Quick-Fixing: Near-Rationality in Consumption and Savings Behavior

Peter Andre* SAFE

Joel P. Flynn† Yale

George Nikolakoudis‡ Princeton

Karthik A. Sastry§ Princeton

February 3, 2025

Abstract

When optimizing consumption-savings decisions is costly, people may instead rely on *quick-fixes*, simple policy functions that avoid these costs. We introduce a model of quick-fixing. To study it empirically, we field a novel survey that measures households' consumption policy functions in response to income shocks. Almost 70% of households follow one of four simple quick-fixes that fully consume or fully save out of small shocks, but they abruptly adjust their behavior for large shocks. This behavior accounts for almost half of the cross-sectional variance in marginal propensities to consume, but is poorly predicted by other demographic and economic information. In an incomplete-markets model calibrated to match our evidence, we find that quick-fixing is *near-rational*: the average opportunity cost of quick-fixing is only \$17 per quarter. Yet, this small, empirically realistic deviation from the rational model significantly alters aggregate consumption responses to income shocks of varying sizes.

*SAFE Leibniz Institute for Financial Research and Goethe University Frankfurt, Theodor-W.-Adorno-Platz 3, 60323, Frankfurt am Main, Germany. Email: andre@safe-frankfurt.de.

†Yale University Department of Economics, 30 Hillhouse Avenue, New Haven, CT, USA, 06511. Email: joel.flynn@yale.edu.

‡Princeton University Department of Economics, Julis Romo Rabinowitz Building, Princeton, NJ, USA, 08544. Email: nikolakoudis@princeton.edu.

§Princeton University Department of Economics, Julis Romo Rabinowitz Building, Princeton, NJ, USA, 08544. Email: ksastry@princeton.edu.

Acknowledgments: We are grateful to Kai Barron, Jonathon Hazell, David Laibson, Chen Lian, Lukas Nord, Ryan Oprea, Michael Peters, Oliver Pfäuti, Charlie Rafkin, Chris Roth, Tony Smith, Matthew White, Johannes Wohlfart, and seminar participants at the New York Fed, the University of Connecticut, the LAEF Workshop, CREI, Oxford, Cambridge, the Philadelphia Fed, Princeton, Columbia, and the 2024 Winter SED Meeting. We thank Zhuokai Huang for excellent research assistance. First posted version: October 25, 2024. Ethics approval: The study obtained ethics approval from the Princeton University Institutional Review Board (#15968). Preregistration: The data collections were preregistered via OSF (https://osf.io/sgw9y/registrations).

1 Introduction

A long-standing hypothesis in economics is that decision-makers may behave suboptimally because the gains from making better decisions are small and not worth the required effort. Akerlof and Yellen (1985b) define such *near-rational behavior* as "behavior that is perhaps sub-optimal but that nevertheless imposes very small individual losses on its practitioners relative to the consequences of their first-best policy."

Household consumption and savings behavior is an important and natural candidate for near-rationality. It is important because consumption choices and corresponding marginal propensities to consume (MPCs) are critical determinants of aggregate demand in macroeconomic models. At the same time, it is a natural candidate for near-rationality because always optimizing consumption and savings plans to keep up with frequent changes in economic circumstances may be both practically challenging and of limited benefit. For example, Cochrane (1989) shows that, if a representative household simply sets its consumption equal to aggregate income in the US, then its welfare losses would be less than $1 per quarter.

While these arguments may be intuitively appealing, we lack both theory and evidence to understand the precise nature and implications of near-rational household behavior. Do households exhibit near-rationality in consumption and savings choices? If so, which forms does near-rationality take? And what are its macroeconomic consequences?

This paper formally studies near-rational consumption and savings behavior in three steps. First, we develop a simple theoretical framework. Instead of costly optimization, households can also rely on quick-fixes, simple policy functions that avoid the cost of optimization. Second, we design a novel survey to elicit the distribution of households' consumption policy functions and evaluate this theory. We find that, consistent with our theory, most households rely on one of four quick-fixes to respond to small income shocks but not large income shocks. Finally, we combine our empirical estimates with a quantitative model to study the macroeconomic implications of this behavior. Even with economically small optimization costs, quick-fixing significantly alters the response of consumption to income shocks. Taken together, our results imply that a small, empirically disciplined deviation from the rational model can have large consequences for shock propagation and policy design.

A Simple Model. We first formulate the near-rationality hypothesis in a textbook two-period model of consumption-savings decisions. Households have simple, potentially sub-optimal consumption policy functions that we call *quick-fixes*. For example, a household's quick-fix might be to keep either consumption or savings at a reference level. Households can depart from this behavior and take the optimal action at a fixed utility cost, which captures the effort required to contemplate their circumstances, deliberate about what to do, or implement their choice. We show that even small costs can sustain large deviations from rational behavior: for example, a household with logarithmic preferences would tolerate a 5% deviation from the rational consumption level if the cost of optimization is as little as 0.25% of consumption.

The model makes sharp predictions for how households respond to income shocks. In particular, households rely on quick-fixes in response to small shocks but abruptly switch to the rational behavior in response to large shocks. Moreover, conditional on having optimized, households behave the same way regardless of the nature of their quick-fix. Quick-fixing is hence crucial to understanding households' MPCs and consumption-savings behavior, but its precise implications clearly depend on which quick-fixes (if any) households use.

Our model implies that the distribution of household-level *consumption policy functions*, *i.e.,* consumption responses to income shocks of different sizes, is the required moment to quantify quick-fixing behavior and evaluate the model against alternatives. To measure this object, observational data are not suitable. Even absent measurement error, these data can reveal how a household responds to only a single shock at a given time. This necessitates a survey-based research design that elicits responses to many shocks of different sizes.

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 is balanced to approximate 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. Relative to existing work collecting survey measures of the marginal propensity to consume (see *e.g.,* Parker and Souleles, 2019; Jappelli and Pistaferri, 2020; Fuster, Kaplan and Zafar, 2021; Colarieti, Mei and Stantcheva, 2024), our survey collects detailed *within-respondent* information on consumption and savings responses to a wide range of shocks.

