

NBER WORKING PAPER SERIES

QUICK-FIXING: NEAR-RATIONALITY IN CONSUMPTION AND SAVINGS BEHAVIOR

Peter Andre
Joel P. Flynn
Georgios Nikolakoudis
Karthik Sastry

Working Paper 33464
<http://www.nber.org/papers/w33464>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2025, Revised May 2026

We thank Alexandre Kohlhas, Morten Ravn, and Matthew Rognlie for excellent discussions which helped improve the paper. We are grateful to Kai Barron, John Becko, Corina Boar, Jonathon Hazell, Christopher House, David Laibson, Chen Lian, Lukas Nord, Ryan Oprea, Michael Peters, Oliver Pfäuti, Laura Pilossoph, Charlie Rafkin, Chris Roth, Tony Smith, Matthew White, Johannes Wohlfart, Michael Woodford, Chris Wolf, and seminar participants at the New York Fed, the University of Connecticut, the LAEF Workshop, CREI, Oxford, the Wadham State of the Nation Conference, Cambridge, the Philadelphia Fed, Princeton, Columbia, the 2024 Winter SED Meeting, the 2025 Spring NBER Monetary Economics Meeting, the University of Chicago, Carnegie-Mellon, the University of Rochester, the University of Michigan, the IZA Beliefs Workshop, the CIGS Conference on Macroeconomic Theory and Policy, the Bank of Portugal Conference on Monetary Economics, the Bank of Canada Behavioral Macro Workshop, and the 2025 NBER Summer Institute. First posted version: October 25, 2024. Ethics approval: The study obtained ethics approval from the Princeton University Institutional Review Board (#15968). Preregistration: The main data collections were preregistered via OSF (<https://osf.io/sgw9y/registrations>). The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Peter Andre, Joel P. Flynn, Georgios Nikolakoudis, and Karthik Sastry. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Quick-Fixing: Near-Rationality in Consumption and Savings Behavior
Peter Andre, Joel P. Flynn, Georgios Nikolakoudis, and Karthik Sastry
NBER Working Paper No. 33464
February 2025, Revised May 2026
JEL No. E21, E62, E70, G51

ABSTRACT

When changing consumption-savings plans is costly, people may rely on quick-fixes: simple rules that avoid these costs. We field a novel survey to elicit the distribution of households' consumption policy functions out of income shocks. Almost 70% of households follow one of four quick-fixes: they fully consume or save small shocks, but abruptly adjust their behavior for large shocks, and thereafter behave similarly. In a calibrated incomplete-markets model, quick-fixing is near-rational: the average opportunity cost of quick-fixing is only \$16 per quarter. Yet, this empirically realistic deviation from benchmark models significantly alters aggregate consumption responses to income shocks.

Peter Andre
Goethe University Frankfurt
Leibniz Institute for Financial Research SAFE
andre@safe-frankfurt.de

Georgios Nikolakoudis
The University of Chicago
Department of Economics
gnikolakoudis@gmail.com

Joel P. Flynn
Yale University
Department of Economics
and NBER
joel.flynn@yale.edu

Karthik Sastry
Princeton University
Department of Economics
and NBER
ksastry@princeton.edu

1 Introduction

Households’ consumption and savings choices, particularly their responses to income shocks, are a critical ingredient in essentially all models of the macroeconomy. Standard theory describes households as frictionless optimizers. However, continually optimizing consumption and savings plans to keep up with frequent changes in economic circumstances is practically challenging and often of limited benefit. For this reason, economists have acknowledged that behavior could be prone to *near-rationality* (Akerlof and Yellen, 1985b): suboptimal behavior when the gains from optimization are too small to justify the effort.

However, assessing the economic consequences of the near-rationality hypothesis is challenging: while there is typically a unique way to frictionlessly optimize, the set of almost optimal policy functions is vast. Moreover, although many consumption–savings plans can come close to replicating frictionless payoffs at the individual level, they can generate sharply different macroeconomic outcomes (Krusell and Smith, 1996). Evaluating the near-rationality hypothesis and its aggregate implications therefore requires bringing discipline to this “wilderness of irrationality” and determining which (if any) specific forms of near-rationality are actually prevalent.

To tackle this challenge, this paper combines novel evidence with a near-rational extension of a benchmark consumption-savings model and explores three questions. Do households’ responses to income shocks exhibit near-rationality? If so, what forms does this near-rationality take? And how does this behavior affect the aggregate response to income shocks?

We proceed in three steps. We first develop a near-rational model in which households either optimize or use a *quick-fix*: a simple policy with a (potentially small) convenience benefit. The model reveals how we can empirically identify households’ near-rational quick-fixes: quick-fixes are prevalent for smaller shocks, where mistakes are less costly, but abandoned for large shocks. Second, motivated by this observation, we design a novel survey to measure households’ consumption responses to a wide range of one-time income shocks. We find that most households rely on one of four distinct quick-fixes to respond to smaller shocks but not large shocks. Finally, we incorporate the empirically identified forms of near-rationality into a quantitative model to study their micro- and macroeconomic implications. Even at negligible utility costs for the individual household, these quick-fixes significantly alter aggregate consumption responses to shocks and can account for the puzzlingly weak empirical relationship between MPCs and (liquid) wealth.

Formalizing Quick-Fixing. To fix ideas and motivate our empirical strategy, we begin with a simple model. We consider a household living for two periods that decides how much to consume and save. A *frictionlessly optimizing* household optimally chooses consumption


and savings to maximize payoffs. This defines an optimal policy function that maps income states to consumption levels.

A *near-rational* household instead begins with a simple, potentially suboptimal policy function that we call a *quick-fix*. For example, the household might set consumption at a reference level, letting savings adjust to shocks, or instead set savings at a reference level, letting consumption adjust to shocks. The household derives a small convenience from using the quick-fix instead of implementing the perfectly optimal plan, perhaps because the quick-fix requires less thinking, allows one to stick to established habits, helps manage psychological urges, or avoids physical or mental transaction costs. The household quick-fixes if the payoff loss from doing so is lower than these costs and frictionlessly optimizes otherwise.

While the model does not specify which quick-fixes households use—this is precisely the problem we seek to address—it delivers three predictions that allow us to identify quick-fixes empirically: (i) quick-fixes are prevalent for small shocks, where payoff losses are low, but (ii) they are abruptly and discontinuously abandoned for large shocks, and (iii) after abandonment, households that would otherwise use different quick-fixes behave similarly. Hence, a key implication of the model is that the distribution of household-level *consumption policy functions*—consumption responses to income shocks of different sizes—is the critical moment to identify quick-fixing behavior and evaluate the model against alternatives. These data can reveal what (if any) quick-fixes people use and what triggers people to abandon them. We further demonstrate that one cannot measure the distribution of policy functions in observational data absent strong assumptions that are implausible in consumption-savings settings. Intuitively, background shocks to spending and income contaminate measures of *individual* propensities to consume, and such data can at most reveal how households respond to a single shock. This necessitates a survey-based research design that elicits responses to many shocks of different sizes while holding background noise fixed.

Empirical Evidence. We therefore design a survey to elicit household consumption policy functions. We administered this survey to about 5,000 US households in October and November of 2023. The sample approximates the US adult population in terms of gender, age, income, education, and region. We ask respondents how they would adjust their spending and saving over the next three months in response to multiple hypothetical shocks in random order: seven unexpected one-time payments ranging from \$50 to \$10,000 and seven unexpected one-time income losses of the same amounts. Collecting this *within-respondent* information differentiates our study from existing work using surveys to measure the marginal propensity to consume and allows us to detect and understand near-rational behavior. Our measured MPCs align with those reported in previous observational studies, consistent with evidence that surveys can accurately capture households’ MPCs (see, *e.g.*, [Parker and Soule-](#)

Figure 1: The “bowtie” MPC distribution across shock sizes



Notes: The alluvial graph summarizes the MPC data of the 4,981 US households in our sample. Each of the 14 columns displays the distribution of MPCs for one particular shock size, with colors indicating the size of the MPC. The streams between bars indicate how households’ MPCs transition between two neighboring shocks (here, we only distinguish between MPCs of 0, MPCs of 1, and interior MPCs to keep the figure readable). Black dots depict the average MPCs for each shock.

les, 2019; Colarieti, Mei and Stantcheva, 2024; Kotsogiannis and Sakellaris, 2025).

We document five new facts. Figure 1 visualizes the first fact: the “bowtie” distribution of MPCs across shock sizes. For small shocks, households frequently report an MPC of 0 (*i.e.*, fully saving a gain or drawing from savings after a loss) or 1 (*i.e.*, fully consuming a gain or cutting back consumption after a loss). As shocks get larger in absolute value, the fraction of extreme responses declines, and the fraction of interior responses increases. This transition from extreme to interior MPCs, visualized by the shaded flows, generates the “bowtie” appearance of Figure 1. Moreover, we observe very similar behavior among households across the wealth distribution. The bowtie pattern is at odds with the predictions of standard consumption-savings models. For example, we find that as shocks become *more* negative, *fewer* households report having a high MPC. Instead, the large difference in MPCs for small and large shocks is reminiscent of quick-fixing.

Our second fact is that the distribution of MPCs is well described by a decomposition of households into four *quick-fixing types*. These households use extreme MPCs of 0 or 1 as heterogeneous quick-fixes for small shocks, but abruptly abandon them for larger shocks, and thereafter behave homogeneously. Motivated by our theoretical framework, we categorize households into types by considering how they respond to the smallest shocks (the \$50 loss and gain), as this is when quick-fixing is most likely to be prevalent. We say that households

are *consumption fixers* if they have a zero consumption response to both shocks; *savings fixers* if they have a zero savings response to both shocks; *consumption prioritizers* if they increase consumption after windfalls but cut savings after losses; and *savings prioritizers* if they increase savings after windfalls but cut consumption after losses. These types span 68% of our respondents. Using within-respondent data, we show that they satisfy all formal properties of quick-fixes: (i) they adopt categorically different policies for small shocks, (ii) abruptly transition to interior MPCs for large shocks, and (iii) have similar MPCs conditional on transitioning to an interior MPC. The remaining 32% of “uncategorized” respondents do not exhibit clear quick-fixing behavior and usually report an interior MPC that is relatively stable for different shock sizes.

Our third fact suggests that quick-fixing helps to explain the empirical puzzle that MPCs are only weakly predicted by households’ financial situation and demographics. In our data, cross-sectional variation in spending, income, income risk, liquid and illiquid wealth, debt, education, age, gender, and household size explains only 13% of the household-level variation in MPCs. This poses a challenge for standard models in which such variables should explain all MPC variation. By contrast, quick-fixing types account for a significantly larger share, 35%, of the household-level variation in MPCs. This comes while quick-fixing types are essentially unpredictable: the same economic and demographic characteristics predict quick-fixes with R^2 values between 2% and 6%. Moreover, consistent with the near-rational theory, the predictive power of traditional economic factors increases with the magnitude of shocks.

Our fourth fact highlights that quick-fixing accounts for a large part of the variation in aggregate MPCs across small and large shocks (size-dependence) as well as gains and losses (sign-dependence). In particular, quick-fixing behavior increases the size-dependence of aggregate MPCs (the maximum absolute difference in aggregate MPCs across different shock scenarios) by a factor of 2–3.5.

Finally, our fifth fact supports the proposed mechanism that households quick-fix because of its convenience. In an additional survey, we elicit households’ responses to shocks alongside measures of how carefully they would consider their decisions, how likely they are to assess their overall financial situation when making a decision, and how likely they are to consult another household member. When households quick-fix and have extreme MPCs, they are substantially less likely to report deliberating.

The Economic Implications of Quick-Fixing. In the final part of the analysis, we leverage our novel data and embed the empirically revealed quick-fixing patterns in an otherwise standard incomplete-markets consumption-savings model. Households in the model are one of five types: the four different quick-fixing types or frictionless optimizers. The fraction of each type in the model corresponds to the fraction of each type that we empir-

ically categorize. We estimate the costs of changing behavior for each type to rationalize the switching patterns from extreme MPCs to the interior that we observe in the data. We calibrate the remaining preference and income process parameters to standard values, and we estimate the discount factor to match the MPCs of uncategorized households.

The model can rationalize observed behavior given very small opportunity costs of quick-fixing: \$16 per quarter per household, on average. Thus, as hypothesized, quick-fixing behavior is near-rational. Moreover, despite the fact that quick-fixing gives rise to very different responses to income shocks, everyone in our economy accumulates an almost identical amount of wealth. That is, quick-fixers who save do not become rich and quick-fixers who spend do not become poor, because large shocks arrive often enough to trigger reoptimizations. This quantitative finding replicates a non-targeted moment from our survey, the inability of financial status to predict quick-fixing types, and stands in stark contrast to the predictions of canonical “spender-saver” models.

Finally, we discuss three macroeconomic implications of empirically disciplined quick-fixing behavior. First, targeting transfers to poor households is less effective than in the nested incomplete markets model with perfect optimization, and can even backfire. For example, the MPC out of a \$500 transfer is larger for households that hold two years’ worth of income in assets than for those that only have one quarter’s worth of income in assets.

Second, compared to the incomplete-markets benchmark, quick-fixing generates substantially more size-dependence (small shocks lead to larger aggregate MPCs than large shocks) and sign-dependence (negative shocks lead to larger aggregate MPCs than positive ones). For example, a \$1,000 income loss generates a static Keynesian (fixed price) multiplier that is 30 cents larger on the dollar with quick-fixing than under the frictionless benchmark.

Third, because of quick-fixing, macroeconomic shocks of the same size but with different *incidence* induce markedly different aggregate responses. For example, the aggregate MPC out of a \$250 shock that affects everybody equally (which is equivalent in aggregate to a 1.6% change in quarterly household income) is more than twice as large as an equally sized aggregate shock that arises from only 2.5% of the population experiencing a \$10,000 shock. This matters because macroeconomic shocks vary widely in their incidence: recessions usually concentrate income losses on a few people who become unemployed, while fiscal stimulus or inflation may be more uniformly experienced.

Taken together, these results suggest that applying a “one-size-fits-all” MPC across these contexts may be misleading when households are near-rational. Moreover, contrary to received wisdom, quick-fixing implies that relatively small and *untargeted* stimulus might be most effective. This is because observable financial status ineffectively screens for near-rational spenders, and smaller amounts are more likely to trigger quick-fixing behavior.

Related Literature. Our article speaks to a large literature on households’ consumption and savings behavior and its aggregate consequences. However, no previous work has used a tailored empirical design to characterize a near-rational model of consumption and savings, contrast it with leading alternatives, empirically discipline the rules that people use, and quantify the macroeconomic implications of the empirically implied near-rational behavior.

Our paper is most closely related to the literature exploring how bounded rationality shapes households’ consumption and savings. Existing studies focus on, for example, mental accounting (*e.g.*, Kőszegi and Matějka, 2020; Lian, 2021), hyperbolic discounting (*e.g.*, Ganong and Noel, 2019; Maxted, Laibson and Moll, 2024), deviations from full-information rational expectations (*e.g.*, Coibion, Georgarakos, Gorodnichenko, Kenny and Weber, 2024; Coibion and Gorodnichenko, 2026; Weber, D’Acunto, Gorodnichenko and Coibion, 2022; Bordalo, Gennaioli, Ma and Shleifer, 2020), sparsity (*e.g.*, Gabaix, 2014, 2023; Eichenbaum, Guerreiro and Obradovic, 2025), inattention (*e.g.*, Maćkowiak and Wiederholt, 2015; Pfäuti and Seyrich, 2024), memory (*e.g.*, Ilut and Valchev, 2023, 2024), and learning from experiences (*e.g.*, Malmendier and Nagel, 2011). Our approach adopts an agnostic stance by allowing data to reveal the dominant near-rational rules that households follow. Moreover, our model of quick-fixing aligns well with recent research in behavioral economics showing that individuals often resort to simple solutions when faced with complex problems (Bordalo, Conlon, Gennaioli, Kwon and Shleifer, 2025; Gabaix and Graeber, 2024; Oprea, 2024).

Surveys are our empirical method because they allow us to simultaneously measure multiple MPCs for each household, which we show is necessary to evaluate near-rational theories. A large literature has used surveys to understand heterogeneity in MPCs (Christelis, Georgarakos, Jappelli, Pistaferri and Van Rooij, 2019; Coibion, Gorodnichenko and Weber, 2020; Eichenbaum et al., 2025; Fuster, Kaplan and Zafar, 2021; Jappelli and Pistaferri, 2014, 2020). Moreover, recent work shows that survey-based measures of MPCs align closely with observational evidence (Parker and Souleles, 2019; Colarieti et al., 2024; Kotsogiannis and Sakellaris, 2025), and our data reproduce many of the basic patterns identified in observational studies. However, no previous study systematically maps households’ consumption policy functions across a wide range of shocks, so the key empirical patterns that we uncover are new. Fuster et al. (2021) identify one similar regularity in their data, a mass of agents with $MPC = 0$, and attribute this to “inertia” in adjusting consumption. Our approach of quick-fixing is conceptually distinct for two reasons. First, inertial consumption behavior (in our language, consumption fixing) is only one of several quick-fixes, accounting for just 14% of households in our data. Second, our identified distribution of quick-fixes gives rise to different macroeconomic implications relative to inertial behavior, such as population-wide MPCs that decline (as opposed to increase) in shock size. More generally, our approach disciplines the “how” of

near-rationality, which is essential for accurately capturing its macroeconomic implications.

Within the literature initiated by Akerlof and Yellen (1985a,b) that studies the near-rationality hypothesis, Cochrane (1989) and Krusell and Smith (1996) quantitatively evaluate the losses from various rule-of-thumb consumption rules and find that they are small. Browning and Crossley (2001) and Kueng (2018) suggest that empirically observed deviations from consumption smoothing could arise due to their low opportunity costs. This work highlights the challenge and starting point of our analysis: the set of near-optimal behaviors is vast, so progress in testing for near-rationality and assessing its implications requires identifying which forms of near-rationality are empirically relevant. We tackle this challenge by empirically measuring prevalent near-rational consumption-savings rules through a tailored survey and studying the macroeconomic implications of the implied behavior.

2 A Simple Model of Quick-Fixing

We describe quick-fixing in a simple two-period model of consumption and savings and use the model to motivate our survey-based strategy for measuring quick-fixing behavior.

2.1 Setup


A household has payoffs $u(c_1) + \beta u(c_2)$ over consumption streams (c_1, c_2) , where u is a concave, increasing, and differentiable utility function, and $\beta > 0$ is a discount factor. The household earns income (y_1, y_2) in each period and earns return R on savings. At $t = 1$, they choose how much to consume and save. This setup lets us describe quick-fixing in the simplest possible way. Our later quantitative model adds borrowing constraints, income risk, and an infinite horizon.

A *frictionlessly optimizing* household implements the consumption and savings plan that maximizes lifetime utility. In particular, they equate marginal utility from consumption in both periods, $u'(c_1) = \beta R u'(c_2)$. We denote this household’s *consumption policy function*, or mapping from income states to consumption choices, as $c_1 = c^*(y_1, y_2)$.

A *near-rational* household might not use the plan described above because they have an alternative that is “good enough.”¹ Such a household, by default, sets its consumption via a *quick-fix*: a policy function $c_1 = c^q(y_1, y_2)$ that differs from c^* . Deviating from this quick-fix

¹We think of this as a reduced-form representation of behavior under psychological constraints, even though it raises the classic problem of infinite regress in how agents should decide “how to decide how to decide how to ...” (Lipman, 1991). This approach is standard in economics (*e.g.*, in models of rational inattention, Maćkowiak, Matějka and Wiederholt, 2023) and in cognitive science (*e.g.*, in the resource-rational framework, Lieder and Griffiths, 2020). As in Krusell and Smith (1996), one simple interpretation of our costs is that they are the costs of changing or implementing behavior.

Figure 2: An illustration of quick-fixing



costs κ_q in utility. This friction captures any conveniences in following the quick-fix: for example, saving cognitive resources, sticking to established habits, managing psychological urges, or avoiding physical or mental transaction costs. Regardless of the precise nature of κ_q or the deep origins of κ_q , behavior follows a common pattern: the household acts “rationally” if the relative benefit from doing so exceeds κ_q ; otherwise, they follow their quick-fix.

2.2 Quick-Fixing Behavior

We illustrate the implications of this theory in an example (Figure 2). We set $u(c) = \log c$, $y_2 = \bar{y}$, and $R = \beta = 1$. We think of $\bar{y}$ as the household’s “usual” income, and any deviation at $t = 1$ as a “shock.” In this example, a household’s quick-fix is to fully spend any shock, corresponding to a marginal propensity to consume (MPC) of 1. Since the household’s savings are held constant, we call this “savings fixing.” This is represented by the policy function $c^q(y_1) = y_1$ (blue dashed line in Panel A). By contrast, the frictionless policy function is $c^*(y_1) = \frac{y_1 + \bar{y}}{2}$ (gray solid line in Panel A).

The payoff loss from following this quick-fix, illustrated in Panel B, is zero when there is no shock and close to flat in small consumption differences. This is a more general result, also noted by [Akerlof and Yellen \(1985a\)](#) and [Cochrane \(1989\)](#): the loss from deviating from an optimal policy function is second-order in the difference in consumption due to the envelope theorem. Thus, the losses from quick-fixing can be economically small, even when the corresponding change in the consumption response is large. For example, in our simple calibration, a 5% shock to income leads to an MPC that is twice as high and a consumption difference of 2.5 percentage points, the equivalent of 0.03% of lifetime consumption.

The household’s behavior, illustrated in Panel C, features a jump in the level of consumption and the MPC. This jump occurs when income deviates sufficiently from $\bar{y}$ for the household to find it worthwhile to abandon their quick-fix.

This basic logic applies to many quick-fixes, including those that imply very different

marginal behavior. In the same figure, we also illustrate a quick-fix of zero spending response to income shocks (“consumption fixing”): that is, $c^q(y_1) = \bar{y}$ (Panel A). Under our simplifying assumptions, the loss is almost identical to the loss from savings fixing because the gap between the optimal and quick-fixing consumption response has the same size (Panel B). Behavior again features a “jump” in consumption and MPCs. Critically, the economic implications are the opposite: large shocks trigger an abrupt *increase* in the MPC (Panel C). This demonstrates that near-rational alternatives can lead to drastically different outcomes while achieving identical payoff losses.

