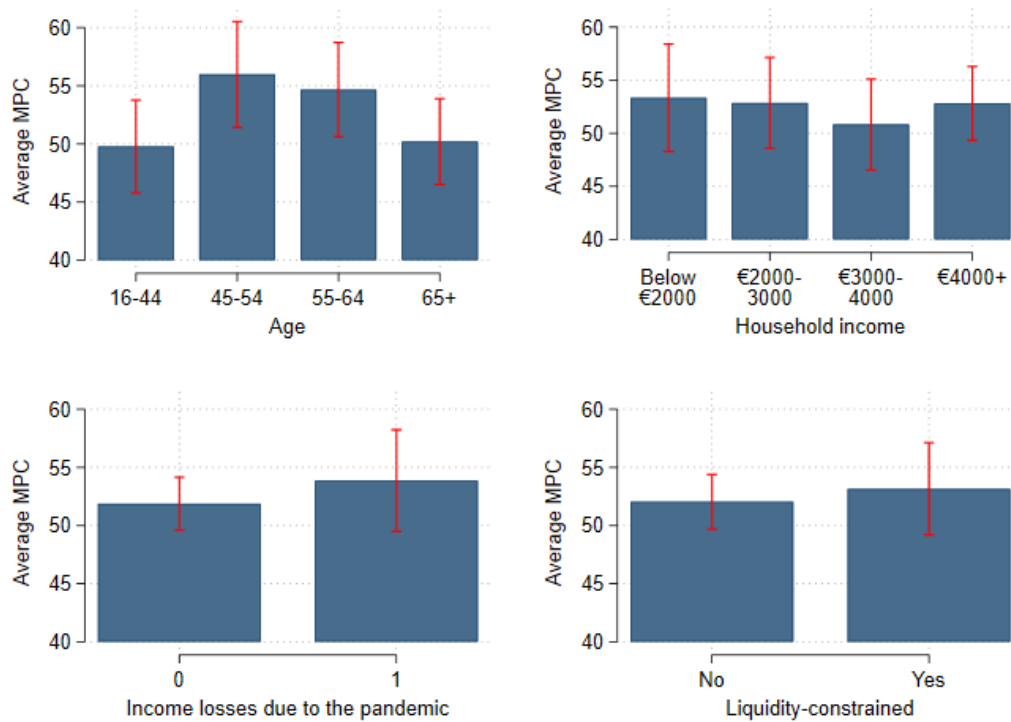




Figure 2: Average MPC out of one month's income shock by covariates

Data source: BOP-HH, December 2020 and June 2021.

Notes: This figure plots the average MPC out of the one month's income shock baseline scenario (BASE-1) by respondents' age, household income, income losses due to the COVID-19 pandemic, and being liquidity-constrained. The latter is defined as lacking a liquidity buffer of at least three months of income.



(a) Shock size effect




(b) Payment mode effect


Figure 3: Distribution of predicted treatment effects

Data source: BOP-HH, December 2020 and June 2021.

Notes: This figure shows how the predicted treatment effects of different income shock scenarios are distributed across individuals. Panel A plots the individualized predictions of the shock size effect, showing the extent to which individuals adjust their MPC in response to an increase in the shock size from one month to three months of income. Panel B plots the individualized predictions of the payment mode effect, showing how much individuals adjust their MPC in response to changes in the payment mode, i.e. windfall paid out in cash or without any specification of the payment mode compared to a windfall deposited in an instant-access savings account. Each distribution is a kernel density of the predicted treatment effect, coming from a separate causal forest.



(a) Shock size effect



(b) Payment mode effect

Figure 4: Bottom vs. top quartile of predicted treatment effects

Data source: BOP-HH, December 2020 and June 2021.

Notes: This figure compares observable characteristics between those adjusting their MPC most strongly in response to changes in the windfall's shock size and payment mode, respectively, and those who hardly react. Panel A (Panel B) reports differences in means of characteristics between respondents located in the bottom and the top quartile of the shock size effect (payment mode effect) distribution. Units are standard deviations of the respective characteristic. Error bars (in red) indicate 95% confidence intervals. Behavioral and household balance sheet characteristics are elicited with the following statements to which respondents can disagree or agree with on a 7-point Likert scale: impatience – "I am generally a very patient person" (inverted scale); act without thinking – "I rarely do anything without thinking about it thoroughly" (inverted scale); myopic decision-making – "I live in the here and now and don't really think about the future"; commitment to goals – "I actively follow through with the plans I make." (inverted scale); financial planning – "I plan major spending and investment decisions more than one year ahead"; low mathematical skills – "I have a lot of confidence in my mathematical skills" (inverted scale); liquidity-constrained – "I have put aside money for a possible emergency so that I can cover expenses for at least three months with no income" (inverted scale).

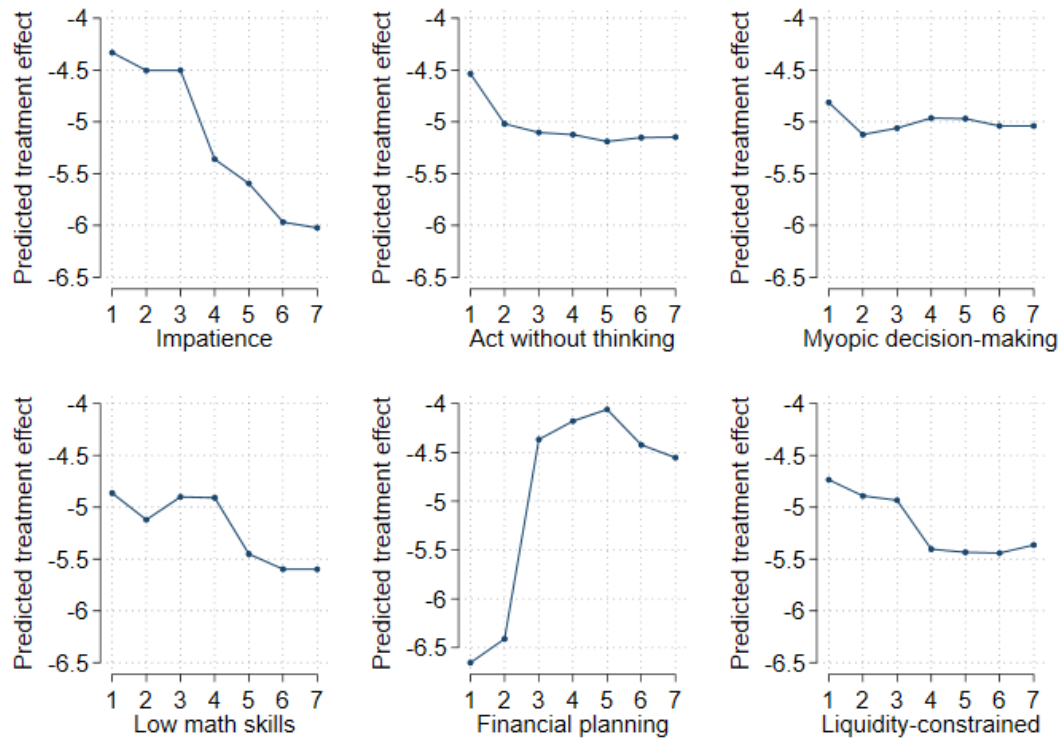


Figure 5: Partial dependence plots for the shock size effect

Data source: BOP-HH, December 2020 and June 2021.

Notes: This figure shows how the effect of increasing the shock size on the MPC varies with respondents' characteristics. Each panel depicts a partial dependence plot, indicating how on average the predicted treatment effect varies when one characteristic is changed at a time, while holding all other characteristics fixed. To generate a plot, I first use the causal forest to predict treatment effects at a particular value of a variable of interest by forcing each observation of the sample to assume that particular value, while holding all other variables constant at their empirical value. I then average over the sample distribution at each value of that variable. Behavioral and household balance sheet characteristics are elicited with an item battery listing the following statements to which respondents can disagree or agree with on a 7-point Likert scale: impatience – "I am generally a very patient person" (inverted scale); act without thinking – "I rarely do anything without thinking about it thoroughly" (inverted scale); myopic decision-making – "I live in the here and now and don't really think about the future"; financial planning – "I plan major spending and investment decisions more than one year ahead"; low mathematical skills – "I have a lot of confidence in my mathematical skills" (inverted scale); liquidity-constrained – "I have put aside money for a possible emergency so that I can cover expenses for at least three months with no income" (inverted scale).




Figure 6: Partial dependence plots for the payment mode effect

Data source: BOP-HH, December 2020 and June 2021.

Notes: This figure shows how the payment mode effect on the MPC varies with respondents' characteristics. Each panel depicts a partial dependence plot, indicating how on average the predicted payment mode effect varies when one characteristic is changed at a time, while holding all other characteristics fixed. To generate a plot, I first use the causal forest to predict treatment effects at a particular value of a variable of interest by forcing each observation of the sample to assume that particular value, while holding all other variables constant at their empirical value. I then average over the sample distribution at each value of that variable. Behavioral and household balance sheet characteristics are elicited with an item battery listing the following statements to which respondents can disagree or agree with on a 7-point Likert scale: impatience – "I am generally a very patient person" (inverted scale); act without thinking – "I rarely do anything without thinking about it thoroughly" (inverted scale); myopic decision-making – "I live in the here and now and don't really think about the future"; financial planning – "I plan major spending and investment decisions more than one year ahead"; low mathematical skills – "I have a lot of confidence in my mathematical skills" (inverted scale); liquidity-constrained – "I have put aside money for a possible emergency so that I can cover expenses for at least three months with no income" (inverted scale).




Figure 7: Variance explained by liquidity constraints and behavioral characteristics

Data source: BOP-HH, December 2020 and June 2021.

Notes: This figure compares the relative importance of liquidity constraints and behavioral characteristics for heterogeneity in the shock size effect. For each characteristic the semi-partial R^2 is reported, indicating the unique variance in the shock size effect explained by the specific characteristic. The semi-partial R^2 s are retrieved from a linear projection of liquidity constraints and behavioral characteristics onto the predicted shock size effect. Behavioral and household balance sheet characteristics are elicited with an item battery listing the following statements to which respondents can disagree or agree with on a 7-point Likert scale: impatience – "I am generally a very patient person" (inverted scale); act without thinking – "I rarely do anything without thinking about it thoroughly" (inverted scale); myopic decision-making – "I live in the here and now and don't really think about the future"; financial planning – "I plan major spending and investment decisions more than one year ahead"; low mathematical skills – "I have a lot of confidence in my mathematical skills" (inverted scale); liquidity-constrained – "I have put aside money for a possible emergency so that I can cover expenses for at least three months with no income" (inverted scale).

Table 1: Descriptive statistics

	Mean	SD	Median	Min	Max	Obs.
Age	55.20	15.26	56.00	16.00	80.00	6427
Secondary education	0.46	0.50	0.00	0.00	1.00	6425
Bachelor	0.26	0.44	0.00	0.00	1.00	6425
Master/PhD	0.28	0.45	0.00	0.00	1.00	6425
Female	0.41	0.49	0.00	0.00	1.00	6427
Unemployed	0.02	0.13	0.00	0.00	1.00	6423
Household size	2.20	0.96	2.00	1.00	4.00	6411
City size 100k+	0.33	0.47	0.00	0.00	1.00	6427
Eastern Germany	0.18	0.39	0.00	0.00	1.00	6427
Homeowner	0.65	0.48	1.00	0.00	1.00	6418
HH income <€2000	0.14	0.35	0.00	0.00	1.00	6098
HH income €2000-3000	0.22	0.42	0.00	0.00	1.00	6098
HH income €3000-4000	0.23	0.42	0.00	0.00	1.00	6098
HH income €4000+	0.40	0.49	0.00	0.00	1.00	6098
Income losses due to the pandemic	0.20	0.40	0.00	0.00	1.00	6414
Liquidity-constrained	2.70	2.15	2.00	1.00	7.00	6409
Long-term financial planning	3.67	1.82	4.00	1.00	7.00	6416
Low math skills	2.99	1.59	3.00	1.00	7.00	6420
Myopic decision-making	2.63	1.54	2.00	1.00	7.00	6422
Commitment to goals	4.94	1.29	5.00	1.00	7.00	6417
Impatience	3.40	1.62	3.00	1.00	7.00	6421
Act without thinking	3.12	1.66	3.00	1.00	7.00	6420

Data source: BOP-HH, December 2020 and June 2021.