We document five new facts. Figure 1 visualizes the first fact, which we refer to as

Figure 1: The "bowtie" MPC distribution across shock sizes

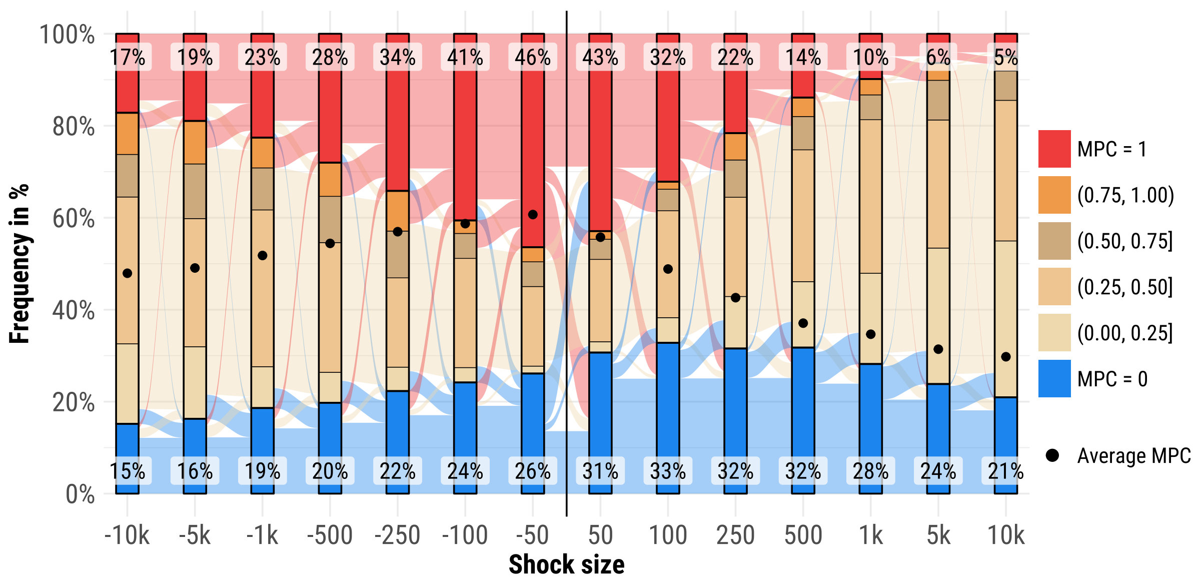

Notes: The alluvial graph summarizes the MPC data of the 4,981 US households in our sample. Each of the 14 columns displays the distribution of MPCs for one particular shock size, with colors indicating the size of the MPC. The streams between bars indicate how households' MPCs transition between two 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.

the "bowtie" shape of the joint distribution of MPCs and 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. This fact 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.

Our second fact is that the distribution of MPCs is well described by a decomposition of households into four *quick-fixing types* who use extreme MPCs of 0 or 1 as quick-fixes for small shocks but abandon them for larger shocks. We only use information from how households respond to the smallest shocks, the $50 loss and gain, to define these types. 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 and cut savings after losses; and *savings prioritizers* if they increase savings after windfalls and cut consumption after losses. These quick-fixing types span 68% of our respondents. Using within-respondent data, we show that quick-fixers (i) adopt categorically different policies for small shocks, (ii) abruptly transition to interior MPCs for large shocks, and (iii) have similar MPCs across quick-fixing types conditional on transitioning to an interior MPC. This validates the three core predictions of the near-rational model. 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 and fourth facts suggest 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 11% of the household-level variation in MPCs. This result, consistent with previous findings in the literature (Lewis, Melcangi and Pilossoph, 2024; Fuster et al., 2021), poses a challenge for standard models in which such variables should explain *all* MPC variation. By contrast, quick-fixing types account for 49% of the household-level variation in MPCs (Fact 3). 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% (Fact 4). We also find that quick-fixing accounts for a large part of the variation in MPCs across small and large shocks as well as gains and losses. Due to the presence of quick-fixes with an MPC of 1, quick-fixing also explains large average MPCs.

Our fifth fact supports the proposed mechanism that households quick-fix to avoid costly deliberation. 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 than when they have interior MPCs.

Quantitative Model. In the final part of the paper, we gauge the macroeconomic implications of this empirically prevalent form of near-rationality for households. To do so, we integrate quick-fixing into a standard incomplete-markets household problem with income risk and borrowing constraints. Motivated by our observation that household characteristics are poor predictors of quick-fixing types, we assume that the population consists of discrete types who adopt each of the four quick-fixes and a fifth group that is rational (a conservative approach to model the uncategorized households). We calibrate the model to match standard external targets for preferences and the stochastic earnings process. From the survey, we match the measured frequency of each quick-fixing type and households' propensities to quick-fix. Here, we highlight three sets of results.

First, we find that small costs of optimization are consistent with the large and macroeconomically important differences in household behavior generated by quick-fixing. The one-time cost of optimization ranges between \$1.50 to \$175 depending on type, and the lifetime loss from near-rationality ranges between \$0.50 and \$75 per quarter. This explains why quick-fixing may persist in spite of its suboptimality: because households lose so little, the scope for learning or "selection pressure" to instill rationality is naturally limited.

Second, we find that different types of quick-fixers as well as rational agents have almost identical long-run wealth distributions despite their markedly different behavior in response to income shocks. This represents a further test of the model: the model matches the non-targeted empirical moment that wealth is a poor predictor of quick-fixing types. This occurs because households reoptimize in response to larger income shocks (*i.e.*, those that likely capture promotions, firings, and life-cycle events) and any "mistakes" along the way wash out in the long run. Moreover, our model generates considerable variation in MPCs that is unexplained by assets and income. This variation is present in the data but absent by construction in the nested rational incomplete-markets model. Thus, quick-fixing helps explain why we often see widely different MPCs among households with identical financial observables (e.g., Lewis et al., 2024; Fuster et al., 2021).