We formally analyze the simple model and illustrate that this basic logic applies to a large class of quick-fixes in Appendix A.1. Quick-fixes have second-order opportunity costs but generate first-order differences in consumption and MPCs. It follows that assessing the macroeconomic implications of near-rationality necessitates identifying *which* quick-fixes are empirically relevant.

2.3 From Theory to Measurement

The model makes three testable predictions that can help identify quick-fixes empirically: (i) households quick-fix in response to small shocks, (ii) they abandon their quick-fix once shocks are big enough, and (iii) they behave the same way conditional on abandoning their quick-fix, regardless of differences in their quick-fixes. Hence, the critical moment that empirically disciplines the near-rational theory is the *distribution of consumption policy functions*: how households’ consumption choices respond to income shocks of different sizes.

The Limits of Observational Data. Despite their economic importance, individual consumption policy functions cannot be recovered from observational data for two reasons.

First, observational data make it difficult to isolate a household’s response to a shock from background noise in income and spending, which can be substantial (Morduch and Schneider, 2017; Ganong, Noel, Patterson, Vavra and Weinberg, 2025b). Even in an ideal experiment, where a household is randomly assigned to receive a transfer or not, an analyst cannot tell how much of the observed consumption change reflects the response to the transfer itself and how much reflects a response to unrelated fluctuations, such as changes in hours worked, side-job income, the timing of vacations, or unexpected car repairs. This is especially problematic for identifying quick-fixes, as background shocks blur mass points in the response distribution and can make a small number of discrete quick-fixes look like a smooth distribution of MPCs. In Appendix A.2, we use a simulation to illustrate this mechanism. We also formalize its connection to a known econometric issue: the distribution of treatment effects (here, MPCs) cannot be non-parametrically identified absent strong economic assump-

tions (Abbring and Heckman, 2007) that are particularly restrictive in consumption-savings settings. For example, the assumptions under which quantile regressions or deconvolution methods identify the MPC distribution are violated in our framework. Thus, state-of-the-art observational methods can fail to detect MPC heterogeneity that does exist and spuriously detect MPC heterogeneity that does not exist.

Second, *even if* we could observe how a household responds to one shock at one point in time, we would still not observe the relevant counterfactual: how that same household would have responded to a larger or smaller shock. This is necessary to recover the household-level consumption policy functions needed to detect near-rational behavior.

Surveys Offer a Way Around These Problems. Surveys have two key advantages. First, they let us directly elicit responses to the shock of interest, effectively holding background fluctuations in income and spending fixed. Second, they let us measure multiple MPCs for different shock sizes within the same household. These properties allow us to elicit individual-level MPCs for income shocks of various sizes. Moreover, prior work shows that survey-based estimates of average consumption responses to transitory income shocks match those obtained from observational data (Parker and Souleles, 2019; Colarieti et al., 2024; Kotsogiannis and Sakellaris, 2025), making surveys a natural tool for our purposes.

3 Empirical Evidence

This section presents the design and results of a novel, large-scale household survey tailored to uncover households’ policy functions and quick-fixes. We find that the majority of households apply simple rules that fully consume or fully save out of sufficiently small shocks, and we identify four empirically prevalent quick-fixes (Facts 1 and 2). The quick-fixes that households use are poorly explained by financial and demographic characteristics, yet highly explanatory of MPC heterogeneity across households (Fact 3). Moreover, the presence of these behaviors critically shapes aggregate MPCs (Fact 4). Finally, consistent with our proposed mechanism, households report that they would be less likely to carefully consider their behavior when facing small shocks and using a quick-fix (Fact 5). We argue that these patterns are consistent with near-rationality but not with the predictions of leading alternative models.

3.1 Survey Design

We conducted our survey with 4,981 US households in October and November 2023, collaborating with the survey company Bilendi. The sample approximates the adult US population

in terms of gender, age, income, education, and region, and broadly captures the wealth distribution across the country (see Appendix C.1 for details).²

Our main goal is to uncover households’ consumption policy functions and, in particular, their responses to small versus large shocks. Hence, we collect detailed *within-household* information on responses to fourteen shocks of different sizes and signs: gains and losses with a magnitude of \$50, \$100, \$250, \$500, \$1,000, \$5,000, and \$10,000. We focus on one-time shocks, rather than other types of income or wealth shocks, to ensure maximum comparability with the literature. We include very small shocks, such as \$50 or \$100, to empirically identify the distribution of quick-fixes.

Our approach builds on a standard and well-studied procedure to elicit spending and savings responses to one-time income shocks (*e.g.*, Jappelli and Pistaferri, 2014, 2020; Christelis et al., 2019; Fuster et al., 2021; Colarieti et al., 2024). We first provide households with short definitions of consumption and saving. We refer to consumption as “spending” to follow common parlance, and we explicitly stress that we consider debt repayment as part of saving. Next, households are asked to think about an unexpected one-time income gain or loss. For example, respondents read:

Consider a hypothetical situation where your household unexpectedly receives a one-time payment of \$1,000 today.

Then, households answer how they would change their spending and saving in response to the shock. They can respond in two numeric open response fields, and we calculate their MPC by dividing their spending response by the income shock.³

How would this one-time extra income cause your household to change its spending and saving over the next three months?

Increase in spending:	\$ _____
Increase in saving:	\$ _____

Appendix C.2 presents the key instructions.⁴

We require consumption and savings responses to sum to the amount of the shock, which ensures that respondents report the cumulative consumption and savings change over the three months after the shock. We randomize both whether respondents first face gains or

²We slightly oversample respondents with college education and respondents with lower total debt and lower illiquid wealth, but results are robust to re-weighting and correcting for these imbalances (Figure B.1).

³We elicit increases in spending and saving for income gains and decreases for income losses. We explain to respondents that they can enter negative numbers to indicate changes in the reverse direction.

⁴The full instructions are available online at <https://osf.io/2s7cf>.

losses and the order of shock magnitudes to avoid any bias or learning effects from the order in which respondents are faced with shocks. We also preclude respondents from going backward in the survey to adjust previous answers. This guards against the possibility that respondents revise the profile of their responses in light of choices made later. In practice, we observe similar results across the randomly assigned question orders (Figure B.1).

The survey closes with a series of detailed demographic and economic background questions, including questions on income, income risk, future income expectations, and wealth (liquid and illiquid). We keep the survey short to avoid response fatigue. The median response duration is approximately 14 minutes, and most respondents complete the survey within 9–24 min (20%–80% quantile range).⁵

Statistical Precision. Due to the large sample of nearly 5,000 respondents and 70,000 MPCs, our estimates are highly precise. For instance, the 95% confidence interval for the estimated average MPC has a width of 0.013. Likewise, the margin of error for population share estimates is 1.4 percentage points, meaning that for any given percentage share (X%) of respondents, a (conservative) 95% confidence interval would be $[X\% - 1.4\%, X\% + 1.4\%]$.

Comparison of Cross-Sectional Results to Previous Work. Prior work shows that survey evidence on the consumption response to transitory income shocks aligns well with observational evidence (Kotsogiannis and Sakellaris, 2025; Colarieti et al., 2024; Parker and Souleles, 2019). In our own data, we also verify that several basic patterns in MPCs line up with prior work, both survey and observational (see Appendix C.4). First, average MPCs are relatively high: for example, the one-quarter MPC out of a \$1,000 gain is 0.35. This is comparable to observational estimates, for example, an MPC of about 0.30 out of the 2008 US tax rebate, as estimated in analysis by Borusyak, Jaravel and Spiess (2024) and Orchard, Ramey and Wieland (2025) that revisits prior analysis by Parker, Souleles, Johnson and McClelland (2013) and Broda and Parker (2014). Second, average MPCs decline in the size of shocks: for example, the MPC out of gains is 0.49 for \$100, 0.35 for \$1,000, and 0.30 for \$10,000 ($p < 0.001$ for all comparisons). This is consistent with observational evidence: Fagereng, Holm and Natvik (2021) find that MPCs out of smaller lottery winnings are higher, and Ganong, Jones, Noel and Sullivan (2025a) find that MPCs out of smaller US stimulus payments in 2008 are higher. Third, the average MPC is larger for losses than for gains ($p < 0.001$), as documented in previous surveys, *e.g.*, Bunn, Le Roux, Reinold and Surico (2018), Christelis et al. (2019), and Fuster et al. (2021). Finally, MPCs vary widely across households, consistent with observational evidence from Lewis, Melcangi and Pilossoph (2024).

⁵We obtain similar results if we restrict attention to only the first or first five MPCs each respondent reports or drop the 50% of the sample who read the preparatory instructions most quickly (App. Fig. B.1).

3.2 Extreme MPCs and the “Bowtie” Pattern

Our first key finding is the “bowtie” pattern in the distribution of MPCs for different shock amounts, which we illustrated in Figure 1 of the introduction. Extreme MPCs of 0 and 1 are common for small shocks but rare for large shocks. For example, 74% of households consume every dollar ($\text{MPC} = 1$) or save every dollar ($\text{MPC} = 0$) of the \$50 gain, but only 26% do so for a \$10,000 gain ($p < 0.001$). Likewise, 72% of households fully reduce their consumption ($\text{MPC} = 1$) or fully reduce their savings ($\text{MPC} = 0$) in response to the \$50 income loss, but only 32% do so for a \$10,000 loss ($p < 0.001$). Consequently, the mass of households with an interior MPC strictly between 0 and 1 increases from small to large shocks, giving Figure 1 its bowtie-like appearance.

This pattern is largely independent of respondents’ current financial situation. Figure 3 displays the distribution of MPCs conditional on different levels of liquid wealth (*i.e.*, assets held in cash, checking or savings accounts, or easily accessible (non-retirement) investment accounts). Even households whose liquid wealth is more than ten times the size of the largest shocks considered (right-most panel) commonly respond to small shocks by fully adjusting consumption or savings, and they become less likely to do so as the shocks become larger. We find the same bowtie pattern across the distribution of illiquid wealth, debt, net wealth, and the ratio of liquid wealth to income (Figure B.2).

Fact 1 (The Bowtie): Many households respond with extreme MPCs of either 0 or 1 to small shocks and transition to interior MPCs for larger shocks. This pattern holds regardless of households’ current liquid and illiquid wealth.


3.3 Four Quick-Fixes

This drastic difference in MPCs for small and large shocks is reminiscent of quick-fixing. Using our within-household data, we identify four quick-fixes that underlie this pattern.

Candidates. To determine candidates for the quick-fixes households use, we isolate their behavior for the smallest shocks: the \$50 loss and gain. The theory suggests that these smallest shocks are most informative about quick-fixing behavior, but we will shortly demonstrate that quick-fixes are prevalent for much larger shock sizes. We find that 68% of respondents report an extreme MPC for *both* of these shocks, and we exhaustively categorize them into four types:

Consumption fixers (14% of households) respond with an MPC of 0 to small gains and losses, keeping consumption fixed and absorbing shocks with their savings.

Figure 3: MPC profiles across the distribution of liquid wealth



Notes: The alluvial graphs summarize the MPC profiles of households with varying liquid wealth (see Appendix C.3 for variable definitions). In each panel, each of the 14 columns displays the distribution of MPCs for one particular shock size, and the streams between bars indicate how households' MPCs transition between two neighboring shocks.

Savings fixers (29% of households) respond with an MPC of 1 to small gains and losses, keeping their savings fixed and absorbing shocks with their consumption.

Consumption prioritizers (11% of households) fully draw on savings to cover a small loss (MPC = 0), but they fully spend a small gain (MPC = 1).

Savings prioritizers (14% of households) fully cut back on consumption when they face a small loss (MPC = 1), but they fully save a small gain (MPC = 0).


The remaining group of **unclassified households (32%)** cannot be assigned to any of the four groups above because they respond with an interior MPC to even the smallest income gains or losses. We do not find clear traces of alternative candidates for quick-fixing behaviors among this group. Unclassified households almost never choose extreme MPCs and instead immediately adopt relatively stable interior MPCs.⁶

Evaluating the Properties of Quick-Fixes. The four widespread behaviors described above—fixing consumption, fixing savings, prioritizing consumption, and prioritizing savings—are intuitively plausible quick-fixes. Now, we show that they satisfy the three formal properties of quick-fixes that we identified in Section 2.

First, households use these behaviors to respond to small shocks but abandon them for larger shocks. To visualize this, Figure 4 decomposes the distribution of MPCs by group. For example, Panel A plots the MPCs of consumption fixers for all fourteen shocks. By construction, 100% of consumption fixers start with an MPC of 0 for \$50 income gains or losses. The figure reveals that respondents are increasingly less likely to have an MPC of 0 as the gains or losses become larger. We see this monotonic decline in adherence to the

⁶Their most common response to the smallest income shocks is an MPC of 0.5, which a total of 12% of households adopt. However, most of these households also choose an MPC close to 0.5 for larger shocks, so we cannot identify a clear transition pattern. No other consumption response is chosen by more than 5% of households, implying that other quick-fixes—to the extent that they exist—are not very prevalent.

Figure 4: Responses to income shocks for different quick-fixing types




Notes: The alluvial graphs summarize the MPC data of four different quick-fixing types, which we define on page 13, and unclassified respondents. In each panel, each of the 14 columns displays the distribution of MPCs for one particular shock size, and the streams between bars indicate how households' MPCs transition between two neighboring shocks.

quick-fix as shocks get bigger for all four types. Nonetheless, we note that an economically significant fraction of households *do* maintain extreme behaviors for relatively large shocks: for instance, 56% of savings prioritizers report an MPC of zero for a \$1,000 income gain.

Second, households tend to abruptly abandon the simple response once a critical shock size is reached. By this, we mean that shifts from extreme to interior MPCs rarely occur gradually. Households that start with an MPC of 1 tend to immediately jump from this extreme MPC to an interior MPC that is typically around 0.25 to 0.60 (the 20%–80% quantile range). Their first interior MPCs are thus not unlike those of households who start from the opposite MPC of 0, whose 20%–80% quantile range is 0.20 to 0.50. Figure 5 visualizes the abrupt transition from extreme MPCs (left column) to first interior MPCs (right column) across types for gains.

Third, once households abandon their simple response, their consumption policies are relatively stable *within* respondent and similar *across* respondents, even though their MPC

Figure 5: Distribution of MPCs for gains before and after switching to interior



Notes: The histograms show conditional distributions of MPCs for gains. The rows correspond to the four quick-fixing types. The first column shows the distribution of MPCs before households switch to an interior MPC, which by construction puts all mass at either $\text{MPC} = 0$ or $\text{MPC} = 1$. The second column shows the conditional distribution of MPCs (given type and shock size) for the first shock for which the respondent reports an interior value. An analogous analysis for losses is reported in Figure B.3.

for the smallest shocks differs radically. In the survey, the average absolute difference between two interior MPCs of adjacent shock sizes is 0.14, while the average conditional on a transition from the extremes (0 or 1) to the interior is 0.41. Moreover, households rarely (if at all) transition *back* from the interior to the extremes: conditional on starting from an interior MPC, households stay in the interior for 93% of shock size increases. In total, we observe a transition from interior to extreme MPCs only for 3% of shock size increases. Finally, interior MPCs are highly similar across respondents compared to MPCs as a whole: the variation in interior MPCs contributes only 16% to the total variance in MPCs. Figure 5 shows that, while MPCs vary widely across types before switching (left column), their interior MPCs after switching are similar (right column).

Thus, we conclude:

Fact 2 (Four Quick-Fixes): The majority of households can be categorized as one of four quick-fixing types—consumption fixers, savings fixers, consumption prioritizers, and savings prioritizers—who vary in their extensive margin response to small shocks. Households of all four quick-fixing types tend to abruptly transition from having extreme MPCs for small shocks to having similar interior MPCs for large shocks.

Discussion of Response Noise. Some response noise is inevitable in survey data, so we examine whether noise could obscure or falsely generate quick-fixing in Appendix C.5. We focus on the prediction that quick-fixing households should transition to an interior MPC at most once and remain interior thereafter, while unclassified households should never report extreme MPCs. We find that 52% of households behave fully consistently, 71% deviate from the one-switch pattern at most once, and 83% deviate at most twice. Given that we elicited MPCs for 14 different shocks *in random order*, the level of consistency in households’ behavior is high and extremely unlikely to be due to chance. Moreover, we illustrate in a simulation that even modest and empirically realistic levels of response error can explain the imperfect consistency levels that we observe across 14 different questions.

3.4 Quick-Fixes Account for MPC Heterogeneity

We next explore how well quick-fixes can account for MPC heterogeneity at the household level. Differences in MPCs across households are notoriously hard to predict using observable characteristics of households’ wealth, income, and demographics (*e.g.*, Fuster et al., 2021; Lewis et al., 2024). To study this in our data, we model respondents’ average MPCs as a function of observables in the following simple regression:


$$\text{MPC}_i = \alpha + \beta' X_i + \varepsilon_i \tag{1}$$

The outcome MPC_i is each respondent’s average MPC across the 12 gain and loss scenarios ranging from \$100 to \$10,000. We exclude the \$50 gain and loss cases, as they are used to define quick-fixing types.

We consider two candidates for the explanatory variables X_i . First, we consider a full suite of demographic and financial characteristics: households’ monthly spending (log), income (log), and income risk; dummies for low, intermediate, or high levels of liquid wealth, illiquid wealth, and debt, respectively; and four additional demographic characteristics (a college dummy, age, gender, and household size). Second, we consider dummy variables for the four quick-fixing types, which only depend on households’ responses to the \$50 shocks.

Echoing findings in the literature, the demographic and financial variables together explain only 13% of the cross-sectional variation in MPCs (see golden bars in Panel A of Figure 6 and Table B.1). In analogous regressions, we find that observable characteristics explain only 21% of variation in the household-level share of MPCs that equal 0 and—what is perhaps most striking—only 3% of the household-level share of MPCs that equal 1. Even though measurement error in our explanatory variables will somewhat attenuate the estimated R^2 , this finding remains surprising since standard theories suggest that these

Figure 6: Quick-fixing accounts for MPC heterogeneity but is hard to predict



Notes: Panel A displays the R^2 from regressions of households' average MPCs and their share of extreme MPCs on (i) their wealth, income, and demographics (golden bars) or (ii) their quick-fixing types (red bars). The average household-level MPC and the MPC shares are derived using all gains or losses ranging from \$100 to \$10,000. See Table B.1 for the complete list of variables and coefficients. Panel B displays the R^2 for MPCs out of small, medium, and large losses or gains. Panel C displays the R^2 from regressions of households' quick-fixing types on economic and demographic characteristics. See Table B.2 for the complete list of variables and coefficients.

characteristics should explain the majority of (or even all) heterogeneity in MPCs.

By contrast, the four quick-fixing types account for about a third of the variation in MPCs by themselves (35%, see red bars in Panel A of Figure 6 and Table B.1). The types also account for 44% and 41% of the variation in the share of extreme MPCs of 0 or 1, respectively. These findings are surprising because, although quick-fixing types are identified using data on MPCs, they are a very low-dimensional summary that uses information only about \$50 shocks (which we drop when we derive households' average MPCs) and only uses information about whether these responses are extreme or interior. This suggests that households' quick-fixing types are an important component of MPC heterogeneity and may help unpack the "black box" of MPC heterogeneity.

The near-rational model implies that quick-fixes predict MPCs best for smaller shocks, while traditional economic factors become more predictive for larger shocks. This is indeed what we find, as Panel B of Figure 6 shows. For example, quick-fixes strongly predict MPCs from small gains (\$100 and \$250) with an R^2 of 41%, but less so ($R^2 = 19\%$) for larger gains (\$5,000 and \$10,000). Observable characteristics of wealth, income, and demographics have negligible explanatory power ($R^2 = 4\%$) for small gains (\$100 and \$250) but a higher explanatory power ($R^2 = 19\%$) for large gains (\$5,000 and \$10,000).

Quick-Fixes Are Hard to Predict. The importance of quick-fixing in accounting for MPCs raises the question: are households' quick-fixing behaviors themselves related to their

economic and demographic characteristics? In particular, different economic circumstances could drive households to adopt different quick-fixes, or different quick-fixes could change households’ economic circumstances (*e.g.*, by affecting wealth accumulation).

To investigate this, we regress indicators for being categorized as one of our four types (consumption fixer, savings fixer, consumption prioritizer, or savings prioritizer) on the same characteristics describing wealth, income, and demographics as before. We find only a weak relationship between household observables and quick-fixing types, with R^2 values ranging from 0.02 to 0.06, as displayed in Panel (C) of Figure 6 (see also Table B.2). These results are also consistent with our earlier observation that the distribution of MPCs looked similar when conditioning on different levels of wealth and debt (Figures 3 and B.2). We observe a clearer relationship between households’ characteristics and whether we can classify them as a quick-fixer at all ($R^2 = 0.17$, Column 5). For example, households with high income risk, intermediate liquid wealth levels, low illiquid wealth, and high education are more likely to adopt interior MPCs even for the smallest shocks.

A related question is whether we can predict households’ switching points, *i.e.*, the smallest shocks for which they switch from an extreme to an interior MPC. We find that neither the quick-fixing types nor the economic characteristics are accurate predictors of switching points ($R^2 = 0.02$ and $R^2 = 0.06$, respectively; Columns 6 and 7). Nonetheless, one consistent *qualitative* finding is that the wealthiest households have higher switching thresholds—that is, they are the most reluctant to give up their quick-fix (Figure B.4). In Appendix A.1, we observe that this is a natural consequence of our model if households have utility curvature that is declining in consumption levels (*prudence*).⁷


We summarize these findings below:

Fact 3 (Quick-Fixes Account for MPC Heterogeneity But Quick-Fixes Are Hard to Predict): The four quick-fixing types account for a large share of variation in average MPCs across households. Yet, quick-fixing behavior is essentially unpredictable from households’ economic and demographic characteristics.