Notes: This table reports summary statistics of variables used in the analysis.

Table 2: Summary statistics on MPCs across income shock scenarios

	Payment mode:					
	Baseline		Cash payment		Savings account	
	Average MPC	N	Average MPC	N	Average MPC	N
One month's income shock	52.32 (33.64)	1067	53.55 (33.75)	1069	44.94 (34.98)	1062
Three months' income shock	46.61 (30.85)	1054	47.33 (31.53)	1059	42.23 (32.36)	1062

Data source: BOP-HH, December 2020 and June 2021.

Notes: This table shows summary statistics on MPCs across the six different income shock scenarios. Each cell reports the average MPC, the corresponding standard deviation (in parenthesis), and the number of respondents (N) that are exposed to a specific income shock scenario, with the shock size indicated in the first column and the payment mode indicated in the super header.

Table 3: Average treatment effects of shock size and payment mode scenarios

	MPC	
	(1)	(2)
<i>Shock size effect (omitted category: one month's income shock)</i>		
Three months' income shock	-4.88*** (0.82)	-5.47*** (0.84)
<i>Payment mode effect (omitted category: savings account scenario)</i>		
Baseline scenario	5.90*** (1.01)	5.55*** (1.04)
Cash payment scenario	6.86*** (1.01)	6.42*** (1.04)
p-value $\beta_{\text{Baseline}} = \beta_{\text{Cash payment}}$	0.33	0.40
Sociodemographics	No	Yes
Wave FE	Yes	Yes
R^2	0.02	0.03
Observations	6373	6051
Mean for BASE-1 scenario	52.32	52.46

Data source: BOP-HH, December 2020 and June 2021.

Notes: This table reports average treatment effects (ATEs) estimated with the following OLS regression equation: $Y_i = \alpha + \gamma \times \text{LargeShock} + \sum_{s=1}^S \beta_s \times \text{PaymentMode}_s + \delta \mathbf{X}_i + \epsilon_i$, where Y_i is the MPC of respondent i . *LargeShock* is an indicator variable equal to 1 if respondent i received a three months' income shock and 0 if respondent i received a one month's income shock. The γ coefficient identifies the ATE of increasing the shock size from one month's to three months' of income. *PaymentMode_s* is an indicator variable that equals 1 if respondent i receives an income shock paid out in cash and without any specification of the payment mode, respectively, and 0 otherwise. The β_s coefficients identify the ATEs of the different payment modes. The second column uses the same specification as in the first column but augmented with respondent-specific controls \mathbf{X}_i . The last row reports the mean of the outcome variable for the group of respondents receiving a one month's income without any specification of the payment mode. Robust standard errors are reported in parentheses. *, **, *** indicate significance at the 10%, 5%, and 1% level, respectively.

Table 4: Extensive and intensive margins

	MPC		I(MPC >0)		MPC MPC >0	
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Shock size effect (omitted category: one month's income shock)</i>						
Three months' income shock	-4.88*** (0.82)	-5.47*** (0.84)	0.02*** (0.01)	0.02** (0.01)	-7.10*** (0.78)	-7.57*** (0.79)
<i>Payment mode effect (omitted category: savings account scenario)</i>						
Baseline scenario	5.90*** (1.01)	5.55*** (1.04)	0.06*** (0.01)	0.06*** (0.01)	3.16*** (0.96)	2.89*** (0.98)
Cash payment scenario	6.86*** (1.01)	6.42*** (1.04)	0.07*** (0.01)	0.07*** (0.01)	3.71*** (0.97)	3.28*** (0.99)
Sociodemographics	No	Yes	No	Yes	No	Yes
Wave FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	6373	6051	6373	6051	5489	5208

Data source: BOP-HH, December 2020 and June 2021.

Notes: This table decomposes the average treatment effects on the MPC into an extensive and an intensive margin. Columns (1) and (2) reproduce results from [Table 3](#), reporting the estimated effect of varying the shock size and payment mode of the income shock on the MPC. Columns (3) and (4) study treatment effects on the extensive margin. Marginal effects from logistic regressions are shown using an indicator variable for whether respondents reported they would spend any money as the dependent variable. Columns (5) and (6) study treatment effects on the intensive margin, defined as the average amount spent conditional on any spending. Coefficients from OLS regressions are shown. Even columns use the same specification as the odd columns but augmented with respondent-specific controls. Robust standard errors are reported in parentheses. *, **, *** indicate significance at the 10%, 5%, and 1% level, respectively.

Table 5: Average treatment effect of income source and payment mode scenarios

	MPC	
	(1)	(2)
<i>Effect of income source (omitted category: lottery win)</i>		
Payment from the government	1.52 (1.37)	1.89 (1.45)
<i>Payment mode effect (omitted category: baseline scenario)</i>		
Cash payment	3.10** (1.37)	3.52** (1.45)
Sociodemographics	No	Yes
R^2	0.00	0.04
Observations	2253	1970
Mean for BASE-1 scenario	41.91	41.79

Data source: Retail bank survey, October 2020.

Notes: This table reports regression results for the effect of varying the source of income shocks and the payment mode of income shocks on the MPC. Average treatment effects (ATEs) are estimated with the following OLS regression equation: $Y_i = \alpha + \beta \times Gov_i + \gamma \times Cash_i + \delta X_i + \epsilon_i$, where Y_i is the MPC of respondent i . Gov_i is an indicator variable equal to 1 if respondent i received a one-off payment from the government and 0 if she received a lottery win. $Cash$ is an indicator variable equal to 1 if respondent i received a transfer paid out in cash and 0 otherwise. The second column uses the same specification as the first column but augmented with respondent-specific controls X_i . The last row reports the mean of the outcome variable for the group of respondents receiving a one month's income shock from the government without any specification of the payment mode. Robust standard errors are reported in parentheses. *, **, *** indicate significance at the 10%, 5%, and 1% level, respectively.

Online Appendix

A Survey Questions

This Appendix provides the survey questions from the Bundesbank Online Panel Households and the retail bank survey that are used in this paper.

A.1 Bundesbank Online Panel Households

The following section lists the survey questions from the BOP-HH December 2020 and June 2021 waves.

A.1.1 MPC Questions

The entire sample is split randomly into six equally-sized groups.

1. **BASE-1:** MPC out of one month's income payment

Imagine that you unexpectedly receive a one-off payment from the government, with no repayment obligation, in the amount of your monthly net household income. What part of this would you spend in the next twelve months and what part would you save or use to repay loans?

Please state the percentage that you would spend. Please respond with a value between 0 and 100. "0" means that you save the entire one-off payment or use it to repay loans; "100" means that you would spend the entire one-off payment over the next twelve months. Use the values 1 to 99 to scale your response.

[Input field]

2. **CASH-1:** MPC out of one month's income payment in cash

Imagine that you unexpectedly receive a one-off payment from the government, with no repayment obligation, in the amount of your monthly net household income – and this one-off payment is paid out in cash. What part of this would you spend in the next twelve months and what part would you save or use to repay loans?

Please state the percentage that you would spend. Please respond with a value between 0 and 100. "0" means that you save the entire one-off payment or use it to repay loans; "100" means that you would spend the entire one-off payment over the next twelve months. Use the values 1 to 99 to scale your response.

[Input field]

3. **SAVINGS-ACCOUNT-1:** MPC out of one month's income payment deposited in an savings account

December 2020 wave:

Imagine that you unexpectedly receive a one-off payment from the government, with no repayment obligation, in the amount of your monthly net household income – and this one-off payment is deposited in an instant-access savings account that pays interest. What part of this would you spend in the next twelve months and what part would you save or use to repay loans?

Please state the percentage that you would spend. Please respond with a value between 0 and 100. "0" means that you save the entire one-off payment or use it to repay loans; "100" means that you would spend the entire one-off payment over the next twelve months. Use the values 1 to 99 to scale your response.

[Input field]

June 2021 wave:

Imagine that you unexpectedly receive a one-off payment from the government, with no repayment obligation, in the amount of your monthly net household income – and this one-off payment is paid into a savings account that is accessible at any time. What part of this would you spend in the next twelve months and what part would you save or use to repay loans?

Please state the percentage that you would spend. Please respond with a value between 0 and 100. "0" means that you save the entire one-off payment or use

it to repay loans; “100” means that you would spend the entire one-off payment over the next twelve months. Use the values 1 to 99 to scale your response.

[Input field]

4. **BASE-3:** MPC out of three months’ income payment

Imagine that you unexpectedly receive a one-off payment from the government, with no repayment obligation, in the amount of three months’ net household income. What part of this would you spend in the next twelve months and what part would you save or use to repay loans?

Please state the percentage that you would spend. Please respond with a value between 0 and 100. “0” means that you save the entire one-off payment or use it to repay loans; “100” means that you would spend the entire one-off payment over the next twelve months. Use the values 1 to 99 to scale your response.

[Input field]

5. **CASH-3:** MPC out of three months’ income payment in cash

Imagine that you unexpectedly receive a one-off payment from the government, with no repayment obligation, in the amount of three months’ net household income – and this one-off payment is paid out in cash. What part of this would you spend in the next twelve months and what part would you save or use to repay loans?

Please state the percentage that you would spend. Please respond with a value between 0 and 100. “0” means that you save the entire one-off payment or use it to repay loans; “100” means that you would spend the entire one-off payment over the next twelve months. Use the values 1 to 99 to scale your response.

[Input field]

6. **SAVINGS-ACCOUNT-3:** MPC out of three months’ income payment deposited in an savings account

December 2020 wave:

Imagine that you unexpectedly receive a one-off payment from the government, with no repayment obligation, in the amount of three months' net household income – and this one-off payment is deposited in an instant-access savings account that pays interest. What part of this would you spend in the next twelve months and what part would you save or use to repay loans?

Please state the percentage that you would spend. Please respond with a value between 0 and 100. "0" means that you save the entire one-off payment or use it to repay loans; "100" means that you would spend the entire one-off payment over the next twelve months. Use the values 1 to 99 to scale your response.

[Input field]

June 2021 wave:

Imagine that you unexpectedly receive a one-off payment from the government, with no repayment obligation, in the amount of three months' net household income – and this one-off payment is paid into a savings account that is accessible at any time. What part of this would you spend in the next twelve months and what part would you save or use to repay loans?

Please state the percentage that you would spend. Please respond with a value between 0 and 100. "0" means that you save the entire one-off payment or use it to repay loans; "100" means that you would spend the entire one-off payment over the next twelve months. Use the values 1 to 99 to scale your response.

[Input field]

A.1.2 Spending Distribution Question

Respondents reporting a positive MPC were asked the follow-up survey question shown below:

How would you split the amount you would like to spend between the following

categories? Please click and drag the slide bar to enter your response. Each answer must be a value between 0 and 100. The sum of all responses must equal 100.

- a General cost of living [Input field]
- b Holidays, travel, restaurant visits or other entertainment/recreation [Input field]
- c Donations/gifts [Input field]
- d Major purchases worth € 1,000 or less (e.g. electrical devices, sports equipment, clothing) [Input field]
- e Larger purchases worth more than € 1,000 (e.g. cars) [Input field]
- f Own training/training for members of the household [Input field]
- g Repairs/renovations (e.g. of vehicles or property) [Input field]
- h Other [Input field]

A.1.3 Reasons for Spending/Saving

Respondents of the June 2021 wave were asked the follow-up survey question shown below:

You stated that you would spend [Value of MPC] % of a one-off payment. What are the main reasons for your decision? Please select all answers that apply.