Third, the model generates considerably more *size-dependent* responses to government transfer shocks than the standard incomplete-markets model with rational consumers. In our model, the aggregate quarterly MPC out of a \$100 transfer is 0.24 higher than that out of a \$1,200 transfer, which is almost three times larger than the difference of 0.09 in the nested rational model. We moreover show that the corresponding dynamic responses or *intertemporal marginal propensities to consume* (Auclert, Rognlie and Straub, 2024) are more front-loaded for small shocks and back-loaded for large shocks. The key mechanism is the high incidence of quick-fixers who transition from a unitary to an intermediate MPC as the shock becomes larger. This implies a delicate trade-off for policymakers who want to use stimulus checks to increase aggregate consumption: small checks will yield a smaller aggregate response but feature much greater and more immediate bang for their buck.

By demonstrating the importance of quick-fixing for iMPCs, our results have direct implications for shock transmission in the large class of heterogeneous-agent models in which iMPCs are sufficient statistics for the general-equilibrium macroeconomic response to shocks (Auclert et al., 2024; Angeletos, Lian and Wolf, 2024). The strong size-dependence of iMPCs highlights a quantitatively important manifestation of the Lucas critique in the use of MPCs in quantitative work: caution should be applied to direct extrapolation of empirical estimates of MPCs that average over shocks of different sizes when the distribution of shocks under a chosen counterfactual is markedly different.

Related Literature. Within the literature initiated by Akerlof and Yellen (1985a,b) that studies the near-rationality hypothesis, Cochrane (1989) and Krusell and Smith (1996) evaluate the losses from various rule-of-thumb consumption rules and find that they are often very small. Kueng (2018) finds that households exhibit high MPCs out of predictable dividend payments, particularly when these payments are small relative to their permanent income, and argues that such behavior aligns with near-rationality. Building on the theoretical observation that near-rational behavior can have quantitatively small opportunity costs, we use survey data to characterize the near-rational model, contrast it with leading alternatives, empirically discipline the rules that people use, and quantify the macroeconomic implications of the empirically implied near-rationality of households.

The most related papers that use surveys to understand heterogeneity in the marginal propensity to consume are Fuster et al. (2021) and Colarieti et al. (2024). Fuster et al. (2021) identify one similar regularity in their data, the mass of agents with MPC = 0, and rationalize this behavior via a fixed "menu cost" of adjusting consumption. Colarieti et al. (2024) use qualitative survey questions and a clustering approach to shed light on the motives and considerations behind households' responses to large income shocks. In contrast to both, we investigate the near-rational model and, to do so, elicit multiple MPCs for each household, varying the sign and magnitude of the shock over a large set of scenarios. Our analysis builds on Fuster et al. (2021) by using more fine-grained variation in shock sizes within respondents. This is what allows us to uncover richer variation in households' policy functions beyond fixing consumption.

Conceptually, the behavior of households in our near-rational model is different from behavior with "menu costs" of adjusting behavior. In menu-cost models, agents cannot change their behavior without paying an adjustment cost. By contrast, in our model, households do change their behavior when they quick-fix—they simply do so sub-optimally and heterogeneously. That is, the cost is an optimization cost of changing how one's actions respond to shocks and not a menu cost of changing one's action itself. In this way, our modeling approach is consistent with the broader notion of costly planning pursued by Reis (2006a,b). It also aligns with foundational research in behavioral economics, which demonstrates that individuals often resort to simple solutions to conserve cognitive resources when faced with complex decision problems (Bordalo, Conlon, Gennaioli, Kwon and Shleifer, 2024; Enke, 2024; Oprea, 2024), and recent theoretical work that explicitly models dual processing in economic decision-making (Ilut and Valchev, 2023, 2024).

Our empirical results are consistent with the near-rational model. By contrast, we argue that no other consumption-savings model of which we are aware is consistent with our empirical findings (see Section 3.7 for more details). These include models with: (i) incomplete markets (Bewley, 1979; Deaton, 1991; Carroll, 1997), (ii) incomplete markets with *ex ante* heterogeneity in preferences (Carroll, Slacalek, Tokuoka and White, 2017; Aguiar, Bils and Boar, 2024) or investment opportunities (Kaplan and Violante, 2022), (iii) incomplete markets with multiple accounts (Kaplan and Violante, 2014; Kaplan, Moll and Violante, 2018), (iv) durables or consumption commitments (Barsky, House and Kimball, 2007; Chetty and Szeidl, 2016), (v) consumption mistakes arising from present bias (Laibson, Maxted and Moll, 2021), temptation (Krusell, Kuruşçu and Smith, 2002), incomplete information or rational inattention (Sims, 2003), finite planning horizons (Boutros, 2022), sparsity (Gabaix, 2014), or misperceptions in wealth (Lian, 2023), or (vi) infrequent optimization (Gabaix and Laibson, 2001; Fuster et al., 2021).

Outline. Section 2 introduces a simple model to formalize the near-rationality hypothesis. Section 3 presents the survey and empirical results. Section 4 presents the quantitative model and calibration. Section 5 presents the quantitative results. Section 6 concludes.

2 A Simple Model of Near-Rationality

We begin our analysis by formalizing quick-fixing and near-rationality in a textbook two-period model of consumption-savings choice. We characterize near-rational behavior and isolate the model's prediction that households use simple, approximate rules to respond to small income shocks, but optimize in response to large shocks. This observation motivates our strategy of empirically measuring consumption policy functions.

2.1 The Model

The Environment. A continuum of households is indexed by $i \in [0,1]$ and lives for two periods $t \in \{1,2\}$. The household is born at period $t = 1$ with assets a_i and receives income $y_{i,t}$ in period t. We collect these household-level state variables as $z_i = (a_i, y_{i,1}, y_{i,2})$. The household can save in a risk-free bond that pays a gross interest rate of $R > 0$. The household has preferences over streams of consumption $(c_{i,1}, c_{i,2}) \in \mathbb{R}^2_+$ that are given by:

$$u(c_{i,1}) + \beta u(c_{i,2}) \tag{1}$$

where $\beta > 0$ and $u : \mathbb{R}_+ \to \mathbb{R}$ is a twice continuously differentiable and strictly concave function. We assume that $\beta R = 1$. Here, we use these simplifying assumptions to clarify the key ideas. In the quantitative model of Section 4, we will consider an environment with an infinite horizon, idiosyncratic consumption risk, and borrowing constraints.