This fact presents an apparent tension. How can quick-fixing simultaneously generate differences in consumption responses without generating a correlation between quick-fixing behaviors and wealth? The ability of our model to resolve this tension will constitute an important “non-targeted moment” in our later quantitative analysis.

⁷This also echoes Kueng’s (2018) finding that households exhibit higher MPCs out of shocks which are small *relative to their wealth*, which he explains using a near-rational model in which households’ default (in our language, quick-fix) is to be hand-to-mouth.

Figure 7: Quick-fixing types account for variation in aggregate MPCs



Notes: The left panel displays the average MPCs of four different quick-fixing types, which we define on page 13, and unclassified respondents for 14 different income shocks. The black line shows the aggregate MPC of the full household sample. The middle panel shows the same statistics after shutting down any variation in interior MPCs to isolate the effect of the extensive margin predicted by households' quick-fixing types. All interior MPCs are fixed at 0.45. Households start with the MPC prescribed by their type. Once they change their MPC, they permanently switch to the interior. The right panel graphs the difference between the left panel and the middle panel.

3.5 Quick-Fixes Matter for Aggregate MPCs

So far, we have established that quick-fixing is prevalent at the household level and helps explain household heterogeneity in spending behavior. We next illustrate that this behavior matters for average, population-wide MPCs.

We first observe that the average MPC measured in our survey varies significantly across shock scenarios (Figure 7, Panel A). In particular, average MPCs (black dots) decline in the magnitude of shocks for both losses and gains, and they are higher for losses than for gains. Underlying this are sharply different patterns for each of the quick-fixing groups and for the unclassified respondents (colored dots).

We next illustrate how quick-fixing affects average MPCs by way of a simple calculation that is visualized in Panel B. Specifically, we construct a predicted average MPC based on households' switches from extreme to interior MPCs: households start with the MPC prescribed by their identified quick-fix type (zero or one). Once these households switch to a different MPC for the first time, they permanently switch to an MPC of 0.45, which is the average interior MPC in our data. Thus, this calculation shuts down any variation in interior MPCs: MPCs only vary if households stop quick-fixing and switch to an interior MPC. Finally, Panel C shows the residual variation in aggregate MPCs, *i.e.*, the difference between Panels A and B due to varying interior MPCs.

The decomposition reveals that a large share of the sign- and size-dependence of aggregate MPCs is driven by households using quick-fixes for small shocks and abandoning them for

large shocks. For instance, the MPC difference across the smallest and largest gain in Panel B is 0.124, while the residual difference in Panel C is 0.136. Under this metric, quick-fixing increases the size-dependence in aggregate MPCs for gains by a factor of $(0.124 + 0.136)/0.136 = 1.9$. For losses, the factor is even larger at 3.5.


In theory, quick-fixes could lead to aggregate MPCs that are either decreasing or increasing in the absolute value of shock sizes. But our data reveal that quick-fixes reduce MPCs for larger shocks for two reasons. First, savings fixers—who transition from $\text{MPC} = 1$ to the interior—are more common than consumption fixers—who transition from $\text{MPC} = 0$ to the interior. Second, since the average interior MPC is below 0.5, abandoning an MPC of 1 has a larger impact than abandoning an MPC of 0. This result also illustrates why our quick-fixes are conceptually distinct from models of “inertial” behavior in which a household’s level of consumption features adjustment frictions, as in [Fuster et al. \(2021\)](#). Such adjustment frictions generate an aggregate consumption function that is increasing in the absolute value of shock sizes, whereas our data generate the opposite pattern. For this reason, our analysis underscores the need to flexibly identify consumption policy functions in the presence of near-rational alternatives to discipline the “wilderness” of irrationality and its emergent macroeconomic implications.

The residual variation driven by the intensive margin (Panel C) is much smaller and features MPCs for gains that are generally declining in shock size. This feature can be rationalized via standard incomplete markets models, which give rise to a concave consumption function. While this is not the focus of our analysis, our quantification in [Section 4](#) will show that our calibrated model can account for this residual variation.

Fact 4 (Quick-Fixing Matters for Aggregate MPCs): Quick-fixes account for a large part of the variation in aggregate MPCs. Moreover, quick-fixing increases the size- and sign-dependence in aggregate MPCs.

Policy and the Stability of the Cross-Sectional Distribution of Quick-Fixes. The preceding discussion indicates that the consequences of quick-fixing for aggregate consumption responses to shocks depend on the cross-sectional distribution of quick-fixing types. This naturally raises the question of how stable this distribution remains as the economic environment changes—a question that connects to [Lucas’s \(1976\)](#) classic critique of reduced-form consumption functions. [Fact 3](#) offers an answer: in the data, quick-fixes are unrelated to standard economic variables such as income and wealth. This suggests that policies could at most weakly influence the cross-sectional distribution of quick-fixing types by altering observable economic characteristics. Therefore, even if policy changes shift the cross-sectional distribution of economic characteristics (for example, by increasing household insurance),

Figure 8: Self-reported deliberation and extreme MPCs



Notes: Data from the additional deliberation study (517 US households). For fourteen different income shocks, the figure shows the average frequency with which households report to consider their response carefully (a score of at least four on a six-point scale, red line), assess their overall financial situation (yellow line), discuss the decision with other household members (purple line), or report extreme MPCs of 0 or 1 (black line). Households either face all seven positive or all seven negative income shocks.

they will not change the cross-sectional distribution of quick-fixing types. We further develop these ideas in our quantitative analysis in Sections 4 and 5.

3.6 Quick-Fixes Require Less Deliberation

Our notion of quick-fixing is based on the idea that households prefer quick and simple solutions for small shocks but are willing to invest more resources to fine-tune their response to large shocks. To more directly test this mechanism, we designed a follow-up survey that digs deeper into *how* households make their spending and savings choices. This survey sampled 517 US households in August 2024. We first measure households’ consumption policy functions for gains or losses, following the same procedure as in the main study. In addition, we ask each respondent to rate for each shock (i) how carefully they would consider how to change their spending and saving (on a six-point scale), (ii) what is the percent chance that they would assess and consider their household’s overall financial situation prior to deciding how to respond, and (iii) what is the percent chance that they would discuss their response with other household members.⁸ Our goal was to capture the multi-faceted ways in which making economic decisions might be costly: for example, through considering different plans, gathering relevant information, or generating agreement within the household. For brevity, we call all such behaviors “deliberation.”

⁸We recruit participants with the survey company Prolific. The sample is not quota-based and does not represent the US population, but we obtain identical results if we correct for sampling imbalances with post-stratification weights (see Table B.3). Appendix C.2 contains the additional survey instructions.

We find that each of these three measures of deliberation strongly increases with shock size, consistent with our interpretation of the quick-fixing model (Figure 8). For example, the likelihood that households assess their overall financial situation when deciding how to respond to an income shock is on average 32% for the smallest shocks but 88% for the largest shocks (yellow line, $p < 0.001$). Likewise, the likelihood that respondents consult other household members increases from 32% to 84% (purple line, $p < 0.001$). These increases in deliberation mirror the *decreasing* likelihood that households choose an extreme MPC of 0 or 1 (black line). At the household level, a one-standard-deviation higher deliberation score (on any of three measures) predicts a 25 pp *lower* chance that a household adopts an extreme MPC of 0 or 1 (Table B.3).⁹

We summarize these findings below:

Fact 5 (Quick-Fixing Requires Less Deliberation): Households more carefully consider choices, assess their finances, and discuss decisions amongst themselves when facing larger income shocks. Lower levels of deliberation come with a higher frequency of quick-fixing.

Additional Qualitative Evidence. We complement our evidence with a smaller qualitative survey to provide some first illustrative insights into why extreme MPCs may serve as convenient quick-fixes and require less deliberation. We ask approximately 500 households how they would respond to both a small shock (\$100) and a large shock (\$1,000). If they report a switch from an extreme to an interior MPC, we follow up by asking them to explain their behavior in their own words. We find four broad patterns in these qualitative data, which we discuss in more detail in Appendix C.6.

First, 86% of households refer to the difference in shock size to explain the difference in their MPCs. For example, one response states matter-of-factly (MPC = 0 for \$100): “One hundred bucks is not that much. It’s great, don’t get me wrong, but it’s something you either spend on a dinner or put away. Where we’re at right now, it’s going right in the bank.”

Second, households commonly mention habits and routines such as a fixed spending budget, a fixed monthly transfer to savings, or the goal to maximize savings. For example, one household summarizes (MPC = 0 for \$100): “I have a budget for a reason and generally stick to it unless there are major changes.”

Third, for small shocks, extreme MPCs appear to be easier to imagine, evaluate, and appreciate. By contrast, interior MPCs lead to small changes that are not perceived to

⁹We find a similar relationship in the response-time data from our main survey. When respondents deliberate longer and spend more time responding to a shock scenario, they are less likely to report extreme MPCs (Table B.4).

“make a dent” in households’ plans. Interestingly, some households make this case in terms of spending—“[\$100] is not really large enough to make an impact on our spending” ($MPC = 0$ for \$100)—while others make the same case in terms of saving—“the \$100 is not really enough to move the needle in saving” ($MPC = 1$ for \$100).

Fourth, many households recognize an income gain as a welcome opportunity to treat themselves or their families. As one respondent puts it, “Why not just use the unexpected \$100 to spend on something you can enjoy or something that can help you in the short-term?” Most balance consumption and saving for the large shock, but they approach the smaller \$100 gain differently. Some conclude that they can give in, “indulge”, and spend everything, while others choose to maintain “discipline” and save everything.

All in all, the qualitative evidence points to a multidimensional explanation of why households rely on quick-fixes. Extreme MPCs are convenient because they require little thought, preserve familiar routines, and discipline (or give in to) psychological urges. Our model of quick-fixing captures the convenience of extreme MPCs for small shocks and the transition pattern from extreme to interior MPCs, thus providing a theoretically plausible and practical description of household behavior.

3.7 Can Existing Models Rationalize These Data?

In this section, we summarize why alternative theories of consumption-savings behavior do not satisfactorily rationalize our empirical findings. The list of models we discuss here exhausts the set of models studied in a recent review article by [Kaplan and Violante \(2022\)](#) and includes many additional models. Our intention is not to dismiss these models, which may succeed in describing other dimensions of consumption-savings decisions. Rather, our point is that these theories cannot account for the novel patterns that we observe.

Incomplete Markets Models. The standard incomplete markets model ([Bewley, 1979](#); [Carroll, 1997](#)) cannot account for the “bowtie” pattern in [Figure 1](#). For example, this model predicts that as shocks become more negative, the fraction of high MPC households should increase as more households become borrowing-constrained. This is the opposite of the pattern that we see in the data: as shocks become more negative, households become less likely to have high MPCs. A similar point applies to many extensions of the model that feature *ex ante* heterogeneity in preferences and investment technologies among agents: for example, among discount factors ([Carroll, Slacalek, Tokunaka and White, 2017](#); [Aguiar, Bils and Boar, 2024](#)), elasticities of intertemporal substitution ([Aguiar et al., 2024](#)), risk aversion ([Kaplan and Violante, 2022](#)), and investment opportunities ([Kaplan and Violante, 2022](#)). While heterogeneity along these margins could partially explain the *aggregate* bowtie pattern,

they cannot explain the shape of the *within-household* policy function for consumption fixers, savings fixers, consumption prioritizers, and savings prioritizers.

Multiple-Asset Models. To account for the presence of wealthy agents with high MPCs, the wealthy hand-to-mouth, [Kaplan and Violante \(2014\)](#) introduce multiple accounts of varying liquidity to the standard incomplete markets model. A variant of this model forms the basis for the highly influential heterogeneous-agent New Keynesian (HANK) model of [Kaplan, Moll and Violante \(2018\)](#). In this model, households are hand-to-mouth if they have low liquid wealth, even if they have high illiquid wealth. Nonetheless, households with high liquid wealth should be unconstrained and therefore not have extreme MPCs. In [Figure 3](#), we document a pronounced “bowtie” pattern of adjustment for all levels of liquid wealth. Even households that have more than \$100,000 of liquid wealth or liquid wealth exceeding ten times their monthly income display the “bowtie” pattern. In [Figure B.2](#), we corroborate this finding for other measures of wealth and debt.

Models with Durables or Consumption Commitments. To account for differences in household consumption-savings decisions between durables and non-durables, many models in the literature explicitly study the role of durable consumption (see, *e.g.*, [Barsky, House and Kimball, 2007](#)). In a similar vein, [Chetty and Szeidl \(2016\)](#) study consumption-savings decisions when households may be pre-committed to certain spending patterns, *e.g.*, because of contracts that they have previously entered into to rent or lease a good. Such models predict that any given household will be increasingly likely to undergo a large increase in consumption and have high MPCs (as households purchase a lumpy durable good) as they experience progressively larger positive shocks. This is at odds with our finding that very few households transition from an interior MPC to an extreme MPC of 1 (or even higher) as positive shocks get larger. These models are also inconsistent with the behavior of consumption prioritizers and savings fixers, who have an MPC of unity for small positive shocks, and account for 40% of respondents.

Models with Mistakes in Consumption. To account for the high MPCs that we see in the data, many papers have introduced behavioral elements to consumption-savings problems. Some prominent such models are those with present bias (see, *e.g.*, [Maxted et al., 2024](#)), temptation (see, *e.g.*, [Krusell, Kuruşçu and Smith, 2002](#)), rational inattention ([Sims, 2003](#)), finite planning horizons ([Boutros, 2025](#)), sparsity ([Gabaix, 2014, 2023](#)), or misperceptions of wealth ([Lian, 2023](#)). While these models generate higher (or lower) MPCs than the incomplete markets model, they once again do not generate the “bowtie” pattern of responses as a function of shock size ([Figure 1](#)) or the stark and discrete heterogeneity in policy functions that we uncover.

Models with Infrequent Optimization. Fuster et al. (2021) generate infrequent optimization through “menu costs” of adjusting consumption. In this model, all households are consumption fixers who do not adjust their consumption for small shocks. However, in our data, 80% of quick-fixing households (savings fixers, consumption prioritizers, and savings prioritizers) adjust their consumption even in response to small shocks. In our model, households differ in whether they adjust along the margin of consumption or saving. Moreover, in the case of consumption and savings prioritizers, this margin also depends on the sign of the shock. In the related analysis of Gabaix and Laibson (2001), households similarly adjust consumption only every so many periods. This model is therefore similarly unable to account for the multiplicity of quick-fixing behavior.

Our Model of Quick-Fixing. We have argued that no existing model can capture the new empirical facts that we have documented. Instead, our empirical evidence on households’ policy functions is consistent with a model of quick-fixing. It suggests that extreme MPCs often constitute simple quick-fixes that reduce physical or mental transaction costs. Our model captures such a decision-maker who prefers quick-fixes as long as the opportunity cost of not implementing the optimal response is sufficiently small. In the remainder of the analysis, we will construct a quantitative model of quick-fixing that is consistent with our new findings as well as more familiar facts about MPCs, which have motivated many of the approaches described above.

4 A Quantitative Model with Quick-Fixing

Our simple framework of Section 2 lacked many empirically relevant features that influence household consumption dynamics. In this section, we enrich our framework with three features present in benchmark quantitative macroeconomic models: more than two periods, uninsurable idiosyncratic income risk, and borrowing constraints. We then demonstrate that a parsimoniously calibrated, quantitative model of quick-fixing can account remarkably well for both the patterns of adjustments and marginal propensities to consume in the data.

4.1 Setup

Time is discrete and indexed by $t \in \mathbb{N}$. There is a unit measure of households indexed by $i \in [0, 1]$. All households have expected discounted utility preferences with discount factor $\beta \in [0, 1]$. Their flow payoff is $u(c) = \frac{c^{1-\gamma}-1}{1-\gamma}$, where γ is the inverse elasticity of intertemporal substitution. In each period t , each household earns a stochastic income y_t which lies in a discrete set $\mathcal{Y} \subset \mathbb{R}_{++}$. Income follows a first-order Markov process with transition matrix

P . Households can save in a risk-free bank account with gross interest rate $R \in \mathbb{R}_+$. We denote their savings as a_t . Due to incomplete markets, households cannot borrow: $a_t \geq 0$. Households differ in their quick-fix types, as we describe below.

Frictionless Optimizers. We first introduce dynamic optimization of households that frictionlessly optimize, type “R” for rational. We let V^R denote their value function:

$$\begin{aligned} V^R(a, y) &= \max_{a'} \{u(c) + \beta \mathbb{E} [V^R(a', y') \mid y]\} \\ \text{s.t. } a' &= Ra + y - c \\ a' &\geq 0 \end{aligned} \tag{2}$$

where the expectation is taken over unknown income states y' . We let c^* denote the optimal consumption policy function for these households.

Quick-Fixers. We now introduce the problem of quick-fixers. As in Section 2, we associate each household i with a *quick-fix consumption function* c^{q_i} and with a quick-fix-specific *utility cost* $\kappa_{q_i} \in \mathbb{R}_+$.

We use the survey to discipline the quick-fixes that households use. To describe these behaviors dynamically, it is necessary to keep track of two additional household state variables: a reference consumption level $\bar{c} \in \mathbb{R}_+$ and a reference income state $\bar{y} \in \mathcal{Y}$. The four quick-fixes are described by four functions, indexed by $q \in \{\text{CF}, \text{SF}, \text{CP}, \text{SP}\}$, that depend on reference consumption $\bar{c}$ and an income deviation $y - \bar{y}$:

$$\begin{aligned} c^{\text{CF}}(\bar{c}, y - \bar{y}) &= \bar{c} & c^{\text{CP}}(\bar{c}, y - \bar{y}) &= \bar{c} + \max\{y - \bar{y}, 0\} \\ c^{\text{SF}}(\bar{c}, y - \bar{y}) &= \bar{c} + (y - \bar{y}) & c^{\text{SP}}(\bar{c}, y - \bar{y}) &= \bar{c} + \min\{y - \bar{y}, 0\} \end{aligned} \tag{3}$$

Consumption fixers (CF) consume a fixed amount. *Consumption prioritizers* (CP) consume a fixed amount plus the positive component of income shocks, while absorbing negative shocks with savings. *Savings fixers* (SF) save a fixed amount, $s^{\text{SF}} = y - c^{\text{SF}}(\bar{c}, y - \bar{y}) = \bar{y} - \bar{c}$. *Savings prioritizers* (SP) save a fixed amount plus the positive component of income shocks, absorbing negative shocks with consumption.¹⁰

In each period, quick-fixers decide whether to adopt their quick-fix, which we denote by $D = 0$, or to pay utility cost κ_{q_i} to adopt the unconstrained, optimal choice, in which case $D = 1$. If the household pays the utility cost, they also reset their reference consumption and income states to $c^*(a, y)$ and y , respectively. For each type q , this behavior is described

¹⁰If c^q is negative, we adopt the convention that the household automatically abandons their quick-fix.

by the dynamic program

$$\begin{aligned}
V^q(a, y, \bar{c}, \bar{y}) = \max_{D \in \{0,1\}} \{ & D(u(c^*(a, y)) + \beta \mathbb{E}[V^q(a', y', c^*(a, y), y) \mid y] - \kappa_q) \\
& + (1 - D)(u(c^q(a, y, \bar{c}, \bar{y})) + \beta \mathbb{E}[V^q(a', y', \bar{c}, \bar{y}) \mid y])\} \\
\text{s.t. } & a' = Ra + y - (D(c^*(a, y)) + (1 - D)(c^q(a, y, \bar{c}, \bar{y}))) \\
& a' \geq 0
\end{aligned} \tag{4}$$

Two modeling choices that are necessary in the dynamic model are the treatment of (i) the persistence of types and (ii) households' sophistication in understanding the future consequences of near-rationality. First, we treat the identity of an individual's quick-fixing function as a permanent characteristic. This is a conservative approach that makes it as hard as possible for us to match our empirical findings that household wealth and financial status are poor predictors of quick-fixing types (Fact 3). Moreover, because we will find that households' quick-fixing type and their financial status are essentially unrelated in the model as well (Section 5.2), type instability would not change our main conclusions as long as the aggregate distribution of types is stable. Second, when households abandon their quick-fix, they adopt the choice of frictionlessly optimizing households. We follow this approach because it is conservative in the sense that it will yield opportunity costs of near-rationality that exceed those that sophisticated households would suffer.

We also emphasize that this combination of assumptions makes the quick-fixing model highly numerically tractable. This arises because each step in the recursive algorithm requires solving only a binary optimization problem (to quick-fix or not) by comparing two numbers in each state. These points notwithstanding, it is of course straightforward to adapt the framework to accommodate transient type membership and sophistication.

Therefore, it would be simple to integrate our framework in larger, general equilibrium models by combining the quick-fixing “demand block” with any “supply block” and “policy block.” As one example, if real interest rates are fixed, then—for any shock of interest—the consumption responses of households in this model can be used to compute general equilibrium responses via the intertemporal Keynesian cross (Auclert, Rognlie and Straub, 2024). However, as will soon see, the appropriate MPCs depend intimately on the size, sign, and incidence of shocks in the cross-section and cannot be taken as shock-invariant sufficient statistics. We now proceed to focus on the novel parts of our analysis, which are the partial equilibrium predictions directly disciplined by our data.

4.2 Calibration

We calibrate the model to match standard facts on US households’ behavior as well as our survey findings. We proceed in four steps.

First, we calibrate the flow utility and income process to match external estimates. We set $\gamma = 1$ (logarithmic preferences) to match standard estimates of the EIS. We calibrate the earnings process to match the frequency and size of quarterly-frequency earnings shocks in US micro data. The process is a 5-state discretization of a Gaussian AR(1) process that targets a variance in log annual earnings of 0.70 and an expected state switch of once every five quarters.¹¹ We scale income such that one unit coincides with the median quarterly income reported in our survey, \$15,625. We set the quarterly interest rate to $R = 1.01$.