- a *It would be too much effort to transfer the money from my savings account to my current account.* [only if **SAVINGS-ACCOUNT-1**| **SAVINGS-ACCOUNT-3**]
- b *It would be too much effort to deposit the money in a bank account.*
[only if **CASH-1**| **CASH-3**]
- c *I never spend money from my savings account.*
[only if **SAVINGS-ACCOUNT-1**| **SAVINGS-ACCOUNT-3**]
- d *I do not like spending money from my savings account.*
[only if **SAVINGS-ACCOUNT-1**| **SAVINGS-ACCOUNT-3**]
- e *I simply enjoy spending money.*
- f *I would rather pay off my debts than spend the money.*

- g *There is nothing specific I would like to buy beyond my usual purchases.*
- h *The one-off payment is not enough to make a major purchase.*
- i *I would save the money for a potential tax increase.*
- j *Other reasons:* [Input field]

A.1.4 Behavioral and Household Balance Sheet Characteristics

Cognitive sophistication, self-control problems, and liquidity constraints are elicited with the following item battery.

We would now like to ask you to evaluate a number of different statements. To what extent do you agree or disagree with the following statements? Please select one answer for each row.

1 = Strongly disagree, 2 -> 6 [no label], 7 = Strongly agree

- a *I plan major spending and investment decisions more than one year in advance.*
- b *I have a lot of confidence in my mathematical skills.*
- c *I live in the here and now and do not really think about the future.*
- d *I actively follow through with the plans that I make.*
- e *I am generally a very patient person.*
- f *I rarely do anything without thinking about it thoroughly.*
- g *I have put aside money for a possible emergency so that I can cover expenses for at least three months with no income.*

A.1.5 Sociodemographic Characteristics

- Education School:

What is your highest level of educational attainment?

- a *Still at school*
- b *Completed lower secondary school*

- c Completed higher secondary school*
- d Polytechnical secondary school certificate, 10th grade*
- e University of applied sciences entrance diploma/completed technical school*
- f General or subject-specific university entrance diploma/senior school-leaving certificate (from a grammar school)/East German secondary school up to 12th grade (also with apprenticeship)*
- g Other school-leaving certificate*
- h No school-leaving certificate (and currently not a student)*

- Professional Education:

What level of vocational training or university degree do you have? Please think of your highest qualification here.

- a Currently in training or studying (bachelor's degree not yet completed)*
- b Completed vocational training (apprenticeship)*
- c Completed vocational training (vocational school or commercial college)*
- d Completed training at a technical or commercial college, school for master craftsmen or engineers or university of co-operative education with shorter preparation time (up to 880 hours)*
- e Completed training at a university of cooperative education with longer preparation time (more than 880 hours)*
- f Bachelor's degree, applied sciences degree, completed training at an engineering college*
- g Diploma or master's degree, completed teacher training course*
- h Doctorate/postdoctoral qualification obtained*
- i Other professional qualification*
- j No vocational training completed (and currently not in training/education)*

- Household income:

What is the total monthly net income of your household? This refers to the total amount, comprising wages, salaries, income from self-employment and pen-

sions, in each case after deducting tax and social security contributions. In this amount, please include any income received through public aid, earnings from rents and leases, housing allowance, child benefits and any other sources of income.

1. Less than €500
2. €500 to €999
3. €1,000 to €1,499
4. €1,500 to €1,999
5. €2,000 to €2,499
6. €2,500 to €2,999
7. €3,000 to €3,499
8. €3,500 to €3,999
9. €4,000 to €4,999
10. €5,000 to €5,999
11. €6,000 to €7,999
12. €8,000 to €9,999
13. €10,000 or more

- Homeowner:

Does your household live in a rented property or an owner-occupied apartment or house? Please select one answer.

- a Rent and do not own any other home(s)*
- b Rent but own other home(s)*
- c Live in own apartment*
- d Live in own house*

- Income losses:

Have you experienced a loss of income or other financial losses (e.g. price losses, etc.) in connection with the coronavirus pandemic? Please select all answers that apply.

- a Yes, lost wage or salary income*
- b Yes, other lost income, e.g. from rents and leases*
- c Yes, other financial losses, e.g. price losses, equities, funds, securities*
- d No, neither income losses nor any other financial losses*

- Employment Status:

Which of the following best describes your employment status?

- a In full-time employment (including apprenticeship)*
- b In part-time employment (including phased retirement)*
- c In casual or irregular employment (including paid internship and integration measures)*
- d On maternity leave/parental leave/longer-term sick leave/other leave, planning to return to work*
- e Unemployed (officially registered)*
- f At school, university or in an unpaid internship*
- g Retiree or pensioner*
- h Retired early or about to retire (including unfit for work or reduced ability to work)*
- i Federal volunteer service/voluntary year*
- j Homemaker*
- k Other form of non-employment*

- Household size:

How many persons live permanently in your household, including yourself?

Please also consider all children living in your household. [Input field]

- Information on age, gender, region, and city size were automatically recorded by Forsa.

A.1.6 Feedback Questions

- Questionnaire interest:

We would now like you to answer a few questions regarding your assessment of the survey. How interesting did you find the survey overall? Please select one answer.

- 1. Very interesting*
- 2. Interesting*
- 3. Partly interesting/partly uninteresting*
- 4. Not so interesting*
- 5. Not interesting at all*

- Questionnaire difficult:

How easy or difficult was it to answer the questions? Please select one answer.

- 1. Very difficult*
- 2. Somewhat difficult*
- 3. Partly difficult/partly easy*
- 4. Somewhat easy*
- 5. Very easy*

- Questionnaire length:

How did you find the length of the survey? Please select one answer.

- 1. Far too long*
- 2. Somewhat too long*
- 3. Just right*
- 4. Somewhat too short*
- 5. Far too short*

A.1.7 Inflation Expectations

The variable measuring the accuracy of quantitative inflation expectations is based on two survey questions.

- Expectation inflation or deflation:

Do you think inflation or deflation is more likely over the next twelve months?

Note: Inflation is the percentage increase in the general price level. It is mostly measured using the consumer price index. A decrease in the price level is generally described as “deflation”. Please select one answer.

- 1. Inflation more likely*
- 2. Deflation more likely*

- Inflation expectations quantitative:

What do you think the rate of inflation/deflation in Germany will roughly be over the next twelve months? Note: Inflation is the percentage increase in the general price level. It is mostly measured using the consumer price index. A decrease in the price level is generally described as “deflation”. Please enter a value in the input field (values may have one decimal place).

[Input field] %

A.2 Online Survey of Retail Customers at a Large German Bank

The following section lists the survey questions from the October 2020 wave.

A.2.1 MPC Module, October 2020

The entire sample is split randomly into four equally-sized groups.

1. **GOV-BASE:** MPC out of one-off payment from the government

Imagine that you unexpectedly receive a one-off payment from the government, with no repayment obligation, in the amount of your monthly net household income. What part of this would you spend in the next twelve months and what part would you save or use to repay loans?

Please state the percentage that you would spend. Please respond with a value between 0 and 100. “0” means that you save the entire one-off payment or use it to repay loans; “100” means that you would spend the entire one-off payment over the next twelve months. Use the values 1 to 99 to scale your response.

[Input field]

2. **GOV-CASH:** MPC out of one-off payment from the government paid out in cash

Imagine that you unexpectedly receive a one-off payment from the government, with no repayment obligation, in the amount of your monthly net household income. Suppose this one-off payment is paid out in cash. What part of this would you spend in the next twelve months and what part would you save or use to repay loans?

Please state the percentage that you would spend. Please respond with a value between 0 and 100. "0" means that you save the entire one-off payment or use it to repay loans; "100" means that you would spend the entire one-off payment over the next twelve months. Use the values 1 to 99 to scale your response.

[Input field]

3. **LOTTERY-BASE:** MPC out of unexpected lottery prize

Imagine you unexpectedly receive money from a lottery, equal to the amount of net income your household receives in a month. What part of this would you spend in the next twelve months and what part would you save or use to repay loans?

Please state the percentage that you would spend. Please respond with a value between 0 and 100. "0" means that you save the entire one-off payment or use it to repay loans; "100" means that you would spend the entire one-off payment over the next twelve months. Use the values 1 to 99 to scale your response.

[Input field]

4. **LOTTERY-CASH:** MPC out of unexpected lottery prize paid out in cash

Imagine you unexpectedly receive money from a lottery, equal to the amount of net income your household receives in a month. Suppose the prize money is paid out in cash. What part of this would you spend in the next twelve months and what part would you save or use to repay loans?

Please state the percentage that you would spend. Please respond with a value between 0 and 100. "0" means that you save the entire one-off payment or use it to repay loans; "100" means that you would spend the entire one-off payment over the next twelve months. Use the values 1 to 99 to scale your response.

[Input field]

A.2.2 MPC Module, September 2019

The entire sample is split randomly into six equally-sized groups. All respondents are asked the LOTTERY-BASE-1 question. Two groups each are asked the three months' income question, i.e. either the LOTTERY-BASE-3 scenario, or the LOTTERY-CASH-3 scenario, or the LOTTERY-SAVINGS-3 scenario. The order of the MPC questions is randomized such that half of the sample is asked first about their consumption response to a one month's income shock and the other half about their response to a three months' shock.

1. **LOTTERY-BASE-1:** MPC out of unexpected lottery prize

Imagine you unexpectedly receive money from a lottery, equal to the amount of net income your household receives in a month. What part of this would you spend in the next twelve months and what part would you save or use to repay loans?

Please state the percentage that you would spend. Please respond with a value between 0 and 100. "0" means that you save the entire one-off payment or use it to repay loans; "100" means that you would spend the entire one-off payment over the next twelve months. Use the values 1 to 99 to scale your response.

[Input field]

2. **LOTTERY-BASE-3:** MPC out of unexpected lottery prize

Imagine you unexpectedly receive money from a lottery, in the amount of three months' net household income. What part of this would you spend in the next twelve months and what part would you save or use to repay loans?

Please state the percentage that you would spend. Please respond with a value between 0 and 100. "0" means that you save the entire one-off payment or use it to repay loans; "100" means that you would spend the entire one-off payment over the next twelve months. Use the values 1 to 99 to scale your response.

[Input field]

3. **LOTTERY-CASH-3:** MPC out of unexpected lottery prize paid out in cash

Imagine you unexpectedly receive money from a lottery, in the amount of three

months' net household income. Suppose the prize money is paid out in cash. What part of this would you spend in the next twelve months and what part would you save or use to repay loans?

Please state the percentage that you would spend. Please respond with a value between 0 and 100. "0" means that you save the entire one-off payment or use it to repay loans; "100" means that you would spend the entire one-off payment over the next twelve months. Use the values 1 to 99 to scale your response.

[Input field]

4. **LOTTERY-SAVINGS-3:** MPC out of unexpected lottery prize paid out in cash

Imagine you unexpectedly receive money from a lottery, in the amount of three months' net household income. Suppose the prize money is deposited in an instant-access savings account that pays interest. What part of this would you spend in the next twelve months and what part would you save or use to repay loans?

Please state the percentage that you would spend. Please respond with a value between 0 and 100. "0" means that you save the entire one-off payment or use it to repay loans; "100" means that you would spend the entire one-off payment over the next twelve months. Use the values 1 to 99 to scale your response.

[Input field]

A.2.3 Sociodemographic Characteristics

Respondents' sociodemographic characteristics were elicited in the initial survey wave using the following survey questions:

- Age:

What year were they born?

- Gender:

I am...

a male

b female

c diverse gender

- Education School/Professional:

What is your highest level of education?

- a No degree*
- b Secondary or high school diploma*
- c High school diploma*
- d Completed vocational training*
- e Master craftsman*
- f Completed studies*
- g Doctorate or higher*
- h No answer*
- i Other*

- Professional status:

Which of the following answer choices best describes your current professional status?

- a In school*
- b In training*
- c In college*
- d Employed*
- e Self-employed*
- f Self-employed, own company with employees*
- g Civil servant*
- h Housewife or househusband*
- i Job-seeking*
- j Not employed*
- k Permanently unable to work due to illness or disability*
- l Retired or taking early retirement*
- m No answer*
- n Other*

- Household size:

How many people live in your household?

- a 1*
- b 2*
- c 3*
- d 4 or more*

- **Household income:**

What do you estimate is the monthly disposable net income of your household, i.e. the money available to the entire household after deduction of taxes and social security contributions to cover expenses? Please consider the following types of income: wages, salary, income from self-employment, pension or annuity, income from public aid, income from renting/leasing, housing allowance, child benefit, other income.