Rational Behavior. We first describe *rational behavior*, which is what the household would optimally do in the absence of any optimization frictions. In this case, each household's optimal consumption follows the Euler equation:

$$u'(c_{i,1}) = \beta R u'(c_{i,2}) \tag{2}$$

Substituting the household's budget constraints into the Euler equation, we have that the household obeys the permanent income hypothesis and fully smooths consumption:

$$c_{i,t} = c^*(z_i) = \frac{R(a_i + y_{i,1}) + y_{i,2}}{1 + R} \tag{3}$$

Near-Rational Quick-Fixing. We now describe *quick-fixing*. Each quick-fixing household has a type $q_i \in \mathcal{Q}$ describing their quick-fix c^{q_i}, a consumption policy function that maps asset-income states to consumption choices and avoids any costs associated with optimization. If a household quick-fixes, its consumption in period 1 is simply

$$c_{i,1} = c^{q_i}(z_i) \tag{4}$$

Consumption in period 2 is the household's residual wealth.[1] For example, households could stick to a fixed consumption budget $\bar{c}$ and have a quick-fixing policy of $c^{q_i}(z_i) = \bar{c}$

[1] As a technical matter, if this implies that $c_{i,2} < 0$, then we set the household's payoff equal to $-\infty$. This ensures that the household will only quick-fix if it is feasible to do so.

for all z_i. Likewise, households could also stick to fixed savings target $\bar{s}$ and have a quick-fixing policy of $c^{q_i}(z_i) = y_{i,1} - \bar{s}$ for all z_i. Which exact quick-fixes households follow is an empirical question. We leave this question open for now, until we turn to the empirical evidence in Section 3. However, we restrict our attention to quick-fixes that have the following natural property: as shocks to income and/or assets become larger, the mistake from quick-fixing becomes larger. Mathematically, this is the statement that the consumption mistake function $|c^{q_i} - c^*|$ is quasi-convex. This assumption is satisfied for fixing either consumption or savings at some reference level.

Households rationally choose whether to follow their quick-fix c^{q_i} or instead to pay a utility cost $\kappa_{q_i} > 0$ to freely optimize. The cost flexibly captures households' difficulty in working out how to optimally change their spending or savings decisions and implementing that action. The value function of following the quick-fix is

$$U^{q_i}(z_i) = u(c^{q_i}(z_i)) + \beta u(R(a_i + y_{i,1} - c^{q_i}(z_i)) + y_{i,2}), \tag{5}$$

and the value function of optimizing is

$$U^*(z_i) = u(c^*(z_i)) + \beta u(c^*(z_i)), \tag{6}$$

This means households quick-fix in state z_i if and only if the payoff loss from quick-fixing is less than the optimization cost:

$$\mathcal{L}^{q_i}(z_i) \equiv U^*(z_i) - U^{q_i}(z_i) \leq \kappa_{q_i} \tag{7}$$

Otherwise, they optimize and behave rationally. To the extent that the optimization costs κ_{q_i} are small, we follow Akerlof and Yellen (1985b) and call households' quick-fixing behavior *near-rational* because it yields a similar payoff to rational behavior.

This model describes consumption behavior under the intuitive assumption that optimization is costly, and cheaper, simple alternatives are available. To isolate the effect of quick-fixing, we otherwise stick to conventional assumptions. In particular, we abstract from other possible deviations from rationality. For example, in our model, households are fully aware of the opportunity costs of quick-fixing behavior, but they simply may not find it worthwhile to optimize. We think of this "meta-rationality" as a reduced-form way to model cognition under resource constraints. This kind of assumption is common in economics (e.g., in the work on rational inattention surveyed by Mackowiak, Matejka and

Wiederholt (2018) and noisy cognition surveyed by Woodford (2020) and Enke (2024)). It is also common in cognitive science (e.g., the resource-rational framework of Lieder and Griffiths, 2020). Likewise, we assume that households are able to flawlessly optimize if they pay the optimization cost. That said, quick-fixing is in principle compatible with many other extensions of the standard model and enriches rather than replaces them.

2.2 The Opportunity Cost of Near-Rational Behavior

We first show that even very small optimization costs can lead near-rational households to behave very differently from rational households. To do this, we describe the opportunity cost of quick-fixing in a given state compared to behaving rationally, 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_i when the asset-income state is z_i is:*

$$\mathcal{L}^{q_i}(z_i) = \frac{1}{2}(1+R)|u''(c^*(z_i))|(c^{q_i}(z_i) - c^*(z_i))^2 + O(|c^*(z_i) - c^{q_i}(z_i)|^3) \qquad (8)$$

Proof. See Appendix A.1. □

This simple result formalizes that there is no first-order loss from deviations from rational 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 (1985b) criterion that there is no "first-order" opportunity cost from quick-fixing behavior. While this simple result abstracts from dynamics, income risk, and higher-order terms, we will later find that it is extremely accurate in our calibrated quantitative model with all of those features (see Figure 7).