Second, we calibrate the discount factor to match the spending behavior of households that do *not* quick-fix in the survey (*i.e.*, are “unclassified”). For the frictionlessly optimizing types in the model, we calculate the average MPC out of transfer shocks of size x as $\text{MPC}_x^R = \frac{1}{x} \int (c^*(a, y+x) - c^*(a, y)) d\Phi^R(a, y)$, where $\Phi^R(a, y)$ is the stationary distribution over assets and income for frictionless optimizers.¹² We choose the discount factor to minimize the sum of squared residuals between these predictions and the measured MPCs of “unclassified” survey respondents. This results in a calibrated value of $\beta = 0.92$. Insofar as this is a low value, we emphasize that this is a conservative calibration so that the standard incomplete markets model can match observed MPCs in the data (see *e.g.*, [Kaplan and Violante, 2022](#)).

Third, we calibrate the fraction of agents of each quick-fixing type to match the categorization in [Figure 4](#). The fraction of frictionless optimizers is matched to the share of unclassified households in the data. As any quick-fixers outside our four types are coded as unclassified in the data, this represents an upper bound on the fraction of frictionless optimizers and is therefore conservative for the near-rational theory.¹³

Fourth, we calibrate the four type-specific utility costs, $(\kappa_{CF}, \kappa_{SF}, \kappa_{CP}, \kappa_{SP})$, to match our main findings about quick-fixing behavior in the survey ([Figure 4](#)). Specifically, for each type and each shock size except the \$50 gain and loss (as these are used to empirically define the types), we calculate in the survey the fraction of respondents who abandon their quick-fix by reporting a propensity to consume that does not coincide with the quick-fix. In the model,

¹¹This calibration matches the variance in log annual earnings estimated by [Kaplan et al. \(2018\)](#) using Social Security Administration data, as well as the variance in the 1-year change in log annual earnings (0.23). See [Kaplan et al. \(2018\)](#) (Table III) for further details.

¹²When computing MPCs, we treat the transfer shock as transitory (*i.e.*, as a shift in households’ liquid asset holdings), but write it as a shift in income to keep the notation consistent with our quick-fixes above.

¹³However, the behavior of frictionlessly optimizing households in the model does a good job of accounting for the average MPCs of unclassified households for *all* income shocks (see [Figure B.5](#)).

we calculate, for each shock x ,

$$\text{ReoptFraction}_x^q = \int D_x^{q*}(a, y, \bar{c}, \bar{y}) d\Phi^q(a, y, \bar{c}, \bar{y}) \quad (5)$$

where $D_x^{q*} \in \{0, 1\}$ denotes the optimal reoptimization policy for type q in response to shock x and Φ^q is the model-implied stationary distribution for those types. To make this experiment most consistent with the survey, we assume that the household contemplates a shock $x \in \mathbb{R}$ in an interim period after initially choosing whether or not to reoptimize in a given period, but before observing income or making decisions of the next period (see Appendix A.3 for the formal details). The utility costs affect reoptimization behavior directly via the optimal policy and indirectly via the stationary distribution of observed and latent states. For each type, we choose the parameter κ_q to minimize the sum of squared residuals of the model versus the data:

$$\kappa_q^* = \arg \min_{\kappa_q > 0} \left\{ \sum_{i=1}^{12} \left(\text{ReoptFraction}_{x_i}^q - \widehat{\text{ReoptFraction}}_{x_i}^q \right)^2 \right\} \quad (6)$$

where the 12 shocks are those asked in the survey, excluding the \$50 gain and loss. We report the calibrated values of κ_q and provide an economic interpretation of their magnitude in Table 1 in Section 5.1.


4.3 Model Fit

Reoptimization. Figure 9(i) compares the model prediction and data for the key moments that discipline the utility costs of not quick-fixing, namely the reoptimization thresholds for different shock sizes. The model fits the overall adjustment pattern quite well, in spite of using only four parameters (the utility costs) to match 48 moments (the adjustment thresholds). Naturally, because the model is heavily overidentified, we do not exactly match all of the measurements. The largest gap between the model's fit with the data occurs in cases in which consumption adjusts and savings are fixed (all of Panel B, the positive shocks in Panel C, and the negative shocks in Panel D). In these cases, the model underestimates adjustment for low shock sizes and overestimates adjustments for large shock sizes.

Average MPCs. Figure 9(ii) shows how the model fits average MPCs and the decomposition of MPCs between the extensive and intensive margin. While the calibration directly targets reoptimization behavior (Figure 9(i)) and the average MPC of unclassified agents (see Figure B.5), it does not target the MPC profiles of quick-fixing agents. The model deviation in average MPCs is small (Panel A). The maximum difference is attained for small


Figure 9: Quantitative model fit

(i) Reoptimization in response to shocks



Notes: The bar graphs compare empirical measurements and model predictions for the propensity of agents to abandon their quick-fix (reoptimize) following unanticipated shocks. Each panel corresponds to one of the four types and therefore to the calibration of the corresponding utility cost parameter. The blue bars denote model predictions, described in Equation 5, and the gray bars are empirical measurements, as reported in Figure 4.

(ii) Average marginal propensities to consume (MPCs)



Notes: The line graphs compare empirical measurements and model predictions for the average marginal propensity to consume out of shocks of different amounts. The decomposition corresponds exactly to that of Figure 7. Panel A shows average MPCs. Panel B shows variation due to the extensive margin, calculated by assuming that quick-fixers follow their associated quick-fix and non-quick-fixers use a reference interior MPC of 0.45. Panel C shows the residual between Panel A and Panel B. The blue dashed line and dots correspond to the model prediction, and the black solid line and dots correspond to the data (and exactly to the “All types pooled” lines of Figure 7).

negative shocks (less than or equal to \$250) and can be largely accounted for by a larger residual component (Panel C). This is driven by the high estimated fraction of agents at the binding borrowing constraint (see also Figure 10). This generates the correct qualitative pattern that residual variation is largest for small negative shocks, but slightly overstates the quantitative magnitude. Finally, as in the data, we note that a large part of the size-dependence in aggregate MPCs arises due to variation in the extensive margin (which is attributed to quick-fixes), as opposed to residual variation in the intensive margin (which is attributed to the concavity of the frictionlessly optimizing consumption function).

The MPC Distribution. We finally re-create our key empirical finding about the shape of the MPC distribution (Figure B.6). The model replicates the “bowtie” pattern, whereby extreme MPCs (0 or 1) are less likely in response to larger shocks.

Summary. Our quantitative model matches our novel empirical findings—that households adopt heterogeneous quick-fixes for small shocks before abruptly switching for large shocks—as well as established facts about the MPC distribution that are familiar from previous observational and survey studies (see Appendix C.4 for a full review of these findings). That is: (i) MPCs are high on average; (ii) MPCs decline in shock sizes; (iii) MPCs are higher for losses than for gains; (iv) MPCs vary widely across households; and (v) many households have an MPC of 0 or 1. Moreover, the model can account well for the observed frequency of reoptimizations and average MPCs observed in the data.

5 The Economic Consequences of Quick-Fixing

Having estimated the model, we next explore its properties. We have six main findings. First, households experience very small losses from quick-fixing, consistent with the hypothesis of near-rationality. Second, there is essentially no relationship between quick-fixing types and wealth accumulation (a non-targeted moment), consistent with our empirical results. Third, quick-fixing generates considerable heterogeneity in MPCs. Fourth, targeting transfers by wealth is much less effective than standard models imply, and it can even backfire. Fifth, quick-fixing rationalizes significant sign- and size-dependence in MPCs, above and beyond the frictionless model. Finally, the aggregate response to income shocks depends markedly on their *incidence*. That is, it matters whether aggregate shocks have a skewed incidence on a few people, which triggers reoptimizations and smaller spending responses per dollar, or a small incidence on many people, which leads to quick-fixing and larger spending responses per dollar.

Table 1: The small costs of near-rationality

Panel A: Utility costs κ_q		
Household type	% reduction in consumption	Average dollar cost
Consumption fixer	1.10	\$174.74
Savings fixer	0.007	\$1.47
Consumption prioritizer	0.006	\$1.44
Savings prioritizer	0.11	\$18.54

Panel B: Value loss due to near-rationality $V^R - V^q$ (per quarter)		
Household type	% reduction in consumption	Average dollar loss
Consumption fixer	0.45	\$71.34
Savings fixer	0.004	\$0.58
Consumption prioritizer	0.003	\$0.54
Savings prioritizer	0.06	\$8.68

Notes: Panel A reports the calibrated utility costs κ_q , in economically interpretable units. Panel B reports “costs of near-rationality” defined as the change in value for frictionless optimizers were they to adopt quick-fixing, on average. See Section 5.1 for details.

5.1 Quick-Fixing Behavior is Near-Rational

Our hypothesis of near-rationality relies on the assumption that the opportunity costs of quick-fixing are small. But how small is small? In Panel A of Table 1, we show that the calibrated utility costs that rationalize the quick-fixing uncovered by our survey are payoff-equivalent to at most a 1% reduction in consumption or \$175 one-time loss. The first column reports these losses for all types in payoff units ($100 \times \kappa_q$, which can be interpreted as percent consumption reduction due to logarithmic preferences), and the second column reports these as dollar equivalents.¹⁴ The costs are, by some margin, highest for consumption fixers. The losses are on the order of *one hundredth of a percent* for savings fixers and consumption prioritizers, or about \$1.50.

We next compute the lifetime losses from near-rationality. Concretely, we compute the average lifetime loss in value for frictionless optimizers if they were to adopt the behaviors and bear the decision costs of quick-fixers:

$$\Delta V^q = \int (V^R(a, y) - V^q(a, y, c^*(a, y), y)) d\Phi^R(a, y) \quad (7)$$

¹⁴For an agent in state z , the dollar-equivalent cost of abandoning their quick-fix when consuming $c^q(z)$ solves $\log(c^q(z) - \Delta^q(z)) - \log(c^q(z)) = -\kappa$, and is therefore $\Delta^q(z) = c^q(z)(1 - e^{-\kappa})$. We compute $\mathbb{E}[\Delta^q(z)|D^*(z) = 1]$, or the average dollar-equivalent cost conditional on reoptimizing.

We express this in units of an equivalent per-period reduction in consumption.¹⁵

The costs of near-rationality in our calibrated model are all less than 0.5% of per-period consumption or \$75 per quarter (Panel B of Table 1). On average, among quick-fixers, the loss is \$16 per quarter. Notably, the losses from quick-fixing for savings fixers are smaller than those of consumption fixers. This is because quick-fixes that fix savings allow consumption to respond (one-to-one) to shocks. For this reason, their response coincides with the optimal one when the borrowing constraint binds strictly. Moreover, away from the borrowing constraint, they can achieve close to the optimal response if the income shock is sufficiently persistent. A similar intuition underscores why consumption prioritizers face smaller losses than savings prioritizers.

Although we have shown that quick-fixes are unpredictable from standard economic and demographic characteristics in the data (Fact 3), this lower latent utility cost could explain why a much larger proportion of households are classified as savings-fixers (29%) relative to consumption fixers (14%). In all cases, however, the small loss from near-rationality helps explain why quick-fixing might persist even in a world of “selection pressure” against suboptimal strategies: the payoff cost of being a quick-fixer is extremely small.


5.2 Quick-Fixing Types Are Unrelated to Wealth

One moment that we did not use for calibration is the limited relationship between quick-fixing types and wealth (Fact 3). This is not a foregone conclusion in the quantitative model. Assets are an endogenous state variable and savings responses to shocks differ across types. It is a reasonable conjecture, for example, that households that default to increasing savings in response to shocks (savings prioritizers) accumulate more wealth than those that default to increasing consumption (consumption prioritizers).

We find instead that the wealth distributions of each of the four quick-fixing types and frictionless optimizers are essentially identical (Figure 10). Most cross-sectional variation in wealth is driven by large and persistent shocks to income. Any short-run differences in savings responses to these shocks wash out in the long run, when households eventually reoptimize. As one concrete example: while savings fixers are as-if “hand-to-mouth” in response to small shocks, they are not actually “hand-to-mouth” in the long run. More succinctly, savers are not rich and spenders are not poor. This represents an important difference from two-agent New Keynesian (TANK) models, in which a fixed fraction of agents consumes their entire income in every period (*e.g.*, Campbell and Mankiw, 1989; Debortoli and Galí, 2024). These

¹⁵The equivalent percentage reduction in a consumption stream $\{c_t\}_{t=0}^{\infty}$ solves $\Delta V^q = \sum_{t=0}^{\infty} \beta^t \log(c_t) - \sum_{t=0}^{\infty} \beta^t \log((1 - \delta^d)c_t)$ and is therefore $\delta^d := 1 - e^{-(1-\beta)\Delta V^q}$. In Panel B of Table 1, we report $100 \times (1 - \beta)\Delta V^q$ in column 1 and $\delta^d \times \mathbb{E}[c^*(z)]$, where $\mathbb{E}[c^*(z)]$ is the average frictionless consumption, in column 2.

Figure 10: Wealth distributions by type



Notes: This figure shows the wealth distribution for each type. Panel (a) shows the frequency of households at $a = 0$, the borrowing constraint. Panel (b) shows the distribution of wealth conditional on $a > 0$.

findings also contrast with those obtained in the class of models of agents with heterogeneous discount factors or present bias (Aguiar et al., 2024; Maxted et al., 2024), which naturally lead to very different long-run rates of saving.

5.3 Quick-Fixing Generates MPC Heterogeneity

We next explore how much quick-fixing contributes to heterogeneity in the marginal propensity to consume. Table 2 reports the percent of variance in this object that can be explained by assets and income in the quick-fixing model and the nested model in which all agents are frictionless optimizers. This is a stronger notion of “predictability” than what we can estimate via regression in the survey data because, in the model, we can calculate exact conditional expectation functions and have no measurement error; thus, this calculation gives a lower bound for what remains to be explained by other sources of heterogeneity. In the model with frictionless optimization, assets and income explain all variation by construction. This is strongly at odds with our findings as well as those from other studies in the literature (*e.g.*, Lewis et al., 2024). In our model, 28% of total MPC variance and 70% of variance conditional on $a > 0$ is unexplained by assets and income, and therefore introduced by quick-fixing behavior.¹⁶ Thus, quick-fixing helps break the tight connection in incomplete markets models between financial observables and MPCs.

¹⁶Note that the overall unexplained variance is lower than the unexplained variance conditional on either $a = 0$ or $a > 0$. Intuitively, this is because the variance of MPCs conditional on assets is sizable, which reduces the overall unexplained variance once one conditions on assets (by the law of total variance).

Table 2: Variance in MPCs unexplained by assets and income

Model	Overall	Conditional on $a = 0$	Conditional on $a > 0$
Quick-fixing	28%	43%	70%
Frictionless optimization	0%	0%	0%

Notes: In each model, we calculate $\text{Var}[\text{MPC}_i|a_i, y_i]$, where i indexes households, MPC_i is the average MPC across the 14 scenarios considered in the survey, a_i is the household’s wealth, and y_i is the household’s income. In the first column, we report $100 \cdot \mathbb{E}[\text{Var}[\text{MPC}_i|a_i, y_i]]/\text{Var}[\text{MPC}_i]$, or the fraction of variance unexplained by wealth and income. In the second column, we report the same conditioning on $a = 0$: $100 \cdot \mathbb{E}[\text{Var}[\text{MPC}_i|a_i, y_i] | a_i = 0]/\text{Var}[\text{MPC}_i | a_i = 0]$. In the third, we report the same conditioning on $a > 0$: $100 \cdot \mathbb{E}[\text{Var}[\text{MPC}_i|a_i, y_i] | a_i > 0]/\text{Var}[\text{MPC}_i | a_i > 0]$. Since (a, y) is the state variable in the problem with frictionless optimization, the fraction unexplained is always 0 in that model.

5.4 Quick-Fixing Makes Targeting Transfers Less Effective


So far, we have shown that quick-fixing dulls the connection between financial circumstances and MPCs. A related question, which is directly policy relevant, is whether targeting transfers based on observable financial characteristics can improve the efficacy of fiscal stimulus. To study this question, Figure 11 shows how the aggregate MPC out of gains depends on wealth in the frictionlessly optimizing benchmark and in our quick-fixing model.

In the benchmark, high MPCs are exclusively concentrated among households with low levels of liquid wealth. This feature is driven by liquidity constraints and uninsurable income risk, which induce a consumption function that is concave in wealth. Nonetheless, this prediction of the frictionlessly optimizing benchmark is inconsistent with observational evidence showing that even households with relatively high liquid wealth have large MPCs (Kueng, 2018; Boehm, Fize and Jaravel, 2025).

By contrast, in the quick-fixing model, MPCs are relatively high even for wealthy households. Strikingly, MPCs can even be *non-monotone* in wealth under our calibration. Intuitively, rich households can have high MPCs in our model because of quick-fixing. Richer households experience a lower utility loss from failing to perfectly smooth consumption, so many of them keep quick-fixing and maintain a high MPC in response to small income shocks. This mechanism is consistent with our empirical finding that richer households are indeed more likely to maintain quick-fixes for large shocks (Figure B.4). And, in our calibration, it implies that a household that holds two years of income ($a \approx 8$) has a higher MPC out of a \$500 transfer than one that holds one quarter of income ($a \approx 1$).

A policy implication is that targeting fiscal stimulus by (liquid) wealth may be only marginally more effective (or even less effective) than broad-based transfers. The appeal of targeting is even further diminished by the next phenomena that we study: as we will soon see, smaller transfers spread across many households elicit a greater spending response

Figure 11: Average MPCs out of gains as a function of assets



Notes: This figure plots aggregate MPCs out of a 250 dollar gain (Panel A) and a 500 dollar gain (Panel B) as a function of assets in the full model (blue line) and the frictionless model (gray line). In each case, we report MPCs conditional on wealth averaged over all other state variables, using the model's stationary distribution.

per dollar than larger transfers concentrated on a small subset of the population, precisely because the former lead to more quick-fixing.


5.5 Quick-Fixing Rationalizes Sign- and Size-Dependent MPCs

To understand how quick-fixing affects MPCs, we now compare the MPCs in the quick-fixing model to those from the nested frictionless model with perfectly optimizing households. Figure 12 plots the difference in MPCs between the two models across shock sizes. Two findings emerge, both of which were foreshadowed by Fact 4 from our empirical analysis.

Size-dependence. First, MPCs are substantially more *size-dependent* with quick-fixing. For example, for \$100 and \$250 shocks, the aggregate MPC with quick-fixing exceeds the frictionless MPC by about 0.1 to 0.2. While these shock sizes may not sound particularly large within the recent historical context of US stimulus checks, they are 0.6% and 1.6% of quarterly GDP, respectively. For reference, Tempalski (2006), in his survey of major US tax bills between 1940 and 2006, finds that the revenue effects of the average major policy change amounts to 0.4% of US GDP within the first year, and only three policies are larger than 1% of GDP. Thus, typical (fiscal and non-fiscal) shocks studied at business-cycle frequencies can have quantitatively very different effects under quick-fixing relative to the frictionless benchmark.

Differences in MPCs across the two models persist also for much larger shocks: for example, quick-fixing increases the aggregate MPC out of a \$1,000 loss by about 0.06. In a

Figure 12: Decomposing how quick-fixing shapes sign- and size-dependent MPCs



Notes: This figure plots the difference in the model-implied marginal propensity to consume between the entire population, which includes quick-fixers, and a counterfactual population of frictionless optimizers. The black line displays the total difference. The colored bars display the contribution of each group of quick-fixers: formally, $w^j \cdot (\text{MPC}^j - \text{MPC}^R)$, where w^j is the fraction of the population in quick-fixing group j , MPC^j is the average MPC of group j , and MPC^R is the average MPC of the frictionless optimizers. The sum of the colored bars equals the black line.

simple, textbook calculation of shock multipliers, such a difference in MPCs would amount to a difference of 0.3, or 30 cents of additional amplification per 1 dollar of shock on impact.¹⁷

The decomposition in Figure 12 shows that savings fixers drive most of these differences. Their MPCs are very high for small shocks but fall sharply for large shocks once they reoptimize. The quick-fixing model therefore generates stronger size dependence for two reasons: (i) larger shocks trigger adjustments that shift MPCs from the extremes toward the interior, and (ii) given our estimated distribution of quick-fixing types, more reoptimization lowers average MPCs. In the frictionless model, by contrast, size-dependence comes entirely from the strict concavity of the consumption function in liquid assets. Larger gains in this model move households to flatter regions of their consumption functions, reducing average MPCs, though by less than in the quick-fixing model; while larger losses in this model move households to steeper regions, raising average MPCs—contrary to the data.

For larger shocks, the quick-fixing and the frictionless model predict aggregate MPCs that are closer to one another. Even then, however, the underlying household behavior differs sharply. For example, 33% of all households (and 49% of all quick-fixing households) still quick-fix after a \$500 gain. In the quick-fixing model, the average response thus reflects

¹⁷In the often studied case of rigid prices and fixed real interest rates (see *e.g.*, Flynn, Patterson and Sturm, 2024), a simple version of the Keynesian cross holds such that the effect of a \$1 income shock on impact is $\$1/(1-\text{MPC})$, where MPC is the aggregate average MPC.

a mix of households that continue to rely on simple rules and households that reoptimize. In the frictionless model, by contrast, high MPCs are driven almost entirely by constrained households. The quick-fixing model can thus generate high MPCs while also matching the micro evidence across shocks and households, including—as we have seen before—the weak relationship between MPCs and wealth.


Sign-dependence. The second pattern we observe is that MPCs are substantially more *sign-dependent* with quick-fixing. For example, we observe that the difference in the aggregate response to a \$500 loss versus a \$500 gain is approximately 0.10 larger with quick-fixing than without. As the decomposition makes clear, most of this difference is driven by the fact that savings prioritizers amplify the aggregate MPC for small and medium-sized losses but shrink it for gains. Sign-dependent MPCs imply asymmetrically large propagation of negative income shocks, which would be consistent with the classic observation that the business cycle features faster crashes and slower recoveries ([Hamilton, 1989](#); [Acemoglu and Scott, 1997](#)).