- a under €500*
- b €500 to under €1,000*
- c €1,000 to under €1,500*
- d €1,500 to under €2,000*
- e €2,000 to under €2,500*
- f €2,500 to under €3,000*
- g €3,000 to under €3,500*
- h €3,500 to under €4,000*
- i €4,000 to under €4,500*
- j €4,500 to under €5,000*
- k €5,000 to under €7,500*
- l €7,500 to under €10,000*
- m €10,000 to under €15,000*
- n €15,000 and more*
- o Don't know*
- p No answer.*

- Homeowner:

Do you own real estate?

- a Yes, residential property*
- b Yes, rental property*
- c Yes, residential and rental property*
- d No*
- e No answer.*

- Region:

Have you lived or do you live in one of the new federal states?

- a Yes*
- b No*
- c No answer.*

How long did you live or work in the new federal states?

From [Input field] to [Input field]

B Causal Forest Algorithm

B.1 Intuition and Implementation

This appendix describes the intuition and implementation of the causal forest algorithm, which proceeds in two steps: a training step and a prediction step.⁵²

During training a number of trees are grown. Individual trees are trained as follows: First, a random training subsample is drawn by sampling a percentage p without replacement from the full sample of respondents N . Further, the training sample of size $n = p * N$ is randomly partitioned into a splitting subsample n_{sp} and an estimation subsample n_e .⁵³ The splitting subsample n_{sp} is split into disjoint groups of observations – so-called nodes. Using a single root node containing the entire subsample n_{sp} as a starting point, the data is recursively split into child nodes subject to a splitting criterion C_{sp} to form a decision tree.

The splitting criterion C_{sp} seeks splits that maximize heterogeneity in the treatment effect:⁵⁴ First it selects a random subset of variables X from all variables used in the causal forest. For each of the selected variables X it considers all candidate splits i.e. it considers all possible values s of a variable X_j to split a parent node into a left child node $L(j, s) = \{X|X_j \leq s\}$ and a right child node $R(j, s) = \{X|X_j > s\}$. From all candidate splits the split (X_j, s) is selected which maximizes the sum of squared predictions:

$$\max_{j,s} \left[\sum_{i: X_i \in L(j,s)} \hat{\tau}_{L(j,s)}^2 + \sum_{i: X_i \in R(j,s)} \hat{\tau}_{R(j,s)}^2 \right], \quad (3)$$

where $\hat{\tau}_{L(j,s)}^2$ and $\hat{\tau}_{R(j,s)}^2$ are the child node-specific ATEs. The splitting criterion C_{sp} thus rewards splits that increase the variance of treatment effects across nodes and penalizes splits that increase within-node variance. Having selected the optimal split,

⁵²The following description of the algorithm is based on [Athey et al. \(2019\)](#). For a more detailed discussion please refer to their paper.

⁵³By default the training sample is equally split at this stage.

⁵⁴Standard regression trees use a splitting rule to optimize for heterogeneity in predicted values \hat{Y} . In contrast, causal trees optimize for finding splits associated with treatment effect heterogeneity in subgroup ATEs.

all observations with values for the split variable X_j greater than the split value s go into the right child node, and all other observations into the left child node. If a node has no valid splits, or if splitting will not result in an improved fit, the node is not split further and forms a leaf of the final tree.⁵⁵ The splitting procedure stops when none of the nodes can be split any further.

Next, the estimation sample n_e is used to fill the tree's leaf nodes by pushing each observation down the tree and adding it to the leaf L in which it falls. Within each filled leaf L the treatment effect, $\hat{\tau}_L = \bar{y}_{Lt} - \bar{y}_{Lc}$, is estimated as the difference in average outcomes between treated t and control observation c within a leaf L . All training steps are repeated multiple times to grow many trees, which then form a causal forest.

In the prediction step the causal forest algorithm predicts an individualized average treatment effect (IATE) for each observation i of the full dataset: For each observation i the algorithm identifies all trees that did not use i during training and pushes i down each of these trees to determine in what leaf i falls. With this information, a weighted list of neighboring observations for each observation i is generated. The weights capture how many times a neighboring observation from the estimation sample n_e falls in the same leaf as the observation i across all identified trees. Lastly, the outcomes and treatment status of the weighted neighbor observations are used to predict the individualized treatment effect for observation i . Specifically, for a set of B trees indexed by $b = 1, \dots, B$, the weight $\alpha_i(x)$ is constructed as follows: $\alpha_i(x) = \frac{1}{B} \sum_{b=1}^B \frac{1(\{X_i \in L_b(x)\})}{|L_b(x)|}$, where $L_b(x)$ is the set of training observations falling in the same leaf as observation i . Thus the predicted treatment effect for each observation is a weighted average of predictions for that particular observation across all trees of the causal random forest.

Note that the algorithm uses one subsample n_{sp} to grow a tree, and another subsample n_e to estimate treatment effects within leaves. This sample-splitting technique, which [Athey et al. \(2019\)](#) refer to as honesty, reduces bias from over-fitting in tree predictions and allows the construction of valid confidence intervals of the estimates by establishing consistency and asymptotic normality.

⁵⁵To improve computational efficiency the causal forest of [Athey et al. \(2019\)](#) in the GRF package uses a gradient-based approximation of this splitting criterion C_t instead of checking each candidate split.

To estimate heterogeneous treatment effects, I use the causal forest algorithm from the GRF package version 1.2.0 in *R* provided by [Athey et al. \(2019\)](#). I grow two causal forests: one for the shock size effect and one for the payment mode effect. To increase statistical power, I pool the treatment groups from the cash payment and the baseline scenario for the analysis of the payment mode effect, i.e. I use a treatment dummy that equals 1 if respondent i receives a windfall either paid out in cash (CASH-1, CASH-3) or without any specification of the payment mode (BASE-1, BASE-3), and 0 if the respondent receives an income shock deposited in an instant-access savings account (SAVINGS-ACCOUNT-1, SAVINGS-ACCOUNT-3). Each causal forest comprises 10,000 trees.⁵⁶ To build the trees I use the algorithm's default sub-sampling rate (50%) and apply the *honest* approach. I set the minimum node size to 10 for the causal forest of the payment mode effect to improve its performance, as tuning for this causal forest suggests using a higher minimum node size. I use the default node size of 5 for the other causal forests.⁵⁷ For all other parameters I use the algorithm's default values. To train the causal forests I use the following variables: age, dummy variables for female, employment status, homeownership status, living in eastern Germany, having lost income due to the COVID-19 pandemic, a wave dummy, as well as categorical variables measuring educational attainment, household income, household size, city size, mathematical skills, financial planning, commitment to goals, impatience, acting without thinking, myopic decision-making, and liquidity constraints. To predict treatment effects for each observation in the sample I use only those trees that did not use the specific observation during training, i.e. where the observation was *out-of-bag*.

B.2 Accuracy of the Causal Forest

This appendix assesses whether the variation in predictions of the causal random forest reflects actual treatment heterogeneity in the data.

First, I follow [Davis and Heller \(2020\)](#) and study graphically how the causal forest

⁵⁶The algorithm's default is 2,000 trees. I follow the recommendations of [Athey et al. \(2019\)](#) and grow a larger forest to reduce excess error stemming from randomization inherent to the ensemble algorithm.

⁵⁷The minimum node size sets a target for the minimum number of observations in each tree leaf.

predictions are linked to actual treatment effects. For each treatment the x-axis in [Figure B.1](#) shows observations binned into twenty groups defined by percentile of the predicted treatment effect. The y-axis depicts OLS estimates of the treatment effect within each bin. The dashed line depicts a linear fit between actual and predicted effects. The 45-degree line (solid line) serves as a benchmark for the evaluation of the causal forest predictions, showing all points where predicted effects would equal actual treatment effects. [Figure B.1](#) illustrates that the causal forest for the shock size and the payment mode effect, respectively, performs quite well on average as the fitted line closely follows the 45-degree line.

Second, I apply a calibration test as described in [Athey and Wager \(2019\)](#) to formally assess the quality of the causal forest estimates. The test computes the best linear projection of the conditional average treatment effect on the forest prediction and the mean forest prediction. A coefficient of 1 for the *mean forest prediction* indicates that the mean forest prediction is correct. A coefficient of 1 for the *differential forest prediction* indicates that the heterogeneity estimates from the forest are well-calibrated. The p-value of the *differential forest prediction* coefficient serves as an omnibus test for the presence of treatment heterogeneity, i.e. the null hypothesis of no heterogeneity can be rejected if the coefficient is significantly greater than zero.

Panel A in [Table B.1](#) shows that the coefficients for the *mean forest prediction* are equal to 1, indication that it is correct for both causal forests. The *differential forest prediction* coefficient is significantly larger than 0 and close to 1 for the predicted shock size effect and the payment mode effect, respectively, indicating the presence of treatment heterogeneity and suggesting that both causal forests are well calibrated.

Lastly, motivated by [Davis and Heller \(2020\)](#) I test whether actual treatment effects differ between respondents with the lowest and highest predicted treatment effects. For each treatment I create indicators for each quartile of predictions. I then estimate separate local average treatment effects (LATEs) using OLS for each quartile and test the null hypothesis that the treatment coefficients for the bottom and top quartile are equal.

Panel B in [Table B.1](#) shows that the group with the most negative prediction (bot-

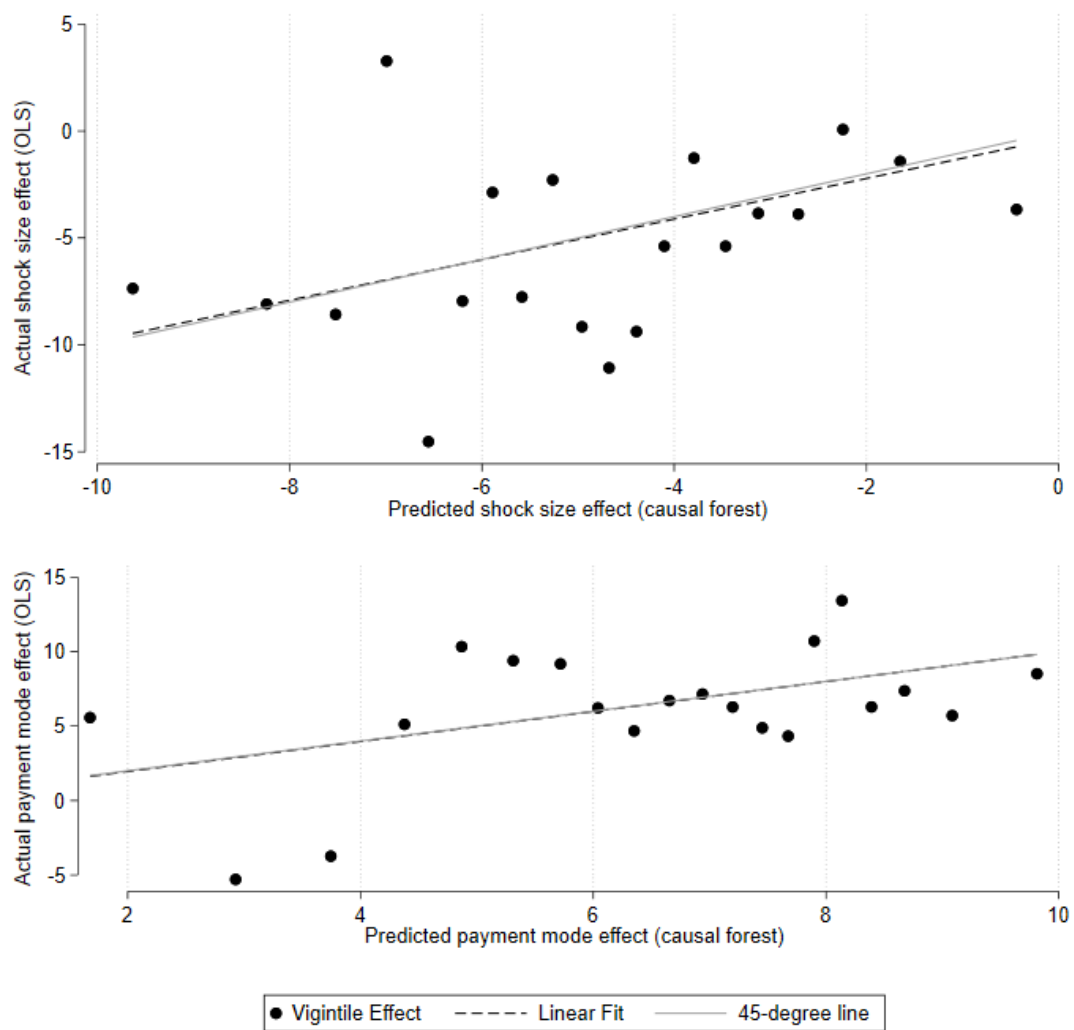


Figure B.1: Predicted vs. actual treatment effects

Data source: BOP-HH, December 2020 and June 2021.