Because of this basic envelope logic, even the presence of small optimization 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 small optimization 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 sub-optimal quick-fix is approximately equal to m^2, where $m = (c^{q_i}(z_i) - c^*(z_i))/c^*(z_i)$ is the household's

Table 1: Small optimization 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 relative consumption mistake of size m: $\tau(m) = 1 - e^{-m^2}$.

consumption "mistake" expressed as a proportional deviation from the rational level. We now ask: what is the smallest optimization 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? This is easy to calculate as $\tau = 1 - e^{-m^2} \approx m^2$. In Table 1, we report this optimization 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 rational model. $\triangle$

2.3 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 $y_{i,1}(x) = y_{i,1} + x$ 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 $z_i(x) = (a_i, y_{i,1} + x, y_{i,2})$, we obtain that:

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

$$c_i(x) = \begin{cases} c^{q_i}(z_i(x)) & \text{if } |c^{q_i}(z_i(x)) - c^*(z_i(x))| \leq \sqrt{\frac{\kappa_{q_i}}{\frac{1}{2}(1+R)|u''(c^*(z_i(x)))|}}, \\ c^*(z_i(x)) & \text{otherwise.} \end{cases} \tag{9}$$

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 rationally for large shocks.[2] 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 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 an optimization 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. △

2.4 From Theory to Measurement

A critical moment according to our theory is the distribution of *household-level consumption policy functions* that map income shocks of different sizes to each household's consumption response (Equation 9). The quick-fixing model predicts that (i) households quick-fix in response to small shocks, (ii) households abruptly abandon their quick-fix once the shock crosses a certain threshold, and (iii) households behave the same way conditional on optimizing, regardless of differences in the quick-fixes that they use. All of the predictions clearly differ from those of the nested model of frictionlessly following the permanent income hypothesis; later, we will clarify how the predictions differ from those of other models.

[2]Income shocks affect households' decision whether to quick-fix via two mechanisms: the size of the mistake $|c^{q_i}(z_i(x)) - c^*(z_i(x))|$ and the curvature of utility $|u''(c^*(z_i(x)))|$. 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 rationally 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.

At the same time, the model highlights a key obstacle to testing the near-rationality hypothesis—and, more generally, disciplining models of household behavior—using observational data. In these settings, we typically observe a given household respond to one shock of a given size at any given time. It is not possible to construct the entire policy function from this single point without strong structural assumptions, which are tantamount to "assuming the result" about what model is appropriate.

This necessitates the design of a novel survey to elicit households' policy functions and shed light on their potentially near-rational behavior. Fortunately, previous research finds that survey-based MPC measures yield estimates that are similar to those obtained from observational data (Parker and Souleles, 2019; Colarieti et al., 2024). A survey also allows us to overcome a key challenge to the practical application of the near-rational model: determining what quick-fixes households use. Many policy functions may be near-rational, including opposite behaviors like "consuming all of a small shock" and "saving all of a small shock." But potentially only some are prevalent. Identifying these behaviors will of course be central to determining the model's consequences for aggregate behavior and MPCs.

3 Empirical Evidence from a Novel Survey

This section presents the design and results of a novel, large-scale household survey tailored to uncover households' policy functions. Our goal is to detect whether households quick-fix and, if so, uncover what quick-fixes they employ. We find that the majority of households quick-fix by applying simple rules of full consumption or savings responses (MPCs of 0 or 1) to sufficiently small shocks. Quick-fixing behaviors are heterogeneous in the population, poorly explained by financial and demographic characteristics, and highly explanatory of MPC heterogeneity across households and shocks. Moreover, we find that households deliberate less about quick-fix consumption decisions. We argue that these patterns are in line with the quick-fixing model (as described in Section 2) but not with leading alternatives.

3.1 Survey Design and Sample

We collect data from 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.[3] Appendix C.1 presents further details on the sample.

We follow a standard procedure to measure households' marginal propensity to consume out of unexpected, one-time income changes (*e.g.*, Jappelli and Pistaferri, 2014, 2020; Christelis, Georgarakos, Jappelli, Pistaferri and Van Rooij, 2019; Fuster et al., 2021). First, we 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 increase their spending and saving in response to the income shock. They can respond in two numeric open response fields, and we calculate their MPC by dividing their spending response by the income shock.[4]

> How would this one-time extra income cause your household to change its spending and saving over the next three months?
>
> 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 <u>de</u>creases in debt-taking over the next three months)? $___________

Appendix C.2 presents the key instructions. The full instructions are available online at https://osf.io/2s7cf.

The key feature of our survey is that we measure MPCs out of *multiple* shocks for each household, fourteen in total. Specifically, we include gains and losses with a magnitude

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

[4]Households' changes in spending and saving need to add up to the income gain. We 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.

of \$50, \$100, \$250, \$500, \$1,000, \$5,000, and \$10,000. We ask households' responses to *very small* shocks, such as \$50 or \$100, to isolate circumstances under which quick-fixing behavior is more likely to be prevalent.

We randomize both whether respondents first face gains or 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 full profile of their responses in light of choices made later. We observe similar results across the randomly assigned question orders (Appendix Figure B.1).

The survey closes with a series of demographic and economic background questions, including questions on income, income risk, future income expectations, and wealth. All in all, 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 and 24 min (20%-80% quantile range). However, we obtain similar results if we restrict attention to only the first or first five MPCs that each respondent reports or drop the 50% of the sample that read the preparatory instructions most quickly (Appendix Figure B.1).

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 (0.47) 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%]. All contrasts we mention achieve statistical significance well above conventional levels. To enhance readability, we do not discuss statistical significance in the main text and present the detailed significance tests in Appendix Table B.1.

3.2 Extreme MPCs and the "Bowtie" Pattern

We first describe the aggregate properties of the survey data. Figure 1, introduced earlier, plots the full distribution of consumption responses to each shock. Each bar presents the distribution of marginal propensities to consume for one particular shock, calculated as the ratio of households' "change in spending" response to the size of the shock. Black dots display the average MPC for each shock.

The figure reveals five broad patterns that are already familiar in the literature in both observational and survey data. First, the average MPC (0.47) is high. Second, MPCs decline with larger shocks: for example, they fall from 0.56 for a $50 gain to 0.30 for a $10,000 gain. Third, the average MPC is larger for losses (0.54) than for gains (0.40). Fourth, MPCs vary widely across households. And, fifth, a significant fraction of households report an MPC of 0 or 1. Appendix C.4 compares our cross-sectional results to related survey and non-survey work in further detail. Our broad conclusion is that our aggregate results on the MPC distribution *conditional on specific amounts* are in line with previous findings, which further fosters confidence in our approach.

Unique to our data is that we can study MPCs across a wide range of shocks and that we can do so at the household level. Our first key finding is the "bowtie" pattern in the distribution of MPCs across shocks: extreme MPCs of 0 and 1 are common for small shocks but rare for large shocks—a pattern that is reminiscent of quick-fixing. The mass of households with an "interior" MPC strictly between 0 and 1, which consistently increases from small to large shocks, gives Figure 1 its bowtie-like appearance.