Intertemporal MPCs. We finally note that the same properties also extend to intertemporal profiles of spending. Concretely, using the dynamic structure of the model, we calculate *intertemporal marginal propensities to consume*, defined as excess spending at date $t + h$ as a fraction of an income shock at date t . As we show in [Figure B.7](#), iMPC profiles vary considerably across types. For example, savings fixers and consumption prioritizers consume almost all of a \$100 shock in one quarter, consumption fixers and savings prioritizers keep about 40% saved after one year, and frictionless optimizers behave somewhere in between (top panel). These differences arise because, for sufficiently small shocks, quick-fixers do not reoptimize even over several quarters (middle panel). The end result is that iMPC patterns are not only highly heterogeneous across households, but also more *front-loaded* compared to the benchmark model (bottom panel). This is economically important as front-loading of iMPCs increases the effects of fiscal shocks in canonical macroeconomic models (see *e.g.*, [Angeletos, Lian and Wolf, 2024](#)). This calculation demonstrates the applicability of the “quick-fixing consumption block” for informing estimates of shock propagation in a large class of Keynesian models for which iMPCs are sufficient statistics ([Auclert et al., 2024](#)).

5.6 Quick-Fixing Generates Incidence-Dependence

Different macroeconomic shocks have different patterns of incidence on the population. For example, income losses during recessions are not experienced uniformly: the vast majority of people experience either no change in income or a small change in income, while some people lose their jobs and experience enormous income losses ([Guvenen, Schulhofer-Wohl,](#)

Figure 13: Aggregate MPCs depend significantly on shock incidence



Notes: This figure plots aggregate MPCs out of transitory income shocks with different incidence distributions. Panel A is a \$250 gain per person and Panel B is a \$250 loss per person. These shocks are 1.6% of the mean income state in the model. In each figure, the horizontal axis varies the incidence scenario (*e.g.*, $x\%$ of the population receives a shock of y , and $(100 - x)\%$ of the population receives a shock of 0). The blue bars show average MPCs under the quick-fixing model and the gray bars show average MPCs under the nested model in which all agents are frictionless optimizers.

Song and Yogo, 2017; Schmidt, 2025). Conversely, untargeted fiscal transfers often affect incomes in ways that are closer to uniform. An important question is whether these different patterns of incidence have first-order effects on macroeconomic propagation.

Existing work argues that, in representative agent theories and in standard calibrations of the incomplete markets model, average MPCs do *not* depend significantly on the exact size or incidence of shocks (Brinca, Faria-e Castro, Ferreira and Holter, 2019; Auclert, Rognlie and Straub, 2025). Therefore, shock propagation depends primarily on a (nearly) context-independent average MPC.

By contrast, the size-dependence of MPCs in our near-rational model implies markedly different aggregate consumption responses to a shock depending on its distribution of incidence. Figure 13 illustrates this point, showing the aggregate MPC for varying levels of incidence under quick-fixing (blue bars) and the frictionless benchmark (gray bars). We fix the aggregate size of the shock at \$250 which, as mentioned above, amounts to a 1.6% change in quarterly income for the average household. Panel (A) shows that the aggregate MPC from a uniform \$250 shock is 0.55, while the aggregate MPC from a \$10,000 shock incident on 2.5% of the population is 0.24, more than half as small. This effect is smaller for the benchmark with frictionless optimization. Panel (B) shows that, for a loss, the near-rational model again features a declining MPC as incidence becomes more concentrated (0.66 vs. 0.52). By contrast, under the rational model, the MPC is very stable and even slightly increases as the incidence becomes more concentrated.

A Lucas Critique for MPCs. The size- and incidence-dependence of the macroeconomic response to income shocks generates a new form of the [Lucas \(1976\)](#) critique, emphasizing the pitfalls of assuming a stable and (nearly) linear aggregate consumption function. This approach has reemerged as a dominant paradigm in the HANK literature (see the review by [Auclert et al., 2025](#)). However, our analysis emphasizes that equally-sized aggregate shocks that close to uniformly affect the whole population (like a fiscal stimulus) may affect aggregate consumption very differently than shocks that are tightly concentrated on a subset of the population (like unemployment during a recession). Thus, the “right” aggregate MPC depends intimately on the counterfactual or shock scenario under consideration. A corollary of this point is that care should be taken in calibrating to MPCs estimated out of shocks of certain sizes when researchers are using such calibrated models to evaluate counterfactuals that feature differently-sized shocks.

6 Conclusion

This paper starts from the idea that, when changing consumption-savings behavior is costly, people may instead rely on simple near-rational policy functions: *quick-fixes* that avoid these costs. We develop a model of quick-fixing to capture this idea. By developing and fielding a novel survey, we can recover the consumption policy functions required to test the theory. We find that many households quick-fix by following simple rules of fully spending or fully saving in response to small income shocks, while pursuing a more moderate strategy for large shocks. This behavior is consistent with near-rationality but inconsistent with alternative models. Our identified quick-fixes discipline the “wilderness” of irrationality and allow us to assess its macroeconomic implications.

Quick-fixing has implications for consumption and savings behavior at both microeconomic and macroeconomic levels. First, quick-fixing is near-rational. Even very small utility costs can account for the economically large deviations in behavior from the benchmark with frictionless optimization. Second, quick-fixing opens the “black box” of latent heterogeneity in the marginal propensity to consume. Quick-fixing types are essentially unpredictable by demographic and financial variables while accounting for a significant fraction of MPC variation. Third, by generating significant size-, sign-, and incidence-dependence together with non-monotonicity in MPCs by wealth, quick-fixing has potentially important implications for business cycle dynamics and policy design. Indeed, given empirical quick-fixing behavior, care should be taken in extrapolating estimates of MPCs into counterfactuals: if the distribution of income shocks under counterfactuals differs from that in the estimation sample, then these MPCs may be incorrect.

Finally, our findings convey a broader lesson about how to “put the near-rationality hypothesis to work.” The premise that changing behavioral rules is costly relative to the utility loss of a simple quick-fix plausibly holds for many other important economic decisions, like portfolio choice or labor supply. But this observation, by itself, does not suffice for making specific predictions for how people actually behave. We still need to understand *which* quick-fixes economic agents use in practice. In these cases, our approach of combining theory, a tailored empirical design to uncover which quick-fixes agents use, and quantitative modeling to study aggregate consequences could prove valuable.

References

- Abbring, Jaap H and James J Heckman**, “Econometric evaluation of social programs, part III: Distributional treatment effects, dynamic treatment effects, dynamic discrete choice, and general equilibrium policy evaluation,” *Handbook of econometrics*, 2007, 6, 5145–5303.
- Acemoglu, Daron and Andrew Scott**, “Asymmetric business cycles: Theory and time-series evidence,” *Journal of Monetary Economics*, 1997, 40 (3), 501–533.
- Aguiar, Mark, Mark Bils, and Corina Boar**, “Who Are the Hand-to-Mouth?,” *The Review of Economic Studies*, 2024, 92 (3), 1293–1340.
- Akerlof, George A and Janet L Yellen**, “Can small deviations from rationality make significant differences to economic equilibria?,” *The American Economic Review*, 1985, 75 (4), 708–720.
- and —, “A near-rational model of the business cycle, with wage and price inertia,” *The Quarterly Journal of Economics*, 1985, 100 (Supplement), 823–838.
- Angeletos, George-Marios, Chen Lian, and Christian K Wolf**, “Can Deficits Finance Themselves?,” *Econometrica*, 2024, 92 (5), 1351–1390.
- Auclert, Adrien, Matthew Rognlie, and Ludwig Straub**, “The intertemporal keynesian cross,” *Journal of Political Economy*, 2024, 132 (12), 4068–4121.
- , —, and —, “Fiscal and Monetary Policy with Heterogeneous Agents,” *Annual Review of Economics*, 2025, 17, 539–562.
- Barsky, Robert B, Christopher L House, and Miles S Kimball**, “Sticky-price models and durable goods,” *American Economic Review*, 2007, 97 (3), 984–998.
- Bewley, Truman**, “The optimum quantity of money,” Discussion Paper 383, Northwestern University 1979.
- Boehm, Johannes, Etienne Fize, and Xavier Jaravel**, “Five Facts about MPCs: Evidence from a Randomized Experiment,” *American Economic Review*, 2025, 115 (1), 1–42.
- Bordalo, Pedro, John Conlon, Nicola Gennaioli, Spencer Kwon, and Andrei Shleifer**, “How People Use Statistics,” *Review of Economic Studies*, 2025, *rdaf022*.
- , **Nicola Gennaioli, Yueran Ma, and Andrei Shleifer**, “Overreaction in Macroeconomic Expectations,” *American Economic Review*, 2020, 110 (9), 2748–2782.
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess**, “Revisiting Event-Study Designs: Robust and Efficient Estimation,” *The Review of Economic Studies*, 2024, 91 (6), 3253–3285.
- Boutros, Michael**, “Windfall Income Shocks with Finite Planning Horizons,” *Working Paper*, 2025.
- Brinca, Pedro, Miguel Faria e Castro, Miguel H Ferreira, and Hans Holter**, “The nonlinear effects of fiscal policy,” in “Proceedings. Annual Conference on Taxation and Minutes of the Annual Meeting of the National Tax Association,” Vol. 112 2019, pp. 1–70.

- Broda, Christian and Jonathan A. Parker**, “The Economic Stimulus Payments of 2008 and the aggregate demand for consumption,” *Journal of Monetary Economics*, 2014, 68, S20–S36.
- Browning, Martin and Thomas F Crossley**, “The Life-Cycle Model of Consumption and Saving,” *Journal of Economic Perspectives*, 2001, 15 (3), 3–22.
- Bunn, Philip, Jeanne Le Roux, Kate Reinold, and Paolo Surico**, “The consumption response to positive and negative income shocks,” *Journal of Monetary Economics*, 2018, 96, 1–15.
- Campbell, John Y and N Gregory Mankiw**, “Consumption, income, and interest rates: Reinterpreting the time series evidence,” *NBER Macroeconomics Annual*, 1989, 4, 185–216.
- Carroll, Christopher D**, “Buffer-stock saving and the life cycle/permanent income hypothesis,” *The Quarterly Journal of Economics*, 1997, 112 (1), 1–55.
- Carroll, Christopher, Jiri Slacalek, Kiichi Tokuoka, and Matthew N White**, “The distribution of wealth and the marginal propensity to consume,” *Quantitative Economics*, 2017, 8 (3), 977–1020.
- Chetty, Raj and Adam Szeidl**, “Consumption commitments and habit formation,” *Econometrica*, 2016, 84 (2), 855–890.
- Christelis, Dimitris, Dimitris Georgarakos, Tullio Jappelli, Luigi Pistaferri, and Maarten Van Rooij**, “Asymmetric Consumption Effects of Transitory Income Shocks,” *The Economic Journal*, 2019, 129 (622), 2322–2341.
- Cochrane, John H**, “The sensitivity of tests of the intertemporal allocation of consumption to near-rational alternatives,” *The American Economic Review*, 1989, 79, 319–337.
- Coibion, Olivier and Yuriy Gorodnichenko**, *Expectations Matter: The New Causal Macroeconomics of Surveys and Experiments*, Princeton, NJ: Princeton University Press, 2026.
- , **Dimitris Georgarakos, Yuriy Gorodnichenko, Geoff Kenny, and Michael Weber**, “The effect of macroeconomic uncertainty on household spending,” *American Economic Review*, 2024, 114 (3), 645–677.
- , **Yuriy Gorodnichenko, and Michael Weber**, “How Did U.S. Consumers Use Their Stimulus Payments?,” Technical Report 27693, NBER 2020.
- Colarieti, Roberto, Pierfrancesco Mei, and Stefanie Stantcheva**, “The How and Why of Household Reactions to Income Shocks,” *Working Paper*, 2024.
- Debortoli, Davide and Jordi Galí**, “Heterogeneity and aggregate fluctuations: insights from TANK models,” Technical Report, National Bureau of Economic Research 2024.
- Eichenbaum, Martin, Joao Guerreiro, and Jana Obradovic**, “Ricardian non-equivalence,” Technical Report, Working Paper 2025.
- Fagereng, Andreas, Martin B. Holm, and Gisle J. Natvik**, “MPC Heterogeneity and Household Balance Sheets,” *American Economic Journal: Macroeconomics*, 2021, 13 (4), 1–54.
- Flynn, Joel P, Christina Patterson, and John Sturm**, “Fiscal policy in a networked economy,” Working Paper 29619, NBER 2024.
- Fuster, Andreas, Greg Kaplan, and Basit Zafar**, “What Would You Do with \$500? Spending Responses to Gains, Losses, News and Loans,” *Review of Economic Studies*, 2021, 88 (4), 1760–1795.
- Gabaix, Xavier**, “A sparsity-based model of bounded rationality,” *The Quarterly Journal of Economics*, 2014, 129 (4), 1661–1710.
- , “Behavioral Macroeconomics via Sparse Dynamic Programming,” *Journal of the European Economic Association*, 2023, 21 (6), 2327–2376.
- **and David Laibson**, “The 6D bias and the equity-premium puzzle,” *NBER Macroeconomics Annual*, 2001, 16, 257–312.

- **and Thomas Graeber**, “The Complexity of Economic Decisions,” *Working Paper*, 2024.
- Ganong, Peter and Pascal Noel**, “Consumer Spending during Unemployment: Positive and Normative Implications,” *American Economic Review*, July 2019, 109 (7), 2383–2424.
- , **Damon Jones, Pascal Noel, and Daniel Sullivan**, “The Marginal Propensity to Consume out of Stimulus: Evidence From Administrative Bank Data,” *Working Paper*, 2025.
- , **Pascal Noel, Christina Patterson, Joseph Vavra, and Alexander Weinberg**, “Earnings Instability,” *Working Paper*, 2025.
- Güvenen, Fatih, Sam Schulhofer-Wohl, Jae Song, and Motohiro Yogo**, “Worker betas: Five facts about systematic earnings risk,” *American Economic Review*, 2017, 107 (5), 398–403.
- Hamilton, James D**, “A new approach to the economic analysis of nonstationary time series and the business cycle,” *Econometrica*, 1989, pp. 357–384.
- Ilut, Cosmin and Rosen Valchev**, “Economic agents as imperfect problem solvers,” *The Quarterly Journal of Economics*, 2023, 138 (1), 313–362.
- **and —** , “Learning Optimal Behavior Through Reasoning and Experiences,” Paper 2024.
- Jappelli, Tullio and Luigi Pistaferri**, “Fiscal Policy and MPC Heterogeneity,” *American Economic Journal: Macroeconomics*, October 2014, 6 (4), 107–136.
- **and —** , “Reported MPC and Unobserved Heterogeneity,” *American Economic Journal: Economic Policy*, 2020, 12 (4), 275–297.
- Kaplan, Greg and Giovanni L Violante**, “A model of the consumption response to fiscal stimulus payments,” *Econometrica*, 2014, 82 (4), 1199–1239.
- **and —** , “The marginal propensity to consume in heterogeneous agent models,” *Annual Review of Economics*, 2022, 14 (1), 747–775.
- , **Benjamin Moll, and Giovanni L Violante**, “Monetary policy according to HANK,” *American Economic Review*, 2018, 108 (3), 697–743.
- Kőszegi, Botond and Filip Matějka**, “Choice simplification: A theory of mental budgeting and naive diversification,” *The Quarterly Journal of Economics*, 2020, 135 (2), 1153–1207.
- Kotsogiannis, Christos and Plutarchos Sakellaris**, “MPCs Estimated from Tax Lottery, Survey and Administrative Data,” *Working Paper*, 2025.
- Krusell, Per and Anthony A Smith**, “Rules of thumb in macroeconomic equilibrium a quantitative analysis,” *Journal of Economic Dynamics and Control*, 1996, 20 (4), 527–558.
- , **Burhanettin Kuruşçu, and Anthony A Smith**, “Time orientation and asset prices,” *Journal of Monetary Economics*, 2002, 49 (1), 107–135.
- Kueng, Lorenz**, “Excess Sensitivity of High-Income Consumers,” *The Quarterly Journal of Economics*, 2018, 133 (4), 1693–1751.
- Lewis, Daniel J., Davide Melcangi, and Laura Pilossoph**, “Latent Heterogeneity in the Marginal Propensity to Consume,” Working Paper 32523, NBER 2024.
- Lian, Chen**, “A theory of narrow thinking,” *The Review of Economic Studies*, 2021, 88 (5), 2344–2374.
- , “Mistakes in future consumption, high MPCs now,” *American Economic Review: Insights*, 2023, 5 (4), 563–581.
- Lieder, Falk and Thomas L. Griffiths**, “Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources,” *Behavioral and Brain Sciences*, 2020, 43, e1.
- Lipman, Barton L**, “How to decide how to decide how to ...: Modeling limited rationality,” *Econometrica*, 1991, 59, 1105–1125.
- Lucas, Robert E**, “Econometric policy evaluation: A critique,” in “Carnegie-Rochester Conference Series on Public Policy,” Vol. 1 North-Holland 1976, pp. 19–46.

- Maćkowiak, Bartosz and Mirko Wiederholt**, “Business Cycle Dynamics under Rational Inattention,” *The Review of Economic Studies*, 2015, 82 (4), 1502–1532.
- , **Filip Matějka, and Mirko Wiederholt**, “Rational Inattention: A Review,” *Journal of Economic Literature*, 2023, 61 (1), 226–273.
- Malmendier, Ulrike and Stefan Nagel**, “Depression babies: do macroeconomic experiences affect risk taking?,” *The Quarterly Journal of Economics*, 2011, 126 (1), 373–416.
- Maxted, Peter, David Laibson, and Benjamin Moll**, “Present Bias Amplifies the Household Balance-Sheet Channels of Macroeconomic Policy,” *The Quarterly Journal of Economics*, 2024, 140 (1), 691–743.
- Morduch, Jonathan and Rachel Schneider**, *The Financial Diaries: How American Families Cope in a World of Uncertainty*, Princeton University Press, 2017.
- Oprea, Ryan**, “Complexity and Its Measurement,” *Working Paper*, 2024.
- Orchard, Jacob D, Valerie A Ramey, and Johannes F Wieland**, “Micro MPCs and Macro Counterfactuals: The Case of the 2008 Rebates,” *The Quarterly Journal of Economics*, 2025, 140 (3), 2001–2052.
- Parker, Jonathan A and Nicholas S Souleles**, “Reported effects versus revealed-preference estimates: Evidence from the propensity to spend tax rebates,” *American Economic Review: Insights*, 2019, 1 (3), 273–290.
- , —, **David S Johnson, and Robert McClelland**, “Consumer Spending and the Economic Stimulus Payments of 2008,” *American Economic Review*, 2013, 103 (6), 2530–2553.
- Pfäuti, Oliver and Fabian Seyrich**, “A Behavioral Heterogeneous Agent New Keynesian Model,” Discussion Paper, DIW Berlin 2024.
- Schmidt, Lawrence DW**, “Climbing and falling off the ladder: Asset pricing implications of labor market event risk,” *Journal of Financial Economics*, 2025, 172, 104131.
- Sims, Christopher A**, “Implications of rational inattention,” *Journal of Monetary Economics*, 2003, 50 (3), 665–690.
- Tempalski, Jerry**, “Revenue effects of major tax bills,” *Office of Tax Analysis Working Paper*, 2006, 81.
- Weber, Michael, Francesco D’Acunto, Yuriy Gorodnichenko, and Olivier Coibion**, “The Subjective Inflation Expectations of Households and Firms: Measurement, Determinants, and Implications,” *Journal of Economic Perspectives*, August 2022, 36 (3), 157–184.

Online Appendix

A Additional Theoretical Analysis

A.1 Full Analysis of the Simple Model

In this section, we fully analyze the simple model of Section 2. This supplements the simple example presented in the main text. For expositional simplicity, we impose (as in the example) that $\beta R = 1$.

A.1.1 The Opportunity Cost of Near-Rational Behavior

We first formalize that even very small utility costs can lead near-rational households to behave very differently from households who frictionlessly optimize. To do this, we define the *loss from quick-fixing* in state $z = (y_1, y_2)$ as

$$\mathcal{L}^q(z) \equiv U^*(z) - U^q(z) \tag{8}$$

where we define the value U^* of optimal behavior and U^q of quick-fixing, respectively, via

$$\begin{aligned} U^*(z) &= u(c^*(z)) + \beta u(c^*(z)) \\ U^q(z) &= u(c^q(z)) + \beta u(R(y_1 - c^q(z)) + y_2) \end{aligned} \tag{9}$$

The following result describes the opportunity cost of quick-fixing in a given state compared to behaving optimally, up to a second-order approximation:

Proposition 1 (Second-Order Losses from Near-Rational Behavior). *The loss from following a quick-fix consumption rule indexed by q when the asset-income state is z is:*

$$\mathcal{L}^q(z) = \frac{1}{2}(1 + R)|u''(c^*(z))|(c^q(z) - c^*(z))^2 + O(|c^*(z) - c^q(z)|^3) \tag{10}$$

Proof. We apply a state-dependent second-order approximation to $\mathcal{L}^q(z)$. Define:

$$U(c, z) = u(c) + \beta u(R(y_1 - c) + y_2) \tag{11}$$

and observe that $U^q(z) = U(c^q(z), z)$. We define the loss function on an enriched domain as

Table A.1: Small utility costs can lead to large errors in consumption

Consumption mistake, m	1%	5%	10%	15%	20%
Consumption-equivalent cost, $\tau(m)$	0.01%	0.25%	1.00%	2.22%	3.92%

Notes: This table reports the maximum percentage of consumption that a household would be willing to give up to prefer quick-fixing to optimization when quick-fixing would lead to a relative consumption mistake of size m : $\tau(m) = 1 - e^{-m^2}$.

$\mathcal{L}^U(c, z) = U^*(z) - U(c, z)$ and observe that $\mathcal{L}^q(z) = \mathcal{L}^U(c^q(z), z)$.