Notes: This figure compares treatment effects from the causal forest with those estimated with OLS. Treatment effects predicted by the causal forest are plotted against actual treatment effects estimated with OLS for the shock size effect (upper panel) and the payment mode effect (lower panel). Observations are binned into twenty groups by percentile of predicted treatment effect. The x-axis shows the average predicted treatment effect for each bin. The y-axis shows the average treatment effect estimated with OLS for each bin (vigintile effect). The 45-degree line (solid line) serves as a benchmark, reflecting a one-to-one relation between predicted and actual treatment effects. The dashed line is a linear fit between the actual and predicted treatment effects.

tom quartile) significantly decreases the MPC by about 8 percentage points in response to an increase in the shock size, whereas the estimated effect for the group with the least negative prediction (top quartile) is not statistically significantly different from zero. The treatment coefficients of these groups statistically differ from each other (p-value of difference = 0.01). In a similar vein, the OLS-estimated treatment effect of changing the payment mode comes to 9% for the subsample of respondents predicted to react most strongly to changes in the payment mode (top quartile). The estimated effect for those predicted to react the least (bottom quartile) is not statistically significantly different from zero. Importantly, the difference in estimated effect across these two groups is statistically significant (p-value of difference = 0.01).

In sum, there is substantial heterogeneity in treatment effects for the shock size and the payment mode effect, and the causal forest algorithm produces valid predictions of these heterogeneous treatment effects.

Table B.1: Quality assessment of causal forest predictions

	Estimate/[p-value]	
	Shock size effect	Payment mode effect
Panel A: Calibration test		
Mean forest prediction	1.00 [0.00]	1.00 [0.00]
Differential forest prediction	0.95 [0.01]	1.04 [0.01]
Panel B: Bottom vs. top quartile of predicted treatment effect		
Bottom quartile	-7.47 [0.00]	2.43 [0.17]
Top quartile	-1.22 [0.46]	9.02 [0.00]
Difference	6.25 [0.01]	6.59 [0.01]

Data source: BOP-HH, December 2020 and June 2021.

Notes: This table assesses the quality of treatment effect predictions of the causal forest. Panel A shows results from a calibration test as described by [Athey and Wager \(2019\)](#). The test computes the best linear fit of the target estimand using the forest prediction (on held-out data) as well as the mean forest prediction as the sole two regressors. A coefficient of 1 for *mean forest prediction* indicates that the mean forest prediction is correct. A coefficient of 1 for the *differential forest prediction* indicates that the heterogeneity estimates from the forest are well calibrated. The p-value of the *differential forest prediction* coefficient (shown in parentheses) acts as an omnibus test for the presence of heterogeneity, i.e. the null hypothesis of no heterogeneity can be rejected if the coefficient is significantly greater than 0. Panel B compares the actual treatment effects of respondents with the lowest and highest predicted treatment effects. Using the subsample of respondents in the bottom and top quartile of the predicted treatment effect distribution, respectively, treatment effects are estimated with OLS separately for each subsample. The last row reports results from testing the null hypothesis that the treatment coefficients for the groups of respondents in the bottom and top quartile are equal. P-values are reported in parentheses.

C Robustness Checks and Additional Analyses

This appendix evaluates the sensitivity of the empirical findings to a variety of robustness checks and provides additional analyses.

Heaping The distribution of the reported MPC variable is characterized by heaping at values of 0, 50, and 100% (Figure 1a). In particular, the high fraction of respondents reporting MPC = 50 might be concerning as it could reflect respondents' uncertainty about the meaning of the question. Heaping may thus be associated with personal characteristics such as low cognitive sophistication. To test this, I use respondents' self-reported mathematical skills as a proxy for cognitive sophistication and run conditional logit regressions for the probability of reporting an MPC equal to 0, 50, and 100%, respectively, in the one month's income shock baseline scenario (BASE-1).⁵⁸ Table D.1 shows that neither the coefficient for low mathematical skills, nor education is statistically significant at the 5% level in the regression for the probability of reporting MPC = 50 (column 1). The same conclusion holds for the probability of reporting MPC = 0 (column 2) and MPC = 100 (column 3), respectively.⁵⁹

Misunderstanding of the question A related problem might be that some respondents do not fully understand the hypothetical scenario in the survey questions. To assess the robustness of the main regression results to respondents' comprehension of the survey question, I rerun Equation 1 conditioning on individuals less prone to misunderstand or misinterpret the MPC question. Table D.3 shows that the findings for the ATEs are confirmed when performing this exercise. The shock size and payment mode effect, respectively, on the MPC becomes even stronger when observations with MPC=50 (column 1) are dropped. Moreover, the regression coefficients remain in the

⁵⁸I generate a dummy that equals 1 if respondents disagree with the statement "*I have a lot of confidence in my mathematical skills*", i.e. scoring 4 or lower on a scale from 1, "*strongly disagree*", to 7, "*strongly agree*."

⁵⁹As an alternative measure for respondents' understanding of the question I consider their perception about the length and difficulty of the questionnaire as well as their reported interest elicited via feedback questions at the end of the interview. I find no systematic relation between these measures and heaping in the MPC variable. Results are available upon request.

same ballpark when those with low mathematical skills (column 2), those reporting the questionnaire was difficult (column 3), too long (column 4), and not interesting (column 5), as well as those with extreme inflation expectations (column 6) are excluded, respectively.⁶⁰

Censoring of the outcome variable Another concern might be that the MPC variable is censored from below, at 0, and above, at 100, which can bias the OLS estimates of the ATEs. To this end, I estimate the specification in [Equation 1](#) with a two-limit Tobit estimator. [Table D.4](#) shows that the estimated effect for the shock size and payment mode treatments on the MPC is almost identical between simple OLS (column 1) and the Tobit estimator (column 2).

Transaction costs The preferred interpretation of the payment mode effect is that the labeling of windfalls induced by different payment modes drives the difference in MPCs between, on the one hand, those receiving the windfall paid out in cash and without any specification in the payment mode, respectively, and, on the other hand, those receiving the windfall deposited in an savings account. An alternative explanation for the difference in MPCs across these groups is that the latter consider it too burdensome to transfer the money from the savings account to their current account in order to spend it. To study whether motives related to (mental) transaction costs drive the payment mode effect, participants in the June 2021 wave of the BOP-HH were directly asked about the main reasons for their decision.⁶¹ [Figure D.2](#) plots the reported reasons across quartiles of the predicted payment mode effect. The share of respondents reporting that *"it would be too much effort to transfer the money from the savings account to the current account"* is close to zero among those predicted to react most strongly to changes in the payment mode (top quartile) and is lower,

⁶⁰Extreme inflation expectations are defined as reporting a point forecast for one-year-ahead inflation above (below) the 95th (5th) percentile of the inflation expectations distribution. Having extreme inflation expectations might be an indication of low financial literacy and hence a low understanding of the MPC question. Inflation expectations as well as feedback on the questionnaire's perceived length, difficulty, and attractiveness, respectively, are elicited in each survey wave. [Appendix A.1](#) reports the exact wording of the survey questions.

⁶¹The survey question is reported in [Appendix A.1](#).

rather than higher, compared to those predicted to react the least (bottom quartile). The results suggest that motives related to transaction costs do not play a role for the payment mode effect. Importantly, the share of respondents reporting that they "*never spend money from their savings account*" or "*do not like spending money from their savings account*" is highest among those predicted to react most strongly to changes in the payment mode (top quartile), in line with respondents associating the savings account with the *current assets* mental account.

Macroeconomic environment Another concern might be that the strong consumption response to changes in the shock size and payment mode is contingent on the macroeconomic environment. Most particularly, the fact that the survey module on the MPC was fielded during the COVID-19 pandemic might limit the external validity of the results. To address this point, I study data from a pilot study administered before the outbreak of the pandemic. In September 2019, participants in the retail bank survey were asked about their consumption response to a lottery win equal to their monthly net income. Additionally, they were asked about their response to a lottery win equal to three months' household income, whereby the payment mode was randomized among respondents in this scenario.⁶² One-third of respondents received the prize money "*paid out in cash*", another third received the prize money "*deposited in an instant-access savings account that pays interest*", and the remainder received the money without any specification of the payment mode.⁶³ To control for potential order effects, a counterbalanced measures design was implemented, i.e. half of the sample was asked first about their response to a one month's income shock and the other half was asked first about a three months' income shock.⁶⁴

Panel A in [Table D.5](#) shows the within differences in MPCs out of a one month's and a three months' income shock. Irrespective of the payment mode, the MPC out of a large shock is significantly lower than the MPC out of a small shock (p-value <0.01). Panel B in [Table D.5](#) reports results from OLS regressions of the MPC out the three

⁶²The payment mode in the one month's income shock scenario was not specified. Accordingly, all respondents received the same questions in this scenario.

⁶³[Appendix A.2](#) reports the exact wording of the survey questions.

⁶⁴[Table D.13](#) shows that the treatment groups are comparable along observable characteristics.

months' income shock on an indicator variable for the payment mode, controlling for order FE effects and sociodemographic characteristics. Respondents who receive a lottery win equal to three months of household income paid out in cash and without any specification of the payment mode, respectively, report a significantly higher MPC than those receiving an equally-sized shock deposited in an instant-access savings account.

In sum, the finding of consumers adjusting their MPC in response to changes in the shock size and payment mode does not seem to depend on the macroeconomic environment.

Payment mode effect by type of savings account One might wonder to what extent the payment mode effect depends on the type of savings account under consideration. To this end, I used a slightly different wording in the June 2021 wave, referring to a classic "*savings account that is accessible at any time*" rather than to an "*instant-access savings account that pays interest*" as in the December 2020 wave. In [Table D.6](#), I estimate separately for each survey wave the treatment effects of varying the shock size and payment mode on the MPC, using [Equation 1](#). Importantly, [Table D.6](#) shows that the payment mode effect is present in both waves, i.e. respondents exposed to an income shock paid out in cash exhibit a significantly higher MPC than those exposed to an income shock either deposited in an instant-access savings account (columns 1 and 2) or deposited in a classic savings account (columns 3 and 4). Therefore, I pool these two savings account scenarios in the main analysis to increase statistical power.

Interestingly, the payment mode effect is larger when referring to an instant-access savings account (columns 1 and 2) rather than to a classic savings account (columns 3 and 4). This result might suggest that consumers are more likely to associate funds deposited in an instant-access savings account with the *current assets* mental account as this type of savings account is used more frequently than classic savings accounts. For example, [Figure D.5](#) illustrates that the amount of money held by German households as transferable deposits (*Sichteinlagen*), comprising instant access accounts (*Tagesgeldkonten*) and current accounts (*Girokonten*), has grown steadily since the

early 2000s. In contrast, the amount held in classic savings accounts (*Spareinlagen*) has fallen by more than half over the same period, such that assets held as transferable deposits were three times greater in value than those held in savings accounts in 2021.

One potential alternative explanation for the difference in effects across waves is that the additional word "*pays interest*" in the December 2020 wave nudges respondents into saving by evoking thoughts related to investment – and not provoking any thinking related to mental accounting. However, the type of account referred to in the June 2021 wave explicitly includes the root word "saving" (*Sparkonto*), which is not the case in the December 2020 wave (*Tagesgeldkonto*). Thus, it is unclear why the scenario in the June 2021 wave should evoke a weaker association with the topic of investment and therefore cause a less negative effect on the MPC.