To quantify this bowtie pattern, we measure how many households transition from extreme MPCs of 0 and 1 to an interior MPC as the shock size increases in absolute value. Since the questions for different amounts are asked in random order, households would not necessarily perceive these transitions when taking the survey. We visualize transitions as the shaded flows between bars in Figure 1.

For the $50 income gain, 43% of households consume every dollar (MPC = 1), and 31% save every dollar (MPC = 0). Likewise, for the $50 income loss, 46% of households only reduce their consumption (MPC = 1), and 26% only reduce their savings (MPC = 0). Starting from these numbers, from one shock to the next larger shock, an average net share of 7% of households transition from an extreme MPC to an interior MPC. In fact, most households— namely 68% for gains and 67% for losses—switch to an interior MPC at most once and stick to interior MPCs thereafter. Only 14% (for gains) and 16% (for losses) of households deviate from this pattern more than once.[5] Consequently, very few households still consume or save every dollar for the largest shocks.

We summarize this finding below:

> **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.

[5]Response error is inevitable in survey data, and simulation results illustrate that small response errors can rationalize most of these deviations. In the simulation, we assume that all households would want to switch to an interior MPC at most once and stick to them thereafter. But households have "trembling hands" and deviate from their desired extensive margin response with an i.i.d. 10% error rate. In this case, 11% would deviate from the pattern more than once.

3.3 Household Behavior is Consistent with Near-Rationality

The prevalence of extreme (0 or 1) values in the raw MPC distribution mostly for small but not for large shocks is suggestive of near-rational behavior. But, to more formally evaluate whether a near-rational model can rationalize this behavior, we need to measure the within-household responses to shocks and then test for the characteristics of near-rational behavior implied by the model.

Our model implies that, for very small shocks, household behavior is entirely determined by the quick-fix that they employ (see Corollary 1). Thus, to measure quick-fixes, we use the information about households' responses to the smallest shocks of \$50 and −\$50. In the survey, 68% of respondents report an extreme MPC for *both* of these shocks, indicative of quick-fixes that feature such extreme MPCs. Motivated by this, we exhaustively categorize these 68% of households into four quick-fixing types based on whether they respond by consuming or saving all of these small shocks.

Consumption fixers (14% of households) default to an MPC of 0 for gains and losses. Their consumption is fixed, and they absorb small shocks—gains or losses—with their savings.

Savings fixers (29% of households) default to an MPC of 1 for gains and losses. Their saving is fixed, and they absorb small shocks—gains or losses—with their consumption.

Consumption prioritizers (11% of households) have a quick-fixing policy that prioritizes their consumption. They only draw on their savings (MPC = 0) when they face a small loss, but they increase only their consumption (MPC = 1) when they face a small gain.

Savings prioritizers (14% of households) have a quick-fixing policy that prioritizes their savings. They cut back only consumption (MPC = 1) when they face a small loss, but they increase only savings (MPC = 0) when they face a small gain.

The remaining group of **unclassified households (32%)** cannot be assigned to either 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 quick-fixing behavior among this group. Unclassified households almost never choose extreme MPCs and instead immediately adopt relatively stable interior MPCs. Their 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

Figure 2: Responses to income shocks for different quick-fixing types

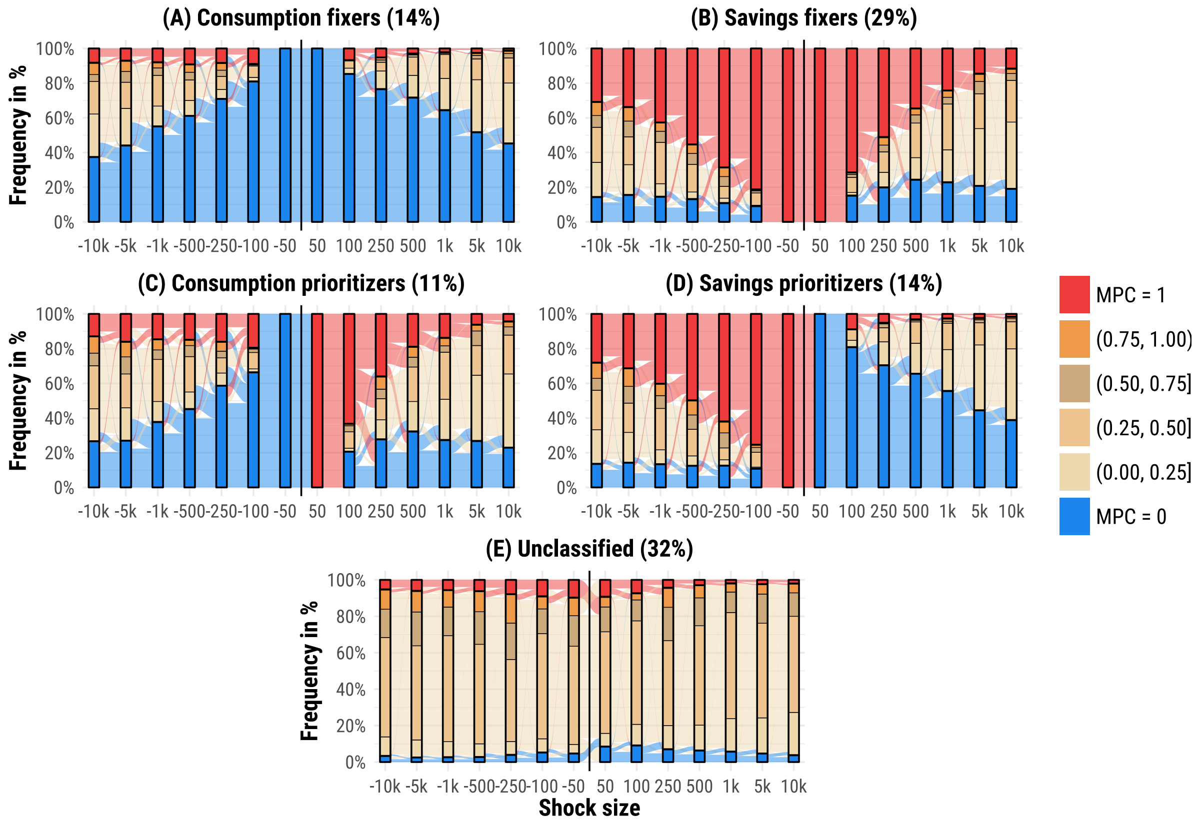

Notes: The alluvial graphs summarize the MPC data of four different quick-fixing types, which we define on page 17, and unclassified respondents. In each panel, each of the 14 columns displays the distribution of MPCs for one particular shock size, and the streams between bars indicate how households' MPCs transition between two neighboring shocks.