To approximate $\mathcal{L}^q$, we therefore approximate $\mathcal{L}^U(\cdot, z)$ to second order around $c^*(z)$ for every state z . This yields the following:

$$\begin{aligned} \mathcal{L}^q(z) &= \mathcal{L}^U(c^*(z), z) + \mathcal{L}_c^U(c^*(z), z)(c^q(z) - c^*(z)) \\ &\quad + \frac{1}{2}\mathcal{L}_{cc}^U(c^*(z), z)(c^q(z) - c^*(z))^2 + O(|c^*(z) - c^q(z)|^3) \end{aligned} \quad (12)$$

We observe however that $\mathcal{L}^U(c^*(z), z) = 0$ by definition and that $\mathcal{L}_c^U(c^*(z), z) = 0$ by the optimality of c^* . Thus,

$$\mathcal{L}^{qi}(z) = \frac{1}{2}\mathcal{L}_{cc}^U(c^*(z), z)(c^q(z) - c^*(z))^2 + O(|c^*(z) - c^q(z)|^3) \quad (13)$$

We can moreover compute $\mathcal{L}_{cc}^U(c^*(z), z)$ as:

$$\mathcal{L}_{cc}^U(c, z) = -\beta R^2 u''(R(y_1 - c) + y_2) - u''(c) = -(1 + \beta R^2)u''(c^*(z)) \quad (14)$$

Completing the proof. □

This result formalizes that there is no first-order loss from deviations from frictionlessly optimizing behavior. This follows from the envelope theorem: when the household makes a small consumption mistake, their first-order condition implies that the slope of their lifetime utility is close to flat in the mistake. Thus, our model generates near-rational behavior by the [Akerlof and Yellen \(1985\)](#) criterion that there is no “first-order” opportunity cost from quick-fixing behavior.

Because of this basic envelope logic, even the presence of small utility costs may motivate a household to follow a quick-fix. To substantiate this, we provide a simple example of how the second-order losses from near-rationality implied by [Proposition 1](#) can lead to small utility costs to generate large differences in behavior.

Example 1 (Small Costs Allow for Large Mistakes). Suppose that $u(c) = \log c$ and $R = 1$. [Proposition 1](#) implies that the payoff loss from following a suboptimal quick-fix is approx-

imately equal to m^2 , where $m = (c^{q_i}(z_i) - c^*(z_i))/c^*(z_i)$ is the household's consumption "mistake" expressed as a proportional deviation from the optimal level. We now ask: what is the smallest utility cost in consumption-equivalent units (*i.e.*, equivalently costly to a τ percent reduction in consumption in one period) that rationalizes making a consumption mistake of a size m ? The utility cost associated with the mistake, up to the second-order approximation of Proposition 1, is $\kappa(m) = m^2$. Putting this into units of a one-period proportionate loss in consumption τ , we have that:

$$\begin{aligned} 2 \log(c^*(z_i)) - \kappa(m) &= \log((1 - \tau)c^*(z_i)) + \log(c^*(z_i)) \\ \implies \tau &= 1 - e^{-m^2} \end{aligned} \tag{15}$$

In Table A.1, we report this utility cost for consumption mistakes ranging from 1% to 20%. Surprisingly, if a household were to have to pay anything less than 0.25% of consumption in order to optimize, then it would be unwilling to optimize even when making a 5% consumption mistake. Thus, even very small costs of optimization can lead to large differences in behavior from the model with frictionless optimization. $\triangle$

A.1.2 The Near-Rational Response to Income Shocks

Having established that quick-fixers may tolerate large deviations in consumption *levels*, we now study how quick-fixers respond to income *shocks*. Formally, we consider a household that is informed that its first-period income will be $y_i(x) = y_i + x$ for some income shock $x \in \mathbb{R}$. What is the optimal near-rational response? By applying Proposition 1, defining the agent's state after the income shock as $z(x) = (y_1 + x, y_2)$, we obtain that:

Corollary 1 (When to Quick-Fix). *Up to a second-order approximation, a household with asset-income state z with quick-fixing type given by q follows the consumption policy function*

$$c(x) = \begin{cases} c^q(z(x)) & \text{if } |c^q(z(x)) - c^*(z(x))| \leq \sqrt{\frac{\kappa_q}{\frac{1}{2}(1+R)|u''(c^*(z(x)))|}}, \\ c^*(z(x)) & \text{otherwise.} \end{cases} \tag{16}$$

Intuitively, households abandon their quick-fix only after a sufficiently large change in circumstances causes the quick-fix to induce a large "mistake." Under natural conditions, this is to say that households quick-fix for small shocks and behave optimally for large shocks. To benchmark the economic significance of this result, we briefly return to the setting of Example 1.

Example 1 (continued). Consider a household whose quick-fix sets consumption to a level appropriate for some "default" state z_0 : that is, $\bar{c} = c^*(z_0)$ and $c^{q_i}(z) \equiv \bar{c}$. Since optimal

consumption c^* is proportional to permanent income, an $x\%$ shock to permanent income starting from z_0 corresponds to an $\frac{x}{1+x}\%$ consumption mistake, which is increasing in the size of the income shock. Moreover, for such a “consumption-fixing” household, Table A.1 can be reinterpreted as giving the maximum shocks to permanent income in response to which a household would persist with the quick-fix. For example, a household with a utility cost that equals 0.25% of their consumption would persist in consuming $\bar{c}$ after a shock to permanent income of less than approximately 5%. Since “large” *transitory* shocks, like generous government stimulus, are potentially “small” shocks to permanent income, quick-fixing can critically shape households’ responses to changes in their economic situation. $\triangle$


What determines whether households quick-fix if and only if shocks are small enough? Income shocks affect households’ decision whether to quick-fix via two mechanisms: the size of the mistake $|c^q(z(x)) - c^*(z(x))|$ and the curvature of utility $|u''(c^*(z(x)))|$. As highlighted before, we restrict our attention to quick-fixes that make mistakes that are increasing in the magnitude of shocks. Under the conventional case with prudence ($u''' \geq 0$), households will quick-fix out of small negative shocks and act optimally for large negative shocks. For large positive shocks, a high degree of prudence can theoretically fight the prediction that households reoptimize: intuitively, they may not care much about mistakes after a large positive shock makes them rich. However, we do not find empirical evidence for these “switch backs,” where households quick-fix for both small and very large positive shocks, and we do not detect this non-convexity of inaction bands in our empirically calibrated model.

A.2 Illustrating why Observational Methods Do Not Suffice

In this section, we illustrate why observational methods can fail to identify the cross-sectional distribution of marginal propensities to consume. At an intuitive level, in observational data, responses to shocks are difficult to disentangle from other noise in spending and income. We show this point directly in a simulation from our simple model from Section 2. Moreover, we describe how this is a specific instance of a classic problem in causal inference: full distributions of treatment effects can be identified only under very specific assumptions (see, *e.g.*, Abbring and Heckman, 2007). In particular, we argue that these assumptions are unlikely to hold in the context we study.

A Simple Example. To make this point transparently, we return to a version of the model in Section 2. A fraction $\alpha = 0.8$ of households are *savings fixers*, whose response to small shocks features an MPC of 1; a complementary fraction $1 - \alpha = 0.2$ are *consumption fixers*, whose response to small shocks features an MPC of 0. The frictionless consumption function is $c^*(y_1) = (y_1 + \bar{y})/2$, where y_1 is the household’s income in the current period and $\bar{y} = 1$

Figure A.1: Observational methods do not measure the MPC distribution



Notes: These figures plot the measured distribution of MPC_i in a simulation from the model. Panel A shows a small income shock with $\tau = 0.04 = 4\%$ and Panel B shows a large income shock with $\tau = 0.10 = 10\%$. The blue dashed lines correspond to households’ true MPCs out of income shocks (zero, one half, one). The solid red line corresponds to the average value of measured MPC_i , $\overline{MPC}$, in each figure. We use the parameters $\alpha = 0.8$, $\bar{y} = 1$, $\sigma = 0.03$, and $\kappa = 0.0006$.

is a “usual” income that is earned in the second period. First-period income deviates from the usual level of 1 by a random shock, $\varepsilon_1 \sim N(0, \sigma^2)$, and a transfer, τ : $y_1 = \bar{y} + \varepsilon_1 + \tau$. The cost of changing behavior is $\kappa = 0.0006 = 0.06\%$, $R = \beta = 1$, and $\sigma = 0.03$ (which is extremely conservative given evidence from [Ganong et al., 2025b](#)). The analyst wants to measure the distribution of household responses to transfers τ , for different values of the transfer. Concretely, they consider a “small transfer” $\tau = 0.04 = 4\%$ and a “large transfer” $\tau = 0.10 = 10\%$.

Panel C of Figure 2 (main text) shows the “true” consumption function for each type of household as a function of total income y_1 , inclusive of both the shock and the transfer. Quick-fixers switch to the frictionless policy function only when income deviates sufficiently from $\bar{y}$. Depending on the type of household and the size of shock, the marginal propensity to consume is either 0, 0.5, or 1.

To measure this distribution of MPCs, an analyst collects data from an experiment in which households are randomly allocated to receive: (i) no transfer, (ii) a small transfer $\tau = 4\%$, or (iii) a large transfer $\tau = 10\%$. The analyst computes individual marginal propensities to consume among the treated as $MPC_i = \frac{c_{1i} - \bar{c}_1}{\tau_i}$, where c_{1i} is individual consumption, $\bar{c}_1$ is the average consumption in the control group, and τ_i is the individually received transfer.

Figure A.1 shows the measured distributions of MPC_i for each transfer size. These data suggest that MPCs are “smoothly” distributed, rather than concentrated at mass points corresponding to the three true MPCs (0, 0.5, and 1). Particularly striking is the fact that the histogram in Panel A lacks a mass point at 1 when, in fact, 80% of households have an MPC of 1 in response to small shocks. Moreover, both histograms place mass on specific

MPCs that no household has. These include MPCs outside of $[0, 1]$ that are not possible given the underlying consumption functions (Figure 2). Thus, these data can fail to detect true MPC heterogeneity and can imply the existence of MPC heterogeneity that does not exist.

The key reason that observational data produce these results is that households receive other shocks that change their spending behavior. For instance, a consumption-fixing household with a true MPC of 0 out of a small shock might receive a positive income shock of $\varepsilon_1 = 6\%$ alongside their transfer of $\tau = 4\%$. Such a household would no longer quick-fix and instead spend half of the shock, or 5% of their average income. The analyst would measure an MPC of $\Delta c_i / \tau_i = 5\% / 4\% = 1.2$, which is not the correct propensity to consume out of shocks around $\bar{y}$ for any shock size. This basic idea explains both why mass points may fail to appear at the “true” MPCs and also why mass may appear at “spurious” MPCs.

The pervasiveness of this phenomenon depends, intuitively, on whether or not analysts can condition on the variables that span what we called ε_1 : shocks to households’ income, wealth, financial situation, and spending needs. Econometricians in practice do have access to other data that may help in this regard, such as current income or past spending for the same household. The latter would allow the analyst to control for a household fixed effect or, equivalently, estimate MPCs from a “difference in differences” estimator. But we argue that these are highly unlikely to remove all background noise in household spending, especially given the complexity of households’ financial situation (Morduch and Schneider, 2017; Ganong et al., 2025b), the known difficulty of accurately measuring wealth and its liquidity, and the near impossibility of fully accounting for all sources of news about future income.

Nonetheless, although observational data cannot detect MPC heterogeneity and near-rationality, these phenomena may still be highly relevant for aggregate behavior. In our example, the aggregate MPC out of the smaller shock (0.6) is 20% higher than that out of the larger shock (0.5). This happens because households are more likely to quick-fix for small shocks than large shocks, and most households’ quick-fixing MPCs are higher than the rational MPC. This can have substantial macroeconomic consequences, as we illustrate in Section 5.

Relationship to Identifying Distributions of Treatment Effects. The issues discussed above are the contextual analog of the general observation that, in causal inference settings, the analyst observes the marginal distributions of potential outcomes with and without treatment, but not the *joint* distribution of these potential outcomes (Abbring and Heckman, 2007). Fundamentally, the joint distribution can never be directly observed because we cannot simultaneously observe the same individual’s potential outcomes with and

without the treatment. Identifying the distribution of treatment effects requires making restrictive assumptions on the structure of this joint distribution.

To describe this phenomenon in the context of consumption behavior, it suffices to ignore quick-fixing and the dependence of MPCs on shock sizes altogether, and to consider the simple linear consumption model,

$$c_i = \bar{c} + m_i(\tau D_i + \varepsilon_i) \quad (17)$$

where $\bar{c}$ is some constant, m_i is a potentially person-specific MPC, τ is a transfer of a fixed size, $D_i \in \{0, 1\}$ is an indicator for whether household i receives a transfer, and ε_i is a background shock to (permanent) income. To build a bridge to the relevant literature in causal inference, we describe household spending with a potential outcomes framework. The potential outcomes for household i 's spending with and without the transfer are, respectively,

$$\begin{aligned} C_{1,i} &= \bar{c} + m_i\tau + m_i\varepsilon_i \\ C_{0,i} &= \bar{c} + m_i\varepsilon_i \end{aligned} \quad (18)$$

The individual level treatment effect is therefore

$$\Delta_i = C_{1,i} - C_{0,i} = m_i\tau \quad (19)$$

The distribution of treatment effects is thus a rescaling of the distribution of marginal propensities to consume m_i .

In a randomized experiment, an econometrician observes the marginal distribution of $C_{1,i}$ and $C_{0,i}$ from two otherwise equivalent populations that are and are not exposed to the treatment, respectively. What can the econometrician learn about the distribution of Δ_i ? At an intuitive level, the key problem is that the econometrician cannot tell if treated households with high spending have a high MPC m_i or happened to draw a high income shock ε_i , since they cannot observe how that *exact* household would have acted absent the transfer. This missing data problem corresponds to not knowing the *joint* distribution of $C_{0,i}$ and $C_{1,i}$ across households i indexed by their MPC and their income shock.

Several methods exist to remedy this problem via structural assumptions. [Abbring and Heckman \(2007\)](#) (Section 2) provide a comprehensive review of these methods and group them into two categories: (i) those that postulate additional structure on the joint distribution of potential outcomes and (ii) those that collect additional choice data that can help reveal (aspects of) the joint distribution. We argue that known methods in category (i), such as those reviewed by [Abbring and Heckman \(2007\)](#), rely on assumptions that may not

hold in a standard consumption-savings setting.

One approach is to study how treatments affect the entire spending distribution via quantile regressions, as pursued by [Misra and Surico \(2014\)](#). These methods identify the distribution of treatment effects only under the assumption that the distribution of potential outcomes is *rank preserving*: that is, letting F_0 and F_1 denote the marginal CDFs of $C_{0,i}$ and $C_{1,i}$, $F_0^{-1}(C_{0,i}) = F_1^{-1}(C_{1,i})$ for each household i (see, *e.g.*, [Heckman, Smith and Clements, 1997](#)). This is violated when there are nontrivial income shocks and MPC heterogeneity: someone with a high MPC who happened to get a low income shock may be a relatively low spender in the F_0 distribution, but may be a relatively high spender in the F_1 distribution.¹⁸

A second approach is to apply non-parametric *deconvolution* methods to the distribution of spending for the treated and untreated, as pursued by [Boehm et al. \(2025\)](#). These methods rely instead on the assumption that treatment effects, $\Delta_i = C_{1i} - C_{0i}$, are independent from consumption if untreated, C_{0i} (see Section 2.5.4 and Appendix A of [Abbring and Heckman, 2007](#)). This boils down to assuming that $m_i \perp m_i \varepsilon_i$. This is essentially impossible due to the dependence of both terms on m_i : whether you spend out of the transfer is not independent from whether you would have spent without the transfer, because of the common dependence on the MPC. Even if, for example, MPCs m_i and income shocks ε_i were independent from one another, the condition $m_i \perp m_i \varepsilon_i$ would hold only if the distributions were degenerate.

A final approach is to project treatment effect heterogeneity onto observables. In practice, this would be done by interacting treatment status in experiments with observable variables like income, wealth, or demographics, and estimating average MPCs conditional on these observables. But this can provide only a lower bound for MPC heterogeneity. Moreover, this bound will be less informative if MPC heterogeneity is not well explained by observables.

Quick-fixing complicates these strategies even further. Under this model, MPCs depend themselves on the size of transfers and background shocks. This further casts doubt on the rank-invariance and treatment-effect-independence conditions that are respectively crucial for the quantile regression and deconvolution approaches. Moreover, according to our empirical findings, observable characteristics are poor predictors of quick-fixing behavior and MPC heterogeneity (see Fact 3 in Section 3.4). This suggests that methods that project onto observables might significantly understate the amount of true MPC heterogeneity.

Our survey-based method can be understood as a method for directly eliciting the dis-

¹⁸Concretely, consider income shocks $\varepsilon_i \in \{0,1\}$ that are independent from MPCs $m_i \in \{0,1\}$, and a transfer $\tau = 1$. There are four types of household, indexed by their income shock (0 or 1) and MPC (0 or 1). Consider a specific household i that has $m_i = 1$ and $\varepsilon_i = 0$. Absent the transfer, they have the lowest possible spending in the distribution: $c_i = C_{0,i} = 0$. In the distribution F_0 , they have the same rank as both types of household with a 0 MPC, who also spend nothing. But, conditional on getting the transfer, the household i spends $c_i = C_{1,i} = 1$. In this distribution, household i spends more than both types of household with a 0 MPC. Thus, $F_0^{-1}(C_{0,i}) \neq F_1^{-1}(C_{1,i})$, violating the rank invariance condition.

tribution of individual treatment effects (or MPCs). Although these methods of course have their own limitations (as discussed in Section 3), they do sidestep the fundamental identification issues described above. Moreover, they allow us to identify the within-household distribution of treatment effects (MPCs) across different treatments (shock sizes), which is important for disciplining our theory.

A.3 One-Time Shocks in the Quantitative Model

In this section, we describe one-time, unanticipated shocks in the quantitative model. We use this experiment to compute reoptimization fractions in the model and match these moments to the data (see Section 4.2 and Equation 5).


The household contemplates a shock $x \in \mathbb{R}$ in an interim period after initially choosing whether to quick-fix or reoptimize in a given period, but before observing income or making decisions for the next period. With some abuse of notation, we write c^{q*} , $\bar{c}^*$, and $\bar{y}^*$ as the optimal choices in the original decision period, suppressing dependence on the household's state. The household's problem, when faced with a shock x , is

$$\begin{aligned}
& \max_{D_x \in \{0,1\}} \{ D_x (u(c^*(a, y + x)) + \beta \mathbb{E} [V^q(a', y', c^*(a, y + x), y) \mid y] - \kappa_q) \\
& \quad + (1 - D_x) (u(c^q(c^{q*}, x)) + \beta \mathbb{E} [V^q(a', y', \bar{c}^*, \bar{y}^*) \mid y]) \} \\
& \text{s.t. } a' = Ra + y - (D_x(c^*(a, y + x)) + (1 - D_x)(c^q(c^{q*}, x))) \\
& \quad a' \geq 0
\end{aligned} \tag{20}$$

If a household reoptimizes, then it follows the optimal, forward-looking behavior embodied in c^* . This entails a utility cost. If the household quick-fixes, then it treats c^{q*} as its reference consumption and the shock x as the income shock (Equation 3).

B Additional Figures and Tables

Figure B.1: Results are robust in a variety of sensitivity tests




Notes: We recalculate key statistics from Section 3 for a variety of robustness specifications:

- *Main sample:* We reproduce the statistics from the main text.
- *Weighted sample:* We use post-stratification weights that correct for possible imbalances across the variables reported in Table C.1.
- *Including MPCs outside [0, 1]:* We add 51 additional respondents whom we drop from the main analysis because they report MPCs outside [0, 1] (see Appendix C.1).
- *Drop potential duplicates:* We drop potential duplicate respondents who submitted similar data on the same day (see Appendix C.1).
- *Excluding the faster half (based on intro):* We exclude the 50% fastest respondents who “speed through” the introductory instructions of the survey.
- *Start with gains:* We restrict the sample to respondents who first respond to gains.
- *Start with losses:* We restrict the sample to respondents who first respond to losses.
- *Start with small shocks (\$50 or \$100):* We restrict the sample to respondents who first respond to a small income shock of \$50, \$100, −\$50, or −\$100.
- *Start with large shocks (\$5,000 or \$10,000):* We restrict the sample to respondents who first respond to a large income shock of \$5,000, \$10,000, −\$5,000, or −\$10,000.
- *Only the first scenario:* We restrict the sample to the first MPC that respondents report.
- *Only the first five scenarios:* We restrict the sample to the first five MPCs that respondents report.


*Small shocks: \$50 and \$100. Large shocks: \$5,000 and \$10,000.

Figure B.2: MPC profiles across the wealth distribution




Notes: The alluvial graphs summarize the MPC profiles of households with varying (i) liquid wealth, (ii) illiquid wealth, and (iii) debt, (iv) net wealth, (v) liquid wealth to income ratio (Kaplan, Violante and Weidner, 2014), and (vi) net wealth to income ratio (see Appendix C.3 for variable definitions). In each panel, each of the 14 columns displays the distribution of MPCs for one particular shock size, and the streams between bars indicate how households' MPCs transition between two neighboring shocks.

Figure B.3: Distribution of MPCs for losses before and after switching to interior



Notes: These histograms show conditional distributions of MPCs for losses. The rows correspond to the four quick-fixing types. The first column shows the distribution of MPCs before households switch to an interior MPC, which by construction puts all mass at either MPC = 0 or MPC = 1. The second column shows the conditional distribution of MPCs (given type and shock size) for the first shock for which the respondent reports an interior value. An analogous analysis for gains is reported in Figure 5.

Figure B.4: Switching thresholds are higher among high-wealth households



Notes: The figure displays the average switching threshold—*i.e.*, the smallest shock for which quick-fixing households switch from an extreme to an interior MPC, averaged across gains and losses—for high-wealth households and other households.