Another interpretation for the larger payment mode effect in the December 2021 wave is that the effect is subject to seasonality. However, the coefficient measuring the effect of increasing the shock size on the MPC is remarkably stable across waves, as shown in [Table D.6](#). Therefore, the question arises why seasonality should play a role only for the payment mode effect and not for the shock size effect. Nonetheless, a seasonality-based explanation is hard to rule out directly in the current setting.

Accumulated local effect plots The construction of the partial dependence (PD) plots, shown in [Figure 5](#) and [Figure 6](#), relies on the assumption that the variable for which the partial dependence is computed is independent of other variables. Specifically, to obtain the partial dependence function at a particular value of a heterogeneity variable of interest, each observation in the data is forced to take on that value, while at the same time the values of all other variables are kept fixed at their empirical value. This might create new data points that are unlikely to occur in reality, e.g. ranging in the top of the income distribution at very young age. An alternative approach to studying the marginal effect of a variable on a predicted outcome would be to assess accumulated local effect (ALE) plots, which effectively control for the influence of correlated variables. To construct an ALE plot, one uses the conditional distribution of the variable of interest to generate augmented data and calculate the average of the

differences in model predictions over the augmented data, instead of the average of the predictions themselves as in the case of constructing PD plots. To be more specific, ALE plots are generated in three steps: First, one divides a variable of interest into intervals. Second, for the observations in an interval, one calculates the difference in the prediction when replacing the variable with the upper and lower limit of the interval. Third, these differences are accumulated and centered, resulting in the ALE curve.⁶⁵ [Figure D.3](#) and [Figure D.4](#) depict ALE plots for the variables examined in [Section 4.2](#). The results are qualitatively similar to those shown in [Figure 5](#) and [Figure 6](#).

CATE using sample-splitting approach [Table D.8](#) studies treatment heterogeneity using a conventional sample-splitting approach. To estimate CATEs, I run separate OLS regressions for the subsamples indicated in the column header, using the specification in [Equation 1](#). The results are consistent with the heterogeneity patterns explored with the causal forest in [Section 4.2](#). Note, however, that relying only on the sample-splitting approach would have made it challenging to detect the heterogeneity patterns shown in [Figure 5](#) and [Figure 6](#).

⁶⁵The y-axis is centered at zero and hence the value at each point of the ALE curve gives the difference to the mean prediction. I use the *FeatureEffects* function from the *iml* package in *R* to generate the ALE plots.

D Additional Tables and Figures

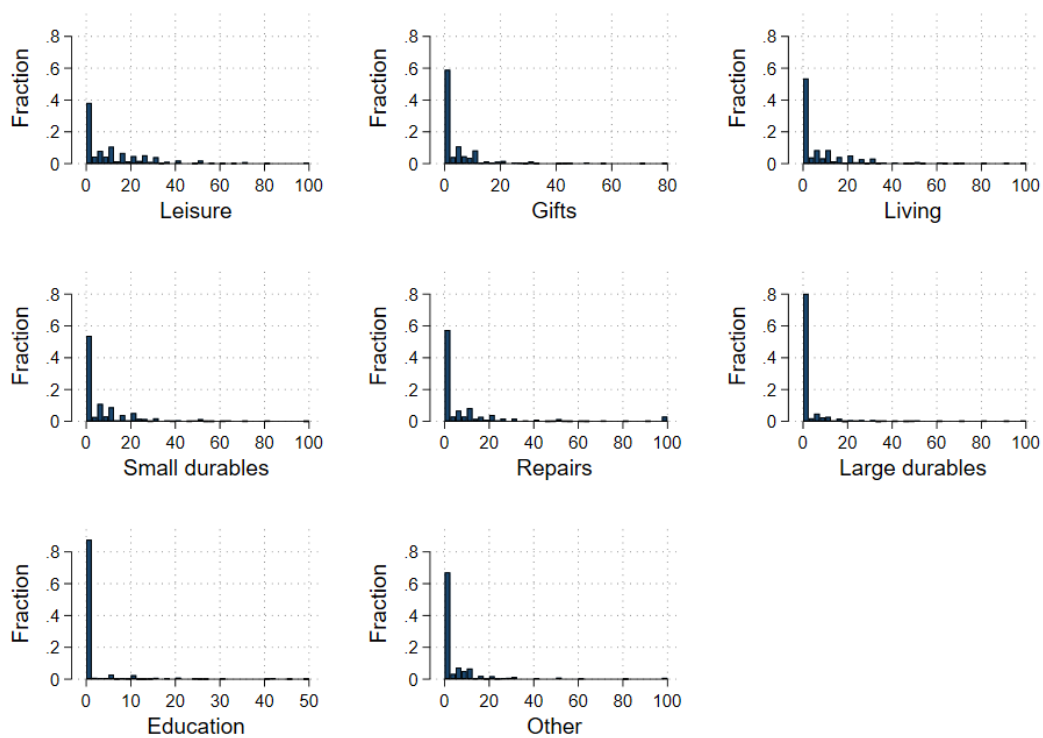


Figure D.1: Histograms of MPCs for categories of consumption

Data source: BOP-HH, December 2020 and June 2021.

Notes: This figure plots the distribution of the marginal propensity to consume out of the one month's income shock in the baseline scenario (BASE-1) for different categories of consumption.



Figure D.2: Reasons for spending by quartiles of predicted payment mode effect

Data source: BOP-HH, June 2021.

Notes: This figure plots the share of respondents reporting that *"it would be too much effort to transfer the money from the savings account to the current account"* as well as the share of respondents reporting that they *"never spend money from the savings account"* or *"do not like spending money from the savings account"* across quartiles of the predicted payment mode effect. The sample is restricted to the June 2021 wave as the survey question was asked only in that survey wave.

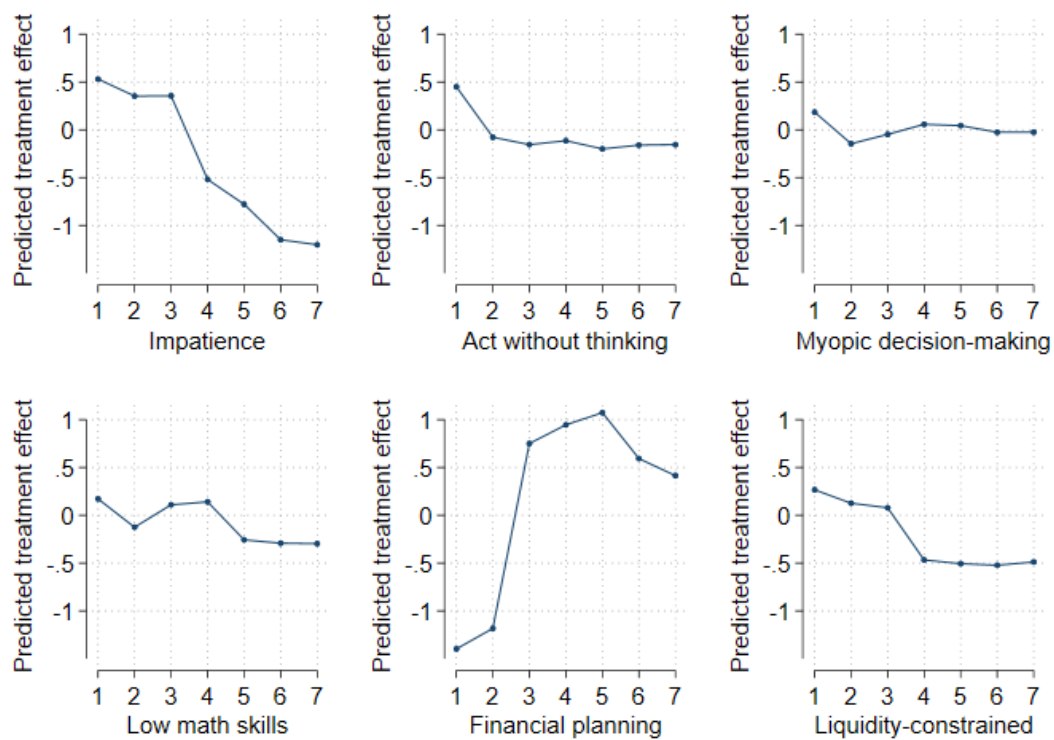


Figure D.3: Accumulated local effects plots for the shock size effect

Data source: BOP-HH, December 2020 and June 2021.

Notes: The figure depicts accumulated local effects (ALE) plots for the variables shown in [Figure 5](#). Each panel shows how the predicted shock size effect varies with a variable of interest conditional on a given value. The y-axis is centered at zero. Accordingly, the value at each point of the ALE curve gives the difference to the mean prediction.

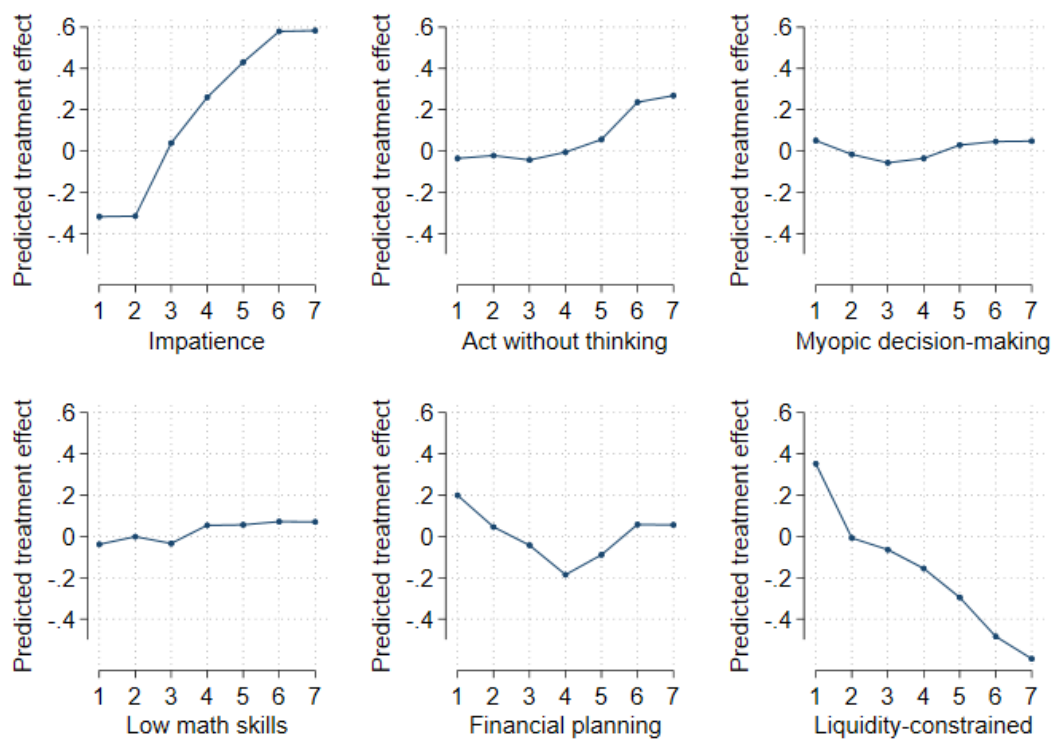


Figure D.4: Accumulated local effects plots for the payment mode effect

Data source: BOP-HH, December 2020 and June 2021.

Notes: The figure depicts accumulated local effects (ALE) plots for the variables shown in [Figure 6](#). Each panel shows how the predicted payment mode effect varies with a variable of interest conditional on a given value. The y-axis is centered at zero. Accordingly, the value at each point of the ALE curve gives the difference to the mean prediction.




Figure D.5: Composition of currency and deposits held by households in Germany

Data source: Financial accounts (Deutsche Bundesbank). Own calculations.