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.

To help visualize the behavioral differences between these types, Figure 2 decomposes the aggregate MPC data and its "bowtie" pattern for each 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. For larger shocks, more consumption fixers adopt and then stick to an interior MPC. Panel B plots the MPCs of savings fixers and reveals a similar logic. Savings fixers respond with an MPC of 1 to small shocks but transition and then stick to interior MPCs for larger shocks. The same story emerges for consumption prioritizers and savings prioritizers: starting from extreme MPCs, they eventually move to interior MPCs. Of course, Figure 2 also reveals that not all households follow their measured type perfectly. For example, moving from \$50 to \$100 and from \$100 to \$250, 8% of households move from an MPC of 0 to 1 or vice versa. But, in light of the inherent response noise in survey data, we do not over-interpret this pattern (see also footnote 5).

Evaluating Near-Rationality. These four widespread response types in our data are intuitively plausible quick-fixes: they prescribe simple responses to sufficiently small shocks. But, to more formally evaluate whether a near-rational model can rationalize this behavior, we now test for the characteristics of near-rational behavior as described in Section 2.4.

First, households are more likely to apply quick-fixes for smaller shocks than for larger shocks. This is apparent from Figure 2: within each type, the fraction of extreme responses declines as the shock size increases. Moreover, within each type, households tend to transition from extreme to interior MPCs, as visualized by the flows in Figure 2.

Second, quick-fixes are abruptly abandoned once a household-dependent critical shock size is reached. In the survey, shifts from extreme to interior MPCs rarely occur gradually. Households who 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 ranges from 0.20 to 0.50. Figure 3 visualizes the abrupt transition from extreme MPCs (left column) to first interior MPCs (right column) across types.

Third, once households abandon their quick-fixes, their consumption policies are relatively stable *within* respondent and similar *across* respondents, even if their quick-fixing type differs. 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: for only 3% of shock size increases, we observe a transition from interior to extreme MPCs, and conditional on starting from an interior MPC, households stay in the interior for 93% 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 3 shows that, while MPCs vary widely across types before switching (left column), their interior MPCs after switching are similar (right column).

We summarize these findings below:

Figure 3: Distribution of MPCs for gains before and after switching to interior

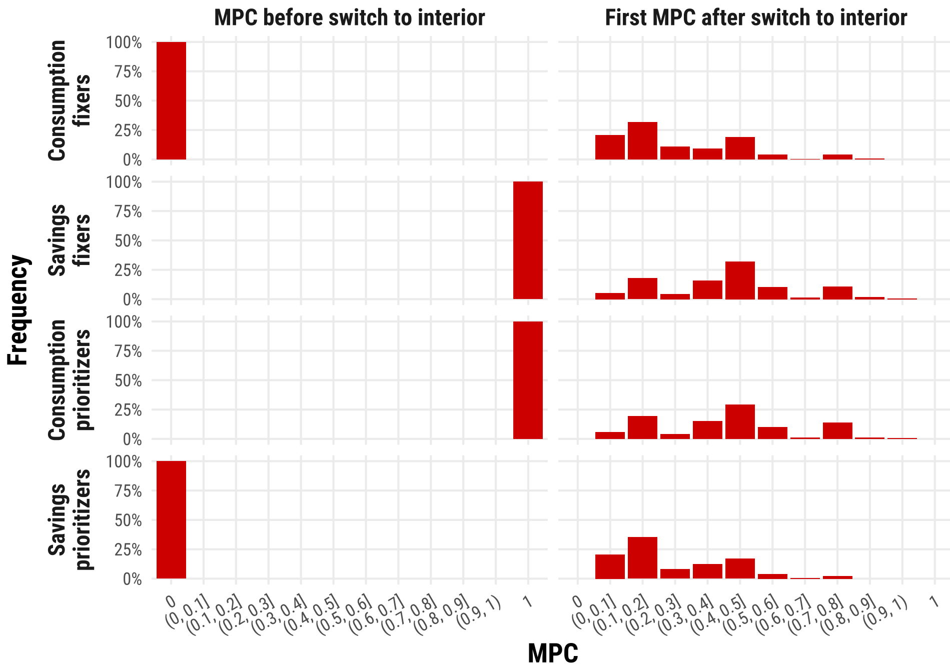

Notes: The histograms show conditional distributions of MPCs for gains. The rows correspond to the four quick-fixing types. The first column shows the distribution of MPCs before households switch to an interior MPC, which by construction puts all mass at either MPC = 0 or MPC = 1. The second column 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 losses is reported in Appendix Figure B.2.

> **Fact 2** (Quick-Fixing and Near-Rationality): 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 to small shocks to having similar interior MPCs for large shocks.

3.4 Quick-Fixes Account for MPC Heterogeneity

Differences in MPCs across households are notoriously hard to predict using observable characteristics of households' wealth, income, and demographics (*e.g.*, Lewis et al., 2024; Fuster et al., 2021). To test this in our data, we model respondents' average MPCs as a function of observables in the following regression equation:

$$\text{MPC}_i = \alpha + \beta' X_i + \epsilon_i \tag{10}$$

where MPC_i is each respondent's average MPC across the 14 scenarios and X_i is a vector of variables that includes: households' monthly spending (log), income (log), and income risk; non-parametric functions of liquid wealth, illiquid wealth, and debt; and four additional demographic characteristics (a college dummy, age, gender, and household size). Echoing findings in the literature, these characteristics together explain only 11% of the cross-sectional variation in MPCs (Appendix Table B.2). 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 characteristics should explain the majority of (or even all) heterogeneity in MPCs.