Table B.1: Exploring the variation in MPCs across households

Households' average MPCs (across shocks +/- \$100 to 10,000)						
	Average MPC		Share with MPC = 0		Share with MPC = 1	
	(1)	(2)	(3)	(4)	(5)	(6)
Consumption and income						
Monthly spending (log.)		0.012*** (0.003)		-0.015*** (0.004)		0.007** (0.003)
Annual income (log.)		0.008 (0.005)		0.005 (0.006)		0.001 (0.006)
Income risk (std.)		0.023*** (0.004)		-0.037*** (0.004)		-0.003 (0.004)
Liquid wealth: dummies with reference group: [0k, 1k)						
[1k, 10k)		-0.054*** (0.009)		0.001 (0.010)		-0.078*** (0.010)
[10k, 100k)		-0.071*** (0.010)		0.014 (0.012)		-0.096*** (0.012)
[100k, more)		-0.071*** (0.013)		0.057*** (0.015)		-0.064*** (0.014)
Illiquid wealth: dummies with reference group: [0k, 10k)						
[10k, 100k)		-0.019* (0.010)		0.009 (0.011)		-0.028** (0.012)
[100k, 500k)		-0.033*** (0.010)		0.050*** (0.011)		-0.000 (0.011)
[500k, more)		-0.078*** (0.012)		0.105*** (0.014)		0.004 (0.014)
Debt: dummies with reference group: [0k, 1k)						
[1k, 10k)		0.028*** (0.008)		-0.060*** (0.010)		-0.008 (0.010)
[10k, 100k)		0.036*** (0.009)		-0.055*** (0.010)		0.022** (0.010)
[100k, more)		0.014 (0.010)		-0.028** (0.012)		0.007 (0.011)
Other characteristics						
College		0.004 (0.007)		-0.022*** (0.008)		-0.013 (0.008)
Age (in 10y)		-0.017*** (0.002)		0.036*** (0.002)		0.004* (0.002)
Female respondent		0.008 (0.006)		-0.013* (0.007)		0.004 (0.007)
Household size		0.010*** (0.003)		-0.012*** (0.003)		0.002 (0.003)
Quick-fixing types						
Consumption fixer	-0.318*** (0.008)		0.573*** (0.012)		0.009 (0.006)	
Savings fixer	0.112*** (0.007)		0.119*** (0.007)		0.383*** (0.008)	
Consumption prioritizer	-0.107*** (0.009)		0.302*** (0.012)		0.147*** (0.009)	
Savings prioritizer	-0.083*** (0.007)		0.313*** (0.009)		0.208*** (0.008)	
Constant	0.487*** (0.003)	0.381*** (0.050)	0.047*** (0.003)	0.147** (0.058)	0.051*** (0.003)	0.177*** (0.058)
Obs.	4,981	4,981	4,981	4,981	4,981	4,981
R ²	0.350	0.129	0.444	0.212	0.406	0.026

Notes: This table reports regressions that explore the heterogeneity of households' MPCs. Columns 1–2 analyze households' mean MPC (averaged across the 12 losses or gains ranging from \$100 to \$10,000), Columns 3–4 analyze households' share of MPCs that equal 0 (among the same 12 shocks), and Columns 5–6 analyze households' share of MPCs that equal 1 (among the same 12 shocks). Appendix C.3 describes how we measure the economic background variables. Robust standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table B.2: Exploring the variation in quick-fixing types across households

	Type membership				Switching point		
	Consumption fixer	Savings fixer	Consumption prioritizer	Savings prioritizer	Unclassified	Average log(shock size)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Consumption and income							
Monthly spending (log.)	−0.012*** (0.004)	0.011** (0.005)	0.005 (0.003)	0.000 (0.004)	−0.005 (0.005)		−0.091*** (0.023)
Annual income (log.)	0.005 (0.007)	0.000 (0.010)	−0.009 (0.007)	−0.018** (0.008)	0.022** (0.010)		0.126*** (0.044)
Income risk (std.)	−0.022*** (0.005)	−0.021*** (0.007)	−0.020*** (0.004)	−0.022*** (0.005)	0.084*** (0.007)		0.057* (0.031)
Liquid wealth: dummies with reference group: [0k, 1k)							
[1k, 10k)	0.012 (0.013)	−0.051*** (0.019)	0.003 (0.013)	−0.010 (0.015)	0.046** (0.018)		−0.438*** (0.082)
[10k, 100k)	0.020 (0.015)	−0.086*** (0.020)	0.005 (0.015)	−0.026 (0.017)	0.088*** (0.019)		−0.205** (0.087)
[100k, more)	0.070*** (0.020)	−0.088*** (0.026)	0.037** (0.019)	−0.045** (0.020)	0.026 (0.023)		0.052 (0.104)
Illiquid wealth: dummies with reference group: [0k, 10k)							
[10k, 100k)	−0.002 (0.015)	−0.000 (0.022)	−0.003 (0.014)	0.015 (0.017)	−0.009 (0.021)		−0.053 (0.100)
[100k, 500k)	0.001 (0.014)	0.050** (0.020)	0.016 (0.014)	0.020 (0.016)	−0.086*** (0.019)		0.111 (0.084)
[500k, more)	0.061*** (0.019)	0.082*** (0.024)	0.003 (0.017)	0.043** (0.019)	−0.189*** (0.021)		0.141 (0.098)
Debt: dummies with reference group: [0k, 1k)							
[1k, 10k)	−0.057*** (0.013)	−0.002 (0.018)	0.008 (0.012)	0.009 (0.014)	0.042** (0.018)		−0.364*** (0.083)
[10k, 100k)	−0.051*** (0.013)	0.046** (0.018)	−0.000 (0.012)	0.017 (0.014)	−0.012 (0.017)		−0.269*** (0.076)
[100k, more)	−0.039** (0.016)	0.045** (0.020)	0.034** (0.015)	0.038** (0.016)	−0.077*** (0.018)		−0.363*** (0.080)
Other characteristics							
College	−0.011 (0.011)	−0.021 (0.015)	0.007 (0.010)	−0.019* (0.011)	0.045*** (0.013)		−0.005 (0.061)
Age (in 10y)	0.016*** (0.003)	0.010** (0.004)	0.011*** (0.003)	0.007** (0.003)	−0.045*** (0.004)		0.100*** (0.018)
Female respondent	−0.009 (0.010)	0.015 (0.013)	−0.017* (0.009)	−0.003 (0.010)	0.014 (0.012)		−0.011 (0.056)
Household size	−0.013*** (0.004)	0.005 (0.005)	0.003 (0.004)	−0.007 (0.004)	0.012** (0.005)		−0.029 (0.025)
Quick-fixing types							
Consumption fixer						0.236*** (0.085)	
Savings fixer						−0.218*** (0.073)	
Consumption prioritizer						−0.555*** (0.093)	
Constant	0.138* (0.077)	0.134 (0.099)	0.089 (0.067)	0.324*** (0.082)	0.315*** (0.103)	7.960*** (0.059)	6.987*** (0.459)
Obs.	4,981	4,981	4,981	4,981	4,981	3,381	3,381
R ²	0.057	0.017	0.020	0.020	0.174	0.024	0.058

Notes: This table reports regressions that explore the heterogeneity of households' quick-fixing types. Columns 1–5 analyze households' type (binary indicators), and Columns 6–7 analyze households' mean log switching threshold (the smallest shock for which they switch from an extreme to an interior MPC, averaged across gains and losses). Columns 6–7 restrict the sample to the four quick-fixing types. Appendix C.3 describes how we measure the economic background variables. Robust standard errors are in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table B.3: Deliberation negatively predicts extreme MPCs (deliberation ratings study)

Extreme MPC of 0 or 1 (binary indicator)						
	(1)	(2)	(3)	(4)	(5)	(6)
Deliberation (std.)	−0.263*** (0.011)	−0.274*** (0.020)	−0.248*** (0.011)	−0.254*** (0.020)	−0.257*** (0.013)	−0.276*** (0.020)
Respondent FE	✓	✓	✓	✓	✓	✓
Weights	–	✓	–	✓	–	✓
Measure	Carefully consider how to change spending		Assess overall financial situation		Discuss with household members	
Observations	3,619	3,619	3,619	3,619	3,080	3,080
R ²	0.740	0.761	0.723	0.744	0.711	0.719


Notes: This table reports regression results and uses data from the deliberation ratings study. We regress a binary indicator for whether a household adopts an extreme MPC of 0 or 1 on different standardized deliberation measures (see row “Measure”). Columns 1, 3, and 5 display unweighted results. Columns 2, 4, and 6 use post-stratification weights that correct for imbalances in the distribution of demographic characteristics (see Table C.2). All regressions use household-level fixed effects. The standard errors in parentheses are robust and clustered at the household level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table B.4: Response time negatively predicts extreme MPCs (main study)

Extreme MPC of 0 or 1 (binary indicator)				
	(1)	(2)	(3)	(4)
Response time (in 10s)	−0.060*** (0.001)	−0.078*** (0.002)	−0.068*** (0.002)	−0.094*** (0.002)
Respondent FE	✓	✓	✓	✓
Order FE			✓	✓
Sample	Full	Only quick-fixers	Full	Only quick-fixers
Observations	69,734	47,334	69,734	47,334


Notes: This table reports regression results and uses data from the main study. We regress a binary indicator for whether a household adopts an extreme MPC of 0 or 1 on the time respondents spent on the shock scenario. We winsorize response time at its 95% quantile. Columns 1 and 3 use the full data, while Columns 2 and 4 report results for quick-fixing households only. Columns 3 and 4 include fixed effects for order, i.e. fourteen dummies that indicate whether respondents make their first, second, third, ... decision. All regressions use household-level fixed effects. The standard errors are robust and clustered at the household level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Figure B.5: Quantitative model fit for frictionlessly optimizing agents' MPCs



Notes: The line graphs compare empirical measurements and model predictions for the average marginal propensity to consume out of shocks of different amounts, among Unclassified respondents in the survey (black line and dots) and frictionlessly optimizing agents in the model (blue line and dots). We calibrate the discount factor β to minimize the sum of squared differences between these model predictions and measurements.


Figure B.6: Quantitative model fit for the MPC distribution



Notes: The alluvial graph summarizes the MPC predictions generated by the quantitative model, following the format of Figure 1. Each of the 14 columns displays the distribution of MPCs for one particular shock size, with colors indicating the size of the MPC. The streams between bars indicate how households' MPCs transition between two neighboring shocks. Black dots depict the average MPCs for each shock. We exclude a few respondents with MPCs outside $[0, 1]$ to facilitate the visual presentation.


Figure B.7: Dynamic consumption responses to transfer shocks

(i) Profiles of intertemporal marginal propensities to consume




Notes: Each panel shows the intertemporal marginal propensity to consume (iMPC) in response to a different stimulus experiment. Each line corresponds to a different household type.

(ii) Profiles of delayed reoptimization



Notes: Each panel shows the fraction of additional households that reoptimize compared to the steady state in response to a different stimulus experiment. The fraction at $t = 1$ is calculated based on whether households reoptimize in response to an income shock *or* the unanticipated stimulus. Each line corresponds to a different household type.

(iii) Comparing iMPC profiles with the model with frictionless optimization



Notes: Each panel shows the intertemporal marginal propensity to consume (iMPC) averaged over households of all types in response to the specified shocks (x-axis) at the specified horizon. Blue dots correspond to the model prediction. Green dots correspond to the subset of frictionlessly optimizing agents.

C Additional Material for the Empirical Studies

C.1 Sample

Sampling. We recruited respondents in October and November 2023 collaborating with the survey company Bilendi. We recruited respondents from different parts of the Bilendi respondent pool in order to approximate the general US population in terms of gender, age, income, education, and region.

Final Sample Characteristics. Table C.1 presents demographic summary statistics for our final sample and compares them to the demographic characteristics of the US adult population.

Exclusion Criteria. Three exclusion criteria are preregistered. The sample does not contain the following responses: (i) incomplete responses, (ii) responses at both extreme 1% tails in the response duration, and (iii) responses with duplicate IDs (very rare cases).

In addition, we exclude 51 respondents who have at least one MPC outside the interval $[0, 1]$. Many of these respondents report just one or a few MPCs outside $[0, 1]$, which could simply reflect response error. Excluding them simplifies the visual presentation of the results and ensures that outliers do not distort our analyses of averages. Unsurprisingly, the robustness check in Figure B.1 confirms that we obtain virtually the same results with the full sample.

Attention Screener. Only participants who pass an attention screener at the beginning of the survey can proceed to the main part of the survey.

Potential duplicate responses. Even though we included a captcha and an attention screener, we observe a couple of very similar respondents who start the survey at a similar time. About 75 respondents have identical answers to 23 different demographic questions and start the survey at a similar time. Fortunately, our results are robust to excluding them. Figure B.1 takes an even more conservative approach and drops roughly 7% of responses with the most similar demographic data within each day. Again, the results are virtually identical.

Table C.1: Demographic characteristics of the sample

Variable	ACS (2022)	Sample
Gender		
Female	50%	50%
Age		
18-34	29%	27%
35-54	32%	33%
55+	38%	40%
Household income		
Below 50k	34%	34%
50k-100k	29%	28%
Above 100k	37%	37%
Education		
Bachelor's degree or more	33%	40%
Region		
Northeast	17%	17%
Midwest	21%	21%
South	39%	39%
West	24%	23%
Sample size	1,980,550	4,981
Variable	SCF (2022)	Sample
Liquid assets		
Below 1k	20%	29%
1k-10k	31%	25%
10k-100k	31%	28%
Above 100k	19%	18%
Illiquid assets		
Below 10k	26%	38%
10k-100k	11%	14%
100k-500k	34%	26%
Above 500k	29%	22%
Debt		
Below 1k	27%	35%
1k-10k	10%	21%
10k-100k	27%	25%
Above 100k	36%	19%
Sample size	4,602	4,981

Notes: This table presents summary statistics for the sample of US households and compares them to benchmark characteristics for the US adult population based on data from the American Community Survey 2022 and the Survey of Consumer Finances 2022. Appendix C.3 describes how we measure the economic background variables.

Table C.2: Demographic characteristics in the additional studies

Variable	ACS (2022)	Deliberation ratings study	Qualitative study
Gender			
Female	50%	50%	49%
Age			
18-34	29%	43%	47%
35-54	32%	48%	46%
55+	38%	10%	7%
Household income			
Below 50k	34%	25%	32%
50k-100k	29%	40%	35%
Above 100k	37%	35%	33%
Education			
Bachelor's degree or more	33%	63%	64%
Region			
Northeast	17%	19%	18%
Midwest	21%	15%	17%
South	39%	49%	41%
West	24%	18%	25%
Sample size	1,980,550	517	502
Variable	SCF (2022)	Deliberation ratings study	Qualitative study
Liquid assets			
Below 1k	20%	26%	26%
1k-10k	31%	32%	27%
10k-100k	31%	32%	37%
Above 100k	19%	10%	10%
Illiquid assets			
Below 10k	26%	38%	38%
10k-100k	11%	21%	19%
100k-500k	34%	24%	25%
Above 500k	29%	18%	19%
Debt			
Below 1k	27%	20%	22%
1k-10k	10%	23%	18%
10k-100k	27%	26%	30%
Above 100k	36%	30%	30%
Sample size	4,602	517	502

Notes: This table presents summary statistics for the sample of US households in the additional studies and compares them to benchmark characteristics for the US adult population based on data from the American Community Survey 2022 and the Survey of Consumer Finances 2022. Appendix C.3 describes how we measure the economic background variables.

C.2 Instructions

The complete instructions are available online at <https://osf.io/2s7cf>. The survey begins with a participation information and informed consent form. Respondents who participate on a mobile device are screened out. Next, respondents have to pass an attention check. Subsequently, respondents fill out a block of demographic questions. Then, the main part of the survey begins (see below). The survey ends with additional questions on households' economic situation.

Introduction

In this survey, we are seeking to understand how your household reacts to unanticipated changes in income. By “household”, we mean everyone who usually lives with you in your primary residence including yourself (but excluding roommates and renters).

You will be presented with various hypothetical scenarios that involve shifts in your income, and we will ask you how such changes would impact your household's spending and saving. Below, we provide a short description of what we mean by “spending” and “saving”. Please read them carefully.

Spending: Spending includes all money spent on goods and services, including rent. Goods include durable goods (such as electronics, furniture, or car maintenance) and nondurable goods (such as groceries, vacations, or gasoline).

Saving: Saving means that, instead of using money today, you reserve it for future use. Examples of savings include cash reserves, money in bank accounts, retirement accounts, financial assets, or real estate. **Repaying debt is also an important form of saving.** By repaying debt today, you owe less money in the future, which means that more money is available for future use.

On the next pages, you will consider hypothetical situations where your household unexpectedly receives a

one-time payment today.

That is, your household's income will be higher for one month due to a one-time payment. The one-time payment comes unexpectedly.

Comment: We randomize whether income losses or gains are displayed first.

A one-time payment

Situation 1

Consider a hypothetical situation where your household unexpectedly receives a

one-time payment of \$250 today.

How would this one-time extra income cause your household to change its spending and saving over the next three months?

Note: Your responses need to add up to \$250.

Enter \$0 if your household's spending/saving would not change.

Enter negative numbers for *decreases* in your household's spending/saving.

Increase in spending

(By how much) would your household increase its monthly spending over the next three months?

\$

Increase in saving

(By how much) would your household increase its monthly saving (which includes increases in debt repayment or decreases in debt-taking) over the next three months?

\$

Total

\$

Comment: Situation 2–7 are analogous. We randomize the order of shock sizes. Each respondent faces seven shocks: \$50, \$100, \$250, \$500, \$1,000, \$5,000, and \$10,000.

On the next pages, you will consider hypothetical situations where your household unexpectedly incurs a

one-time income loss today.

That is, your household's income will be lower for one month due to a one-time income loss. The one-time income loss comes unexpectedly.

A one-time income loss

Situation 1

Consider a hypothetical situation where your household unexpectedly incurs a

one-time income loss of \$100 today.

How would this one-time income loss cause your household to change its spending and saving over the next three months?

Note: Your responses need to add up to \$100.

Enter \$0 if your household's spending/saving would not change.

Enter negative numbers for *increases* in your household's spending/saving.

Decrease in spending

(By how much) would your household decrease its monthly spending over the next three months?

\$

Decrease in saving

(By how much) would your household decrease its monthly saving (which includes decreases in debt repayment or increases in debt-taking) over the next three months?

\$

Total

\$

Comment: Situation 2–7 are analogous. We randomize the order of shock sizes. Each respondent faces seven shocks: \$50, \$100, \$250, \$500, \$1,000, \$5,000, and \$10,000.

Deliberation study

The complete instructions are available online at <https://osf.io/2s7cf>. Below, we show the example screen for a \$50 income gain.

A one-time payment

Situation 1

Consider a hypothetical situation where your household unexpectedly receives a

one-time payment of \$50 today.

How would this one-time extra income cause your household to change its spending and saving over the next three months?

Note: Your responses need to add up to \$50.

Enter \$0 if your household's spending/saving would not change.

Enter negative numbers for *decreases* in your household's spending/saving.

Increase in spending

(By how much) would your household increase its monthly spending over the next three months?

\$

Increase in saving

(By how much) would your household increase its monthly saving (which includes increases in debt repayment or decreases in debt-taking) over the next three months?

\$

Total

\$

In response to the unexpected one-time payment of \$50 ...

On a scale from 1 (not at all) to 6 (very carefully), how carefully would your household consider how to change its spending and saving?

Would not consider it at all. 1	2	3	4	5	Would consider it very carefully. 6
------------------------------------	---	---	---	---	--

In response to the unexpected one-time payment of \$50 ...

What is the percent chance that you would discuss with other household members (like your partner) how your household should change its spending and saving?

0%	1-20%	21-40%	41-60%	61-80%	81-99%	100%
----	-------	--------	--------	--------	--------	------

In response to the unexpected one-time payment of \$50 ...

What is the percent chance that you would assess and consider your household's overall financial situation prior to deciding how to change your household's spending and saving?

0%	1-20%	21-40%	41-60%	61-80%	81-99%	100%
----	-------	--------	--------	--------	--------	------

C.3 Definition of Additional Variables

Age: Age of the respondent.

Education: Highest education level of the respondent.

Gender: Gender of the respondent.

Household size: Size of the respondent’s household, capped at 10 to account for outliers.

Income, annual: Household income in 2022 before taxes and transfers.

Income risk: Households indicate whether their monthly household income varies by less than 5% (1), between 5% and 10% (2), between 10% and 25% (3), or by more than 25% (4). We derive a standardized index based on ordinal responses (1–4).

Monthly spending: Household spending (in contrast to saving and debt repayment) in a typical month, capped at the 95% quantile to account for outliers.

Region: Census region.

Wealth, Liquid: The total value of a household’s financial savings and investments, such as cash holdings, checking and savings accounts, money market funds, government/municipal bonds or treasury bills, stocks and bonds in publicly held corporations, stock and bond mutual funds.

Wealth, Illiquid: The sum of (i) the total value of the land and real estate a household owns, including primary residence, second homes and other real estate, and (ii) the total value of a household’s currently non-withdrawable financial savings and investments, such as the value of your retirement accounts (401(k)s, IRAs, thrift accounts, and future pensions), the cash value of life insurance policies, certificates of deposit, and saving bonds.

Wealth, Debt: Total household debt including credit card debt, mortgages, and other debt, such as student loans, auto loans, and personal loans.

Wealth, Net: Liquid wealth + illiquid wealth – debt.

C.4 Comparison of Cross-Sectional Results to Previous Work

In the cross-section of MPCs, our data replicate many patterns that are familiar to the literature. This appendix section compares our cross-sectional results to related work. It is important to keep in mind that we estimate households’ marginal propensity to consume over a three-month horizon in response to unexpected one-time income shocks and that our survey-based consumption measure includes both nondurable and durable consumption. This means we do not measure *notional consumption* as defined by [Laibson, Maxted and Moll \(2022\)](#) but *consumption expenditures*, which is common in the literature.

High average MPCs. The average MPC in our data is 0.47, but the comparison to other estimates becomes easier if we focus on the MPC to a larger income gain, *e.g.*, the \$1,000 shock, for which we estimate an MPC of 0.35 (Figure 1).