Notes: This figure shows the composition of currency and deposits held by households in Germany, using end-of-quarter data from the financial accounts. The y-axis depicts assets held by households in the form of currency (*Bargeld*), savings deposits (*Spareinlagen*), savings certificates (*Sparbriefe*), time deposits (*Termineinlagen*), and transferable deposits (*Sichteinlagen*), as a share of total assets held in the form of currency and deposits (expressed in %). Transferable deposits, comprising current accounts (*Girokonten*) and instant access accounts (*Tagesgeldkonten*), are deposits with monetary financial institutions (MFIs) that are available without notice.

Table D.1: The role of heaping for the reported MPC Variable

	I(MPC = 50)	I(MPC = 0)	I(MPC = 100)
Age 45-64	-0.01 (0.04)	0.01 (0.02)	0.07** (0.03)
Age 65+	-0.01 (0.05)	0.09*** (0.03)	0.06 (0.04)
Female	0.07** (0.03)	0.02 (0.02)	-0.02 (0.03)
Bachelor	0.03 (0.04)	0.02 (0.03)	-0.05* (0.03)
Master and above	-0.06* (0.03)	0.04 (0.03)	0.01 (0.03)
HHsize 2	-0.02 (0.04)	-0.03 (0.03)	0.07** (0.03)
HHsize 3	-0.07 (0.06)	-0.02 (0.05)	0.04 (0.05)
HHsize 4+	-0.06 (0.06)	-0.03 (0.05)	0.19*** (0.06)
HH income €2000-3000	0.00 (0.05)	0.03 (0.04)	0.03 (0.05)
HH income €3000-4000	-0.01 (0.05)	0.06 (0.04)	-0.02 (0.05)
HH income €4000+	0.01 (0.05)	0.05 (0.04)	-0.04 (0.06)
Homeowner	0.03 (0.03)	-0.01 (0.03)	0.02 (0.03)
Unemployed	0.06 (0.09)	0.01 (0.09)	0.00 (0.09)
City size 20-100k	-0.03 (0.03)	-0.01 (0.03)	0.05 (0.03)
City size 100k+	0.06* (0.04)	-0.02 (0.03)	-0.05 (0.03)
Eastern Germany	-0.01 (0.04)	-0.02 (0.03)	0.04 (0.03)
Low math skills	0.01 (0.04)	-0.05 (0.03)	-0.01 (0.03)
Wave FE	Yes	Yes	Yes
Observations	1013	1013	1013

Data source: BOP-HH, December 2020 and June 2021.

Notes: The table studies the role of heaping in the reported MPC variable. Columns (1) to (3) report marginal effects from conditional logit regressions for the probability of reporting an MPC out of a one month's income shock equal to 0, 50, and 100%, respectively, using the sample exposed to the one month's income baseline scenario (BASE-1). Low math skill is a dummy that equals 1 if respondent i disagrees with the statement "*I have a lot of confidence in my mathematical skills*", i.e. scores 4 or lower on a scale from 1, "*strongly disagree*", to 7, "*strongly agree*". Robust standard errors are reported in parentheses. *, **, *** indicate significance at the 10%, 5%, and 1% level, respectively.

Table D.2: Average treatment effects of shock size and payment mode scenarios, consumption categories

	Total Consumer Spending	Nondurables			Durables				Other Consumer Spending
		Leisure	Gifts	Living	Large durables	Small durables	Repairs	Edu- cation	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
<i>Shock size effect (omitted category: one month's income shock)</i>									
Three months' income shock	-5.47*** (0.84)	-2.82*** (0.40)	-0.49** (0.21)	-1.85*** (0.32)	1.21*** (0.34)	-0.66** (0.27)	0.15 (0.38)	-0.09 (0.14)	-0.91*** (0.33)
<i>Payment mode effect (omitted category: savings account scenario)</i>									
Baseline scenario	5.55*** (1.04)	0.14 (0.48)	0.43* (0.26)	1.43*** (0.38)	0.56 (0.41)	0.92*** (0.33)	1.31*** (0.44)	0.01 (0.17)	0.75* (0.40)
Cash payment scenario	6.42*** (1.04)	0.90* (0.50)	0.50* (0.27)	1.34*** (0.38)	0.95** (0.42)	0.85** (0.33)	1.45*** (0.46)	-0.05 (0.18)	0.48 (0.40)
p-value $\beta_{\text{Baseline}} = \beta_{\text{Cash payment}}$	0.40	0.12	0.79	0.83	0.36	0.82	0.78	0.74	0.51
Sociodemographics	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.03	0.02	0.04	0.03	0.02	0.02	0.05	0.03	0.01
Observations	6051	6051	6051	6051	6051	6051	6051	6051	6051
Mean for BASE-1 scenario	52.46	12.37	4.52	8.36	4.35	7.18	8.64	1.54	5.51

Data source: BOP-HH, December 2020 and June 2021.

Notes: This table reports average treatment effects (ATEs) on different categories of consumption estimated with the following OLS regression equation: $Y_i = \alpha + \gamma \times \text{LargeShock} + \sum_{s=1}^S \beta_s \times \text{PaymentMode}_s + \delta \mathbf{X}_i + \epsilon_i$, where Y_i is the share (scaled to be between 0 and 100) of individual i 's income shock that went to the category of consumption indicated in the column header. *LargeShock* is an indicator variable equal to 1 if respondent i received a three months' income shock and 0 if respondent i received a one month's income shock. The γ coefficient identifies the ATE of increasing the shock size from one month's income to three months' of income. *PaymentMode_s* is an indicator variable that equals 1 if respondent i receives an income shock paid out in cash and without any specification of the payment mode, respectively, and 0 otherwise. The β_s coefficients identify the ATEs of the different payment modes. \mathbf{X}_i is a vector of sociodemographic controls. The last row reports the mean of the outcome variable for the group of respondents receiving a one month's income without any specification of the payment mode. Column (1) replicates results from Table 3. Robust standard errors are reported in parentheses. *, **, *** indicate significance at the 10%, 5%, and 1% level, respectively.

Table D.3: The role of measurement errors

	MPC					
	Excluding MPC=50	Excluding low math	Excluding questionnaire too difficult	Excluding questionnaire too long	Excluding questionnaire not interesting	Excluding extreme inflation expectations
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Shock size effect (omitted category: one month's income shock)</i>						
Three months' income shock	-7.20*** (1.14)	-5.01*** (0.94)	-5.09*** (0.96)	-5.71*** (1.07)	-5.41*** (0.88)	-5.70*** (0.89)
<i>Payment mode effect (omitted category: savings account scenario)</i>						
Baseline scenario	7.13*** (1.37)	5.54*** (1.15)	6.24*** (1.18)	5.02*** (1.32)	5.32*** (1.09)	5.48*** (1.09)
Cash payment scenario	8.34*** (1.39)	6.21*** (1.15)	6.80*** (1.18)	6.64*** (1.32)	6.10*** (1.09)	6.70*** (1.09)
Sociodemographics	Yes	Yes	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.04	0.03	0.03	0.03	0.03	0.03
Observations	4478	4958	4733	3847	5462	5472
Mean for BASE-1 scenario	53.42	51.89	53.19	52.95	52.70	52.88

Data source: BOP-HH, December 2020 and June 2021.

Notes: This table assesses the robustness of the main regression results to respondents' comprehension of the survey question. The baseline regression from column (2) in Table 3 is replicated for subsamples of respondents less prone to measurement error. Column (1) excludes respondents reporting a MPC=50. Column (2) excludes respondents reporting having low confidence in their mathematical skills. Column (3) excludes respondents reporting that the questionnaire was "very difficult" or "somewhat difficult". Column (4) excludes respondents reporting that the questionnaire was "too long" or "somewhat too long". Column (5) excludes respondents reporting that the questionnaire was "not interesting" or "less interesting". Column (6) excludes respondents reporting extreme inflation expectations, defined as reporting a point forecast for one-year-ahead inflation above (below) the 95th (5th) percentile of the inflation expectations distribution. Robust standard errors are reported in parentheses. *, **, *** indicate significance at the 10%, 5%, and 1% level, respectively.

Table D.4: Average treatment effects, using a Tobit model

	MPC			
	OLS		Tobit	
	(1)	(2)	(3)	(4)
<i>Shock size effect (omitted category: one month's income shock)</i>				
Three months' income shock	-4.88*** (0.82)	-5.47*** (0.84)	-4.72*** (0.85)	-5.38*** (0.87)
<i>Payment mode effect (omitted category: savings account scenario)</i>				
Baseline scenario	5.90*** (1.01)	5.55*** (1.04)	6.06*** (1.04)	5.78*** (1.07)
Cash payment scenario	6.86*** (1.01)	6.42*** (1.04)	7.14*** (1.03)	6.74*** (1.06)
Sociodemographics	No	Yes	No	Yes
Wave FE	Yes	Yes	Yes	Yes
Observations	6373	6051	6373	6051

Data source: BOP-HH, December 2020 and June 2021.

Notes: This table compares average treatment effects (ATEs) of the payment mode and shock size scenarios on the MPC using different estimators. Columns (1) and (2) replicate the OLS estimation results from [Table 3](#). Columns (3) and (4) report regression results using a two-limit Tobit estimator, showing marginal effects at mean. Robust standard errors are reported in parentheses. *, **, *** indicate significance at the 10%, 5%, and 1% level, respectively.

Table D.5: Test of means for shock size scenarios and regression for payment mode scenarios

Panel A: Test of Means for MPC Differences		
	Three months' minus one month's shock	
	Mean	p-value
Baseline scenario	-7.38	2.951×10^{-09}
Cash payment scenario	-7.81	2.981×10^{-11}
Savings account scenario	-8.71	4.401×10^{-12}
Panel B: Regression for Payment Mode Scenarios		
	MPC	
	(1)	(2)
<i>Payment mode effect (omitted category: savings account scenario)</i>		
Baseline scenario	5.13*** (1.71)	5.33*** (1.75)
Cash payment scenario	7.14*** (1.73)	7.26*** (1.77)
Sociodemographics	No	Yes
Order FE	Yes	Yes
R^2	0.01	0.05
Observations	1884	1744

Data source: Retail bank survey, September 2019.

Notes: This table reports results from a pilot study administered in the retail bank survey in September 2019. Column (1) in Panel A reports the means of the within differences in MPCs out of a one month's and a three months' income shock for different payment modes. Column (2) reports p-values from a t-test of the null hypothesis that the difference in means equals 0. Panel B reports results from OLS regressions of the MPC out the three months' income shock on an indicator variable for the payment mode. The average treatment effects (ATEs) of the different payment modes is estimated with the following OLS regression equation: $MPC_i = \alpha + \sum_{s=1}^S \beta_s \times PaymentMode_s + \delta X_i + \epsilon_i$, where Y_i is the MPC of individual i . $PaymentMode_s$ is an indicator variable that equals 1 if respondent i receives an income shock paid out in cash and without any specification of the payment mode, respectively, and 0 otherwise. The β_s coefficients identify the ATEs of the different payment modes. The second column uses the same specification as the first column but augmented with respondent-specific controls X_i . All regressions include a dummy that equals 1 if respondent i was asked first about her response to a three months' income shock and 0 otherwise (Order FE). Robust standard errors are reported in parentheses. *, **, *** indicate significance at the 10%, 5%, and 1% level, respectively.

Table D.6: Shock size and payment mode effect by survey wave

	MPC			
	December 2020		June 2021	
	(1)	(2)	(3)	(4)
<i>Shock size effect (omitted category: one month's income shock)</i>				
Three months' income shock	-5.02*** (1.04)	-5.69*** (1.07)	-4.69*** (1.34)	-5.08*** (1.37)
<i>Payment mode effect (omitted category: savings account scenario)</i>				
Baseline scenario	9.14*** (1.28)	8.35*** (1.32)	0.64 (1.64)	1.41 (1.67)
Cash payment scenario	9.08*** (1.28)	8.47*** (1.32)	3.28** (1.65)	3.36** (1.69)
p-value $\beta_{\text{Baseline}} = \beta_{\text{Cash payment}}$	0.97	0.93	0.11	0.24
Sociodemographics	No	Yes	No	Yes
R^2	0.02	0.03	0.01	0.03
Observations	3943	3718	2430	2333
Mean for BASE-1 scenario	51.81	51.65	53.14	53.74

Data source: BOP-HH, December 2020 and June 2021.

Notes: This table reports average treatment effects (ATEs) of the shock size and payment mode scenarios estimated with the following OLS regression equation for each wave of the BOP-HH separately: $Y_i = \alpha + \gamma \times \text{LargeShock} + \sum_{s=1}^S \beta_s \times \text{PaymentMode}_s + \delta \mathbf{X}_i + \epsilon_i$, where Y_i is the MPC of individual i . *LargeShock* is an indicator variable equal to 1 if respondent i received a three months' income shock and 0 otherwise. The γ coefficient identifies the ATE of increasing the shock size from one month's income to three months' income. *PaymentMode_s* is an indicator variable that equals 1 if respondent i receives an income shock paid out in cash and without any specification of the payment mode, respectively, and 0 otherwise. The β_s coefficients identify the ATEs of the different payment modes. Even columns uses the same specification as odd columns but augmented with respondent specific controls \mathbf{X}_i . The last row reports the mean of the outcome variable for the group of respondents receiving a one month's income without any specification of the payment mode. Robust standard errors are reported in parentheses. *, **, *** indicate significance at the 10%, 5%, and 1% level, respectively.

Table D.7: Average MPCs by survey wave

	MPC		
	Savings account scenario	Cash payment scenario	Baseline scenario
	(1)	(2)	(3)
Panel A: one month's shock			
December 2020 Wave	41.65	51.51	51.81
June 2021 Wave	50.33	56.80	53.14
Panel B: three months' shock			
December 2020 Wave	37.85	46.14	45.94
June 2021 Wave	49.20	49.29	47.70

Data source: BOP-HH, December 2020 and June 2021.

Notes: This table shows average MPCs out of different income shock scenarios separately by survey wave.

Table D.8: Conditional average treatment effects of shock size and payment mode scenarios, using OLS

	Impatience		Act without thinking		Myopic decision-making		Low math skills		Long-term financial planning		Liquidity-constrained	
	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes	No	Yes
<i>Shock size effect (omitted category: one month's income shock)</i>												
Three months' income shock	-4.88*** (0.98)	-6.97*** (1.65)	-5.31*** (0.95)	-6.32*** (1.81)	-5.46*** (0.90)	-5.48** (2.29)	-5.01*** (0.94)	-7.79*** (1.98)	-5.98*** (1.03)	-4.62*** (1.48)	-5.30*** (0.97)	-6.13*** (1.73)
<i>Payment mode effect (omitted category: savings account scenario)</i>												
Baseline scenario	4.99*** (1.20)	7.18*** (2.07)	4.73*** (1.16)	7.87*** (2.28)	5.26*** (1.11)	7.25** (2.91)	5.54*** (1.15)	5.79** (2.41)	4.73*** (1.27)	7.16*** (1.80)	6.45*** (1.19)	2.96 (2.13)
Cash payment scenario	5.48*** (1.21)	9.17*** (2.05)	6.19*** (1.18)	6.95*** (2.24)	6.24*** (1.12)	7.86*** (2.89)	6.21*** (1.15)	7.47*** (2.46)	6.39*** (1.27)	6.32*** (1.83)	6.95*** (1.19)	4.72** (2.13)
Sociodemographics	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4480	1570	4748	1301	5192	859	4958	1091	4022	2027	4616	1426
Adjusted R^2	0.02	0.04	0.03	0.02	0.02	0.02	0.02	0.02	0.03	0.02	0.02	0.03
Unconditional MPC	47.44	48.75	47.64	48.34	46.87	53.32	47.54	48.79	47.92	47.51	47.00	50.34

Data source: BOP-HH, December 2020 and June 2021.

Notes: The table reports conditional average treatment effects (CATEs) estimated with OLS for the respective subsample indicated in the column header, using the specification in Equation 1. Behavioral and household balance sheet characteristics are elicited with the following statements to which respondents can disagree or agree with on a 7-point Likert scale. Being impatient is defined as disagreeing with the statement "I am generally a very patient person", i.e. scoring 3 or lower. Acting without thinking is defined as disagreeing with the statement "I rarely do anything without thinking about it thoroughly." Myopic decision-making is defined as agreeing with the statement "I live in the here and now and don't really think about the future", i.e. scoring 5 or higher on a 7-point Likert scale. Having low mathematical skills is defined as disagreeing with the statement "I have a lot of confidence in my mathematical skills." Financial planning is defined as agreeing with the statement "I plan major spending and investment decisions more than one year ahead." Being liquidity-constrained is defined as disagreeing with the statement "I have put aside money for a possible emergency so that I can cover expenses for at least three months with no income." The last row reports the unconditional average MPC for the subgroup of respondents indicated in the column header. Robust standard errors are reported in parentheses. *, **, *** indicate significance at the 10%, 5%, and 1% level, respectively.

Table D.9: Balance of treatment groups: BOP–HH

	Full sample	BASE-1	CASH-1	SAVINGS- ACCOUNT-1	BASE-3	CASH-3	SAVINGS- ACCOUNT-3	P-value
Age	55.20	55.90	54.89	55.30	54.95	55.32	54.86	0.09
College and more	0.54	0.52	0.53	0.55	0.51	0.55	0.56	0.07
Female	0.41	0.39	0.39	0.44	0.40	0.42	0.39	0.08
Household size	2.05	2.00	2.07	2.03	2.07	2.08	2.08	0.25
HH income <€2000	0.14	0.15	0.14	0.14	0.14	0.15	0.14	0.62
HH income €2000-3000	0.22	0.25	0.21	0.25	0.22	0.21	0.20	0.10
HH income €3000-4000	0.23	0.25	0.25	0.22	0.22	0.23	0.24	0.32
HH income €4000+	0.40	0.35	0.40	0.39	0.42	0.41	0.42	0.15
Homeowner	0.65	0.64	0.66	0.65	0.65	0.64	0.65	0.94
Unemployed	0.02	0.02	0.01	0.02	0.02	0.01	0.02	0.86
City size 100k+	0.33	0.35	0.32	0.32	0.32	0.32	0.32	0.43
Eastern Germany	0.18	0.19	0.19	0.20	0.17	0.17	0.19	0.35
Dummy for June 2021	0.38	0.38	0.39	0.38	0.38	0.38	0.39	0.00
Observations	6427	1074	1078	1072	1066	1068	1069	

Data source: BOP–HH, December 2020 and June 2021.

Notes: This table compares respondents' observable characteristics across the different treatment groups. Columns (1) to (7) report sample averages of characteristics for the group indicated in the column header. Column 8 reports the p-value of a one-way ANOVA test of equality of each row variable across the six treatment groups.

Table D.10: Summary statistics: retail bank survey

	Mean	SD	Median	Min	Max	Obs.
Age	49.91	15.14	51.00	19.00	80.00	2198
College	0.57	0.49	1.00	0.00	1.00	2147
Female	0.35	0.48	0.00	0.00	1.00	2205
HHsize	2.37	0.99	2.00	1.00	4.00	2205
Unemployed	0.01	0.10	0.00	0.00	1.00	2152
Income <€2000	0.15	0.36	0.00	0.00	1.00	2068
Income €2000-3000	0.23	0.42	0.00	0.00	1.00	2068
Income €3000-4000	0.24	0.43	0.00	0.00	1.00	2068
Income €4000+	0.38	0.48	0.00	0.00	1.00	2068
Homeowner	0.61	0.49	1.00	0.00	1.00	2062
Living in East Germany	0.10	0.30	0.00	0.00	1.00	1714

Data source: Retail bank survey, October 2020.

Notes: This table reports summary statistics of variables used in the analysis for the October 2020 wave of the retail bank survey.

Table D.11: Balance of treatment groups: retail bank survey, October 2020

	Full sample	LOTTERY-BASE	GOV-BASE	LOTTERY-CASH	GOV-CASH	P-value
MPC	43.72	41.95	42.42	44.07	46.43	0.09
Age	49.91	49.09	49.28	50.73	50.51	0.17
College	0.57	0.56	0.60	0.58	0.54	0.24
Female	0.35	0.36	0.35	0.33	0.35	0.67
HHsize	2.37	2.39	2.41	2.37	2.29	0.20
Income <€2000	0.15	0.18	0.13	0.15	0.16	0.18
Income €2000-3000	0.23	0.22	0.25	0.21	0.25	0.26
Income €3000-4000	0.24	0.20	0.27	0.24	0.24	0.10
Income €4000+	0.38	0.41	0.36	0.40	0.35	0.13
Unemployed	0.01	0.02	0.00	0.01	0.01	0.24
Observations	2198	520	572	564	542	

Data source: Retail bank survey, October 2020.

Notes: This table compares respondents' MPC and observable characteristics across the different treatment groups in the October 2020 wave of the retail bank survey. Columns (1) to (5) report sample averages of characteristics for the group indicated in the column header. Column (6) reports the p-value of a one-way ANOVA test of equality of each row variable across the 4 treatment groups.

Table D.12: Comparison of survey samples: BOP-HH and retail bank survey

	Retail bank survey	BOP-HH	Difference	P-value
Age	49.91	55.20	5.29	0.00
College and more	0.57	0.54	-0.03	0.01
Female	0.35	0.41	0.06	0.00
Household size	2.37	2.20	-0.16	0.00
Unemployed	0.01	0.02	0.01	0.02
HH income <€2000	0.15	0.14	-0.01	0.31
HH income €2000-3000	0.23	0.22	-0.01	0.49
HH income €3000-4000	0.24	0.23	-0.01	0.58
HH income €4000+	0.38	0.40	0.02	0.07
Homeowner	0.61	0.65	0.04	0.00
Eastern Germany	0.10	0.18	0.08	0.00
Observations	2198.00	6427.00	.	.

Data sources: BOP-HH, December 2020 and June 2021. Retail bank survey, October 2020.

Notes: This table compares sociodemographic characteristics of participants in the retail bank survey and in the BOP-HH. Columns (1) and (2) report means for different sociodemographic characteristics for the sample indicated in the column header. Differences in means across samples are reported in column (3). Column (4) reports a p-value from a t-test for the null hypothesis that the two respective sample means are equal.

Table D.13: Balance of treatment groups: retail bank survey, September 2019

	Full sample	One month shock first			Three months shock first			P-value
		Base	Cash	Savings account	Base	Cash	Savings account	
MPC out of one month shock	48.42	49.86	52.71	44.74	48.07	49.58	45.59	0.08
MPC out of three months shock	40.47	43.04	44.97	34.82	39.97	42.08	37.98	0.00
Age	49.08	49.00	50.21	48.57	49.89	49.37	47.46	0.27
College	0.54	0.52	0.53	0.56	0.54	0.52	0.53	0.91
Female	0.32	0.36	0.32	0.33	0.29	0.29	0.34	0.41
HHsize	2.39	2.47	2.34	2.39	2.39	2.43	2.29	0.29
Income <€2000	0.17	0.19	0.18	0.17	0.17	0.17	0.17	0.98
Income €2000-3000	0.25	0.23	0.22	0.22	0.25	0.26	0.30	0.23
Income €3000-4000	0.22	0.26	0.22	0.20	0.22	0.23	0.21	0.53
Income €4000+	0.35	0.33	0.37	0.41	0.36	0.33	0.33	0.23
Unemployed	0.01	0.00	0.01	0.01	0.01	0.01	0.02	0.39
Observations	1854	308	314	311	304	308	309	

Data source: Retail bank survey, September 2019.

Notes: This table compares respondents' observable characteristics across the different treatment groups in the September 2019 wave of the retail bank survey. Columns (1) to (7) report sample averages of characteristics for the group indicated in the column header. Column (8) reports the p-value of a one-way ANOVA test of equality of each row variable across the treatment groups. The order of the MPC questions was randomized such that the treatment groups in column (2) to (4) were asked first about their consumption response to a one month's income shock and second about their response to a three months' income shock, while the opposite is the case for the treatment groups in columns (5) to (7). The payment mode was varied only in the three months' income shock scenarios, i.e. respondents received the three months' income shock either without any specification of the payment mode (baseline), or paid out in cash, or as a payment deposited in an instant-access savings account. The payment mode was not specified in the one month's income shock scenario for all respondents.