This estimate is within the range of typical estimates in the literature. Using a survey-based approach, Jappelli and Pistaferri (2014) estimate an MPC of 0.48 in Italy, Christelis et al. (2019) estimate an MPC of 0.39 in the Netherlands, Drescher, Fessler and Lindner (2020) find MPCs ranging from 0.33 to 0.57 in 17 European countries over the first twelve months, and Colarieti et al. (2024) estimate an MPC of 0.16 over the first quarter, which continues to increase over subsequent months. An exception is Fuster et al. (2021) who observe an MPC of 0.07 for \$500, mainly because 74% of respondents report an MPC of 0. Studying consumption responses to the 2008 US tax rebate, Borusyak et al. (2024) and Orchard et al. (2025) estimate an MPC of 0.30, correcting earlier higher estimates by Parker et al. (2013) and Broda and Parker (2014). Estimates for the consumption response to the 2020 Economic Stimulus Payment in the US range from 0.08–0.28 (Parker, Schild, Erhard and Johnson, 2022), to 0.25–0.30 (Baker, Farrokhnia, Meyer, Pagel and Yannelis, 2023), or 40% (Coibion et al., 2020). In a randomized experiment, Boehm et al. (2025) observe a one-month MPC of 0.23 in response to an unanticipated 300 Euro transfer. Ganong, Jones, Noel, Farrell, Greig and Wheat (2023) study responses to typical income shocks and find an MPC of 0.21 for nondurable consumption on a monthly basis and 0.29 on a quarterly basis. Fagereng et al. (2021) estimate a within-year MPC of around 0.50 out of lottery winnings. Kotsogiannis and Sakellaris (2025) estimate an MPC of 0.43 out of \$1,000 tax lottery winnings in Greece.

MPCs decline for larger shocks. MPCs decline with larger shock size. This has been observed, *e.g.*, by Kueng (2018), Fagereng et al. (2021), Colarieti et al. (2024), and Ganong et al. (2025a). An exception are Fuster et al. (2021) who find that MPCs increase with shock size, though they also find a negative relationship on the intensive margin.

MPCs are larger for losses. An asymmetry between equally-sized gains and losses has been observed, *e.g.*, by Bunn et al. (2018), Christelis et al. (2019), Fuster et al. (2021), and Colarieti et al. (2024).

Heterogeneity in MPCs. MPCs vary widely in the cross-section of households (see, *e.g.*, Jappelli and Pistaferri, 2014; Lewis et al., 2024; Fuster et al., 2021).

Extreme MPCs of 0 or 1 are common. Identifying extreme MPCs requires identifying MPCs on the household level. Due to the inherent noise in households’ consumption processes, most studies only estimate and report average MPCs or average MPCs in a subgroup of the population. Here, survey-based methods are at an advantage because they can directly

elicit household-level MPCs from each respondent. These studies typically find many households who report an MPC of either 0 or 1, *e.g.*, [Drescher et al. \(2020\)](#) who use HFCS data from 17 European countries, [Andreou, Demetriadou and Tryphonides \(2024\)](#) who work with the NielsenIQ Consumer Panel 2008 tax rebate survey in the US, [Coibion et al. \(2020\)](#) who study consumer responses to the 2020 Economic Stimulus Payment in the US, or [Jappelli and Pistaferri \(2020\)](#) who use survey data from Italy. An exception among survey work is [Fuster et al. \(2021\)](#) who observe a large share of MPCs of 0 (74% for a \$500 gain) but few MPCs close to 1.

Observational studies do not always find spikes at MPCs of 0 or 1 in estimated MPC distributions. However, as we argued in Section 2 and we demonstrate in Appendix A.2, observational methods are not well-suited to identifying the MPC distribution, and so we should not expect observational methods to detect such spikes (see Figure A.1). To re-emphasize the point, it is well known in econometrics that even perfect experiments do not identify the treatment effect distribution absent very strong independence assumptions that are not guaranteed by the mere fact that treatments are randomized (see *e.g.*, [Abbring and Heckman, 2007](#), for a thorough discussion). Intuitively, in the consumption-savings context, one cannot detect if a household had a high MPC out of the randomized transfer or merely experienced an income or expenditure shock. We formalize this point in a simulation in Appendix A.2.

C.5 Discussion of Response Noise

Can noise falsely generate quick-fixing in our data? To find out, we test whether quick-fixing can arise from chance. We focus on the prediction that quick-fixing households should transition to an interior MPC at most once and remain interior thereafter, while unclassified households should never report extreme MPCs.

In our data, 52% of households behave fully consistently with this hypothesis, 71% deviate from the one-switch pattern at most once, and 83% deviate at most twice.

In a randomization test, we derive the distribution of consistency levels under the null hypothesis that there is no link between MPCs, shock sizes, and valence. Specifically, we break any systematic quick-fixing transition in our data by reshuffling MPCs within households (separately for gains and losses, 10,000 permutations). On average, we find much lower consistency rates (29% full consistency, 39% deviate at most once, 53% deviate at most twice). Moreover, these low consistency levels vary little across permuted datasets due to the large sample size. Hence, the frequency of quick-fixing observed in our data is extremely unlikely to arise from chance.

Can noise conceal quick-fixing in our data? To find out, we test whether noise can obscure true quick-fixing in a simulation. We assume that households are truly quick-fixing but have “trembling hands” and erroneously report a wrong extensive margin response with probability $\alpha=0.1$. We set the share of each type in our simulation equal to the estimated type shares in our data and the distribution of switching points equal to the actually observed distribution of switching points in our data. This simulation, based on a modest amount of response noise, yields consistency rates that are comparable to our data: 67% deviate at most once; though only 29% behave fully consistently, much less than in our data, and 90% deviate at most twice, a bit more than in our data. The key takeaway of this analysis is that even modest levels of response error can explain inconsistencies in households’ responses across our 14 different MPC questions.

Test-retest study. We conduct an empirical “test-retest” study to gauge the likely response error directly. For 138 households recruited on Prolific, we elicit our full MPC module twice—a few hours apart. Unlike for other data collections reported in the paper, we missed the chance to preregister the test-retest study, though its objective should be clear nonetheless. Households report an identical extensive margin response to identical shocks in 82% of cases (random benchmark: 39%), and we assign the same household to the same type in 69% of cases (random benchmark: 24%). The random benchmark indicates which consistency level could be expected if we had randomly matched wave 1 and wave 2 responses.

Averaged across the three extensive margin responses (MPC of 0 or 1 or in interior), weighted by their respective frequency, we find a test-retest correlation of 0.71. Averaged across the five types, weighted by their respective frequency, we find a test-retest correlation of 0.61. Hence, the test-retest correlations are close to the test-retest correlations ranging from 0.71–0.86 reported for economic references such as risk or time preferences in [Falk, Becker, Dohmen, Huffman and Sunde \(2023\)](#).

C.6 Additional Qualitative Evidence

Sample and Design. How do households explain their extreme MPCs for small shocks and their transition to interior MPCs for large shocks? We survey 502 additional US households and ask them for their consumption-savings responses to \$100 and \$1,000 income shocks. We recruit households with the survey company Prolific. The demographic characteristics of the sample are summarized in [Table C.2](#).

57% of households adopt an extreme MPC for the \$100 shocks, but only 23% do so for the \$1,000 shocks. We ask respondents who switch from an extreme MPC to an interior MPC to explain why they do so. This qualitative approach complements our quantitative

evidence and sheds light on why households prefer extreme MPCs for small shocks.

For example, households who report an MPC of 1 for a small income gain of \$100 but an MPC below 1 for an income gain of \$1,000 are asked:

[Q1] You responded that your household would not increase its saving in response to a \$100 one-time payment. You would spend everything.

Please explain why your household would spend everything and would not increase its saving.

[Q2] However, you responded that your household would increase its saving in response to a \$1,000 one-time payment.

Please explain why your household would respond differently in these two situations.

We ask analogous questions for households who report an MPC of 0 for the small income gain but an MPC above 0 for a large income gain. We also ask these questions for losses.

Results. We manually identify common themes in households' responses, develop a coding scheme, and assign each response to the themes it contains. Table C.3 provides an overview of the resulting coding scheme. We discuss the results below.

We focus on gains first. For gains, the coded text data reveal that almost all households (86%) explicitly refer to the contrast in shock size (\$100 versus \$1,000) when explaining their extreme MPC for the small shock or their transition from extreme MPCs to an interior MPC.

Many households view small shocks as insignificant and conveniently addressed with an extreme MPC of 0 or 1. The following respondent expresses it directly:

“One hundred bucks is not that much. It’s great, don’t get me wrong, but it’s something you either spend on a dinner or put away. Where we’re at right now, it’s going right in the bank.” (*MPC=0 for \$100, MPC=0.2 for \$1,000*)

By contrast, the large shock of \$1,000 is often described as a significant change to their household finances, and households realize that this requires a more balanced approach.

“Since the amount of \$1000 is fairly significant, and we are increasing our savings by a good amount, I think taking \$100 dollars out and saving the other \$900 is fair and feels rewarding from both a long-term and short-term perspective.” (*MPC=0 for \$100, MPC=0.1 for \$1,000*)

Why do households adopt extreme MPCs for small shocks? The reasons for this behavior can be multifaceted. For example, some households (16%) refer to **habits** or rules such as fixed spending budgets (leading to an MPC of 0) or saving targets (leading to an MPC of 1) that they do not want to overturn for small shocks.

“I have a budget for a reason and generally stick to it unless there are major changes.”
($MPC=0$ for \$100, $MPC=0.2$ for \$1,000)

“\$100 is not such a big amount that it will make me change my spending habits.”
($MPC=0$ for \$100, $MPC=0.1$ for \$1,000)

“My wife and I already contribute regularly to our savings.” ($MPC=1$ for \$100, $MPC=0.25$ for \$1,000)

Many households recognize the unexpected income as a welcome opportunity to **treat themselves**. Some households are ready to immediately spend the \$100 for themselves or their families, but they view it as “irresponsible” to not save a good part of the larger windfall.

“\$100 is not all that much when it comes down to it. It will cover one or maybe 2 utility bills. Why not just use the unexpected \$100 to spend on something you can enjoy or something that can help you in the short-term?” ($MPC=1$ for \$100, $MPC=0.75$ for \$1,000)

“The \$1,000 is a larger amount so I would be overindulging if I did spend it all and increased my spending instead of saving. I could have done a percentage and saved the \$100 but I felt like \$100 was an appropriate gift for myself. When the dollar amounts get much larger, the impact is much bigger if I don’t save anything.” ($MPC=1$ for \$100, $MPC=0.2$ for \$1,000)

Other households instead choose to maintain their **household discipline**. They seek to avoid “frivolous consumption” for small shocks, but, in case of a larger payment, they “feel comfortable” to treat themselves and spend a part of the larger income shock.

“\$100 is not a lot of money and since it came unexpectedly, I would put it in savings. That way I could use it later. I think if I spent the money now, it would be spent frivolously.” ($MPC=0$ for \$100, $MPC=0.5$ for \$1,000)

“An extra \$1000 feels like it is a lot more extra than an extra \$100. While I would still want to save the majority of it, it feels more comfortable to be able to use some of the larger sum of money for extra spending right now versus saving it.” ($MPC=0$ for \$100, $MPC=0.2$ for \$1,000)

We code 41% of households as talking about the desire to treat themselves and 17% as referring to household discipline. As illustrated above, the two arguments often occur together.

Another prominent argument that 25% of households express is that the \$100 would **not have any meaningful impact** if split between spending and saving. To avoid such a “drop in the bucket”, they choose an extreme, one-sided response. For example, the following household cannot think of a meaningful way to spend a small amount of money and hence opts to save the entire amount ($MPC = 0$).

“Our bills are mostly covered and we do not have significant debt. This amount of money is not really large enough to make an impact on our spending. It would be put into our savings as we typically save extra money.” ($MPC=0$ for \$100, $MPC=0.3$ for \$1,000)

Other households make the opposite case, arguing that it “would not make a dent” in their savings if they save part of the \$100 ($MPC = 1$), hence preferring to spend everything.

“It [\$100] is not enough money to make a real dent in any debt payments. We would use this money like a “treat” to go out to dinner or the movies.” ($MPC=1$ for \$100, $MPC=0.4$ for \$1,000)

“The \$100 is not really enough to move the needle in saving. It is a very small amount that spending it would actually provide more joy and benefit from it then it would to save.” ($MPC=1$ for \$100, $MPC=0.1$ for \$1,000)

The results are similar for losses where households who absorb the small \$100 loss with their savings do not view the loss as substantial enough to disrupt their regular spending habits. This strategy becomes infeasible or undesirable for the larger \$1,000 loss. Households who absorb the \$100 loss with their spending provide a mirror image. They do not want to disrupt their savings routines and find it easy to cut back on discretionary expenses like dining out, entertainment, or non-essential purchases. But for the \$1,000 loss, households want or need to draw on their savings or even increase debt (through loans or credit) to manage this larger loss without cutting important expenditures.

Multiple factors appear to make extreme MPCs convenient solutions. First, households refer to habits and routines, *e.g.*, a fixed spending budget, a fixed monthly transfer to savings, or the goal to maximize savings, and deviating from such default rules could come at a cost. Second, for small shocks, extreme MPCs appear to be easier to imagine, evaluate, and appreciate. By contrast, interior MPCs lead to two small, seemingly imperceptible changes that are not perceived to “make a dent” in households’ savings or spending. Third, many households recognize an income gain as a welcome opportunity to treat themselves or

Table C.3: Overview of the coding scheme

Theme (and detected freq.)	Description
Gains	
<i>199 cases where respondents choose an extreme MPC for a small shock but not for a large shock.</i>	
Shock size (86%)	Respondent mentions the difference in the shock sizes, e.g. contrasts the two shocks or says that \$100 is little or \$1000 a lot.
Habit (16%)	Respondent mentions that they generally try to save/spend in situations with small income gains.
Does not make a difference (25%)	Respondent mentions that spending/saving the money would not make a meaningful difference to their spending or savings.
Household discipline (17%)	• MPC of 0 for \$100: Only in case of a larger amount, respondent feels comfortable to spend part of the amount, but they avoid “frivolous” spending for the small amount. • MPC of 1 for \$100: Respondent is fine with spending the small amount, but they argue it would be “irresponsible” to fully spend the larger amount.
Treat oneself (41%)	• MPC of 0 for \$100: Only in case of a larger amount, respondent wants to use a part to treat themselves. • MPC of 1 for \$100: Respondent wants to use the \$100 to treat themselves.
Need (15%)	• MPC of 0 for \$100: Respondent argues that they do not need additional purchases. • MPC of 1 for \$100: Respondent immediately needs the money for essential purchases.
Lumpy consumption plans (6%)	Respondent has a specific spending plan or need, but \$100 is not yet enough to realize it.
Losses	
<i>184 cases where respondents choose an extreme MPC for small shock but not for large shock.</i>	
Shock size (84%)	Respondent mentions the difference in the shock sizes, e.g. contrasts the two shocks or says that \$100 is little or \$1000 a lot.
Habit (13%)	Respondent mentions that they generally try to cut saving/spending in situations with small income losses.
Buffer (49%)	• MPC of 0 for \$100: Respondent can easily draw on a buffer of savings. • MPC of 1 for \$100: Respondent can easily cut discretionary, non-essential consumption.
Balance required (34%)	Interior MPC for large loss because respondents do not want to or simply cannot afford to reduce their spending/savings by the full \$1000.
Budget already tight (8%)	Respondent reports having such a tight spending budget they prefer to not reduce spending any further in response to a \$100 loss.

their families. Most balance consumption and saving for the large shock, but they approach the smaller \$100 gain differently. Some conclude that they should “indulge” and spend everything, while others choose to maintain “discipline” and save everything. Of course, it seems plausible that further psychological forces are at work, which are harder for households to explicitly articulate. For example, finding a compromise between consumption and saving could require more computational effort.

Our model of quick-fixing captures the convenience of extreme MPCs for small shocks and the transition pattern from extreme to interior MPCs, thus providing a plausible representation of households’ introspection.

References

- Abbring, Jaap H and James J Heckman**, “Econometric evaluation of social programs, part III: Distributional treatment effects, dynamic treatment effects, dynamic discrete choice, and general equilibrium policy evaluation,” *Handbook of econometrics*, 2007, 6, 5145–5303.
- Akerlof, George A and Janet L Yellen**, “A near-rational model of the business cycle, with wage and price inertia,” *The Quarterly Journal of Economics*, 1985, 100 (Supplement), 823–838.
- Andreou, Elena, Maria Demetriadou, and Andreas Tryphonides**, “The Cross Section of Household Preferences: MPCs and the HtM Status,” *Working Paper*, 2024.
- Baker, Scott R., Robert A Farrokhnia, Steffen Meyer, Michaela Pagel, and Constantine Yannelis**, “Income, Liquidity, and the Consumption Response to the 2020 Economic Stimulus Payments,” *Review of Finance*, 2023, 27 (6), 2271–2304.
- Boehm, Johannes, Etienne Fize, and Xavier Jaravel**, “Five Facts about MPCs: Evidence from a Randomized Experiment,” *American Economic Review*, 2025, 115 (1), 1–42.
- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess**, “Revisiting Event-Study Designs: Robust and Efficient Estimation,” *The Review of Economic Studies*, 2024, 91 (6), 3253–3285.
- Broda, Christian and Jonathan A. Parker**, “The Economic Stimulus Payments of 2008 and the aggregate demand for consumption,” *Journal of Monetary Economics*, 2014, 68, S20–S36.
- Bunn, Philip, Jeanne Le Roux, Kate Reinold, and Paolo Surico**, “The consumption response to positive and negative income shocks,” *Journal of Monetary Economics*, 2018, 96, 1–15.
- Christelis, Dimitris, Dimitris Georgarakos, Tullio Jappelli, Luigi Pistaferri, and Maarten Van Rooij**, “Asymmetric Consumption Effects of Transitory Income Shocks,” *The Economic Journal*, 2019, 129 (622), 2322–2341.
- Coibion, Olivier, Yuriy Gorodnichenko, and Michael Weber**, “How Did U.S. Consumers Use Their Stimulus Payments?,” Technical Report 27693, NBER 2020.
- Colarieti, Roberto, Pierfrancesco Mei, and Stefanie Stantcheva**, “The How and Why of Household Reactions to Income Shocks,” *Working Paper*, 2024.
- Drescher, Katharina, Pirmin Fessler, and Peter Lindner**, “Helicopter money in Europe: New evidence on the marginal propensity to consume across European households,” *Economics Letters*, 2020, 195, 109416.
- Fagereng, Andreas, Martin B. Holm, and Gisle J. Natvik**, “MPC Heterogeneity and Household Balance Sheets,” *American Economic Journal: Macroeconomics*, 2021, 13 (4), 1–54.

- Falk, Armin, Anke Becker, Thomas Dohmen, David Huffman, and Uwe Sunde**, “The Preference Survey Module: A Validated Instrument for Measuring Risk, Time, and Social Preferences,” *Management Science*, 2023, 69 (4), 1935–1950.
- Fuster, Andreas, Greg Kaplan, and Basit Zafar**, “What Would You Do with \$500? Spending Responses to Gains, Losses, News and Loans,” *Review of Economic Studies*, 2021, 88 (4), 1760–1795.
- Ganong, Peter, Damon Jones, Pascal Noel, and Daniel Sullivan**, “The Marginal Propensity to Consume out of Stimulus: Evidence From Administrative Bank Data,” *Working Paper*, 2025.
- , —, —, **Diana Farrell, Fiona Greig, and Chris Wheat**, “Wealth, Race, and Consumption Smoothing of Typical Labor Income Shocks,” *Working Paper*, 2023.
- , **Pascal Noel, Christina Patterson, Joseph Vavra, and Alexander Weinberg**, “Earnings Instability,” *Working Paper*, 2025.
- Heckman, James J, Jeffrey Smith, and Nancy Clements**, “Making the most out of programme evaluations and social experiments: Accounting for heterogeneity in programme impacts,” *The Review of Economic Studies*, 1997, 64 (4), 487–535.
- Jappelli, Tullio and Luigi Pistaferri**, “Fiscal Policy and MPC Heterogeneity,” *American Economic Journal: Macroeconomics*, October 2014, 6 (4), 107–136.
- and —, “Reported MPC and Unobserved Heterogeneity,” *American Economic Journal: Economic Policy*, 2020, 12 (4), 275–297.
- Kaplan, Greg, Giovanni L Violante, and Justin Weidner**, “The wealthy hand-to-mouth,” *Brookings Papers on Economic Activity*, 2014, pp. 77–153.
- Kotsogiannis, Christos and Plutarchos Sakellaris**, “MPCs Estimated from Tax Lottery, Survey and Administrative Data,” *Working Paper*, 2025.
- Kueng, Lorenz**, “Excess Sensitivity of High-Income Consumers,” *The Quarterly Journal of Economics*, 2018, 133 (4), 1693–1751.
- Laibson, David, Peter Maxted, and Benjamin Moll**, “A Simple Mapping from MPCs to MPXs,” Technical Report, Harvard University 2022.
- Lewis, Daniel J., Davide Melcangi, and Laura Pilossoph**, “Latent Heterogeneity in the Marginal Propensity to Consume,” Working Paper 32523, NBER 2024.
- Misra, Kanishka and Paolo Surico**, “Consumption, Income Changes, and Heterogeneity: Evidence from Two Fiscal Stimulus Programs,” *American Economic Journal: Macroeconomics*, 2014, 6 (4), 84–106.
- Morduch, Jonathan and Rachel Schneider**, *The Financial Diaries: How American Families Cope in a World of Uncertainty*, Princeton University Press, 2017.
- Orchard, Jacob D, Valerie A Ramey, and Johannes F Wieland**, “Micro MPCs and Macro Counterfactuals: The Case of the 2008 Rebates,” *The Quarterly Journal of Economics*, 2025, 140 (3), 2001–2052.
- Parker, Jonathan A, Nicholas S Souleles, David S Johnson, and Robert McClelland**, “Consumer Spending and the Economic Stimulus Payments of 2008,” *American Economic Review*, 2013, 103 (6), 2530–2553.
- Parker, Jonathan, Jake Schild, Laura Erhard, and David Johnson**, “Household Spending Responses to the Economic Impact Payments of 2020: Evidence from the Consumer Expenditure Survey,” *Working Paper*, 2022.