By contrast, the four quick-fixing types account for almost *half* of the variation in MPCs by themselves. That is, when we replace X_i with four dummy variables indicating the quick-fixing types, we obtain an R^2 of 49%. This may be surprising because, although quick-fixing types are identified using data on MPCs, they are a relatively low-dimensional summary that use information about only \$50 shocks. This suggests that households' quick-fixing types are a potentially major driver of MPC heterogeneity and could go a long way toward opening this "black box." Households' quick-fixing types also account for 56% of the variation in the share of extreme MPCs of 0 or 1.

Quick-fixing types not only account for much of the cross-sectional dispersion in MPCs, they also explain variation in the aggregate MPCs out of shocks of varying signs and sizes. To demonstrate this fact, Figure 4 plots the average MPC for each shock, pooled across all households and separately for each type. Panel A shows the raw average MPCs for all respondents (black line and dots) and separately for each type (colored lines and dots). Panel B shows the portion of MPCs that is predicted by the "extensive margin" of type membership and switching to the interior via a simple calculation. Suppose that households consistently follow their estimated type: they start with the MPC prescribed by their type, and once they switch to a different MPC for the first time in the data, they permanently switch to the average interior MPC of 0.45. Thus, MPCs only vary if

Figure 4: The extensive margin explains MPC means and heterogeneity

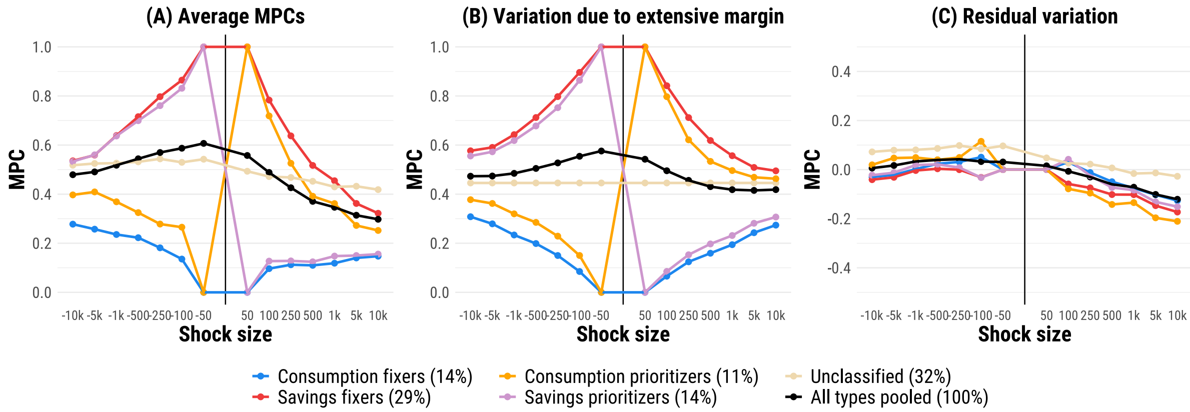

Notes: The left panel displays the average MPCs of four different quick-fixing types, which we define on page 17, and unclassified respondents for 14 different income shocks. The black line shows the average MPC of the full household sample. The middle panel shows the same statistics after "enforcing consistency" to isolate the effect of the extensive margin. 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.

households transition from extreme to interior MPCs. Panel B shows that this simple calculation accounts for many of the key patterns in the MPC data: (i) significant heterogeneity across the five groups, (ii) the high average MPC, (iii) the higher MPCs for losses versus gains, and (iv) the size-dependence of MPCs, that is, the decline in MPCs as shock sizes increase in absolute value.

Since the size-dependence in aggregate MPCs will be an important macroeconomic implication of quick-fixing, we briefly explain what generates this pattern. In the raw data, the MPC out of a \$100 gain is 0.06 higher than the MPC out of a \$250 gain and 0.14 higher than that out of a \$1,000 gain. The extensive margin (Panel B of Figure 4) accounts for roughly two-thirds of these differences (0.04 and 0.08, respectively). Quick-fixing via our four types could, in principle, lead to MPCs that are decreasing or increasing in shock size. But we observe that quick-fixing *reduces* MPCs for larger shocks for two reasons. First, our data features more savings fixers and consumption prioritizers who transition from MPC $= 1$ to the interior than consumption fixers and savings prioritizers who transition from MPC $= 0$ to the interior. Second, due to an average interior MPC below 0.5, abandoning an MPC of 1 has a larger impact than abandoning an MPC of 0.

Finally, Panel C of Figure 4 shows the residual variation, *i.e.*, the difference between Panel A and B, which is much smaller and similar across all quick-fixing types. This is a different way of showing that quick-fixing types behave similarly conditional on adopting an interior MPC, consistent with our model. The residual variation also illustrates what a simple framework with constant interior MPCs does not capture: (i) falling interior MPCs for higher shock sizes and (ii) lower interior MPCs for gains than for losses. Although this is not the focus of our analysis, our full model in Section 4 does not impose constant interior MPCs (i.e., allows for a concave consumption function out of income).

We summarize these findings below:

> **Fact 3** (Quick-Fixes Explain MPCs): The four quick-fixing types explain both a large share of variation in average MPCs across households and aggregate MPCs out of shocks of varying sizes and signs.

3.5 Household Characteristics Do Not Explain Quick-Fixing

The importance of quick-fixing for explaining 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). Theoretically, our framework in Section 2 imposes no restriction on these relationships.

To investigate this, we estimate a series of linear probability models that try to predict our type classification across households i:

$$1_{i \text{ has type } j} = \alpha + \beta' X_i + \epsilon_i \tag{11}$$

where the outcome is an indicator for being categorized as one of our four types (consumption fixer, savings fixer, consumption prioritizer, or savings prioritizer) and the regressors X_i are the same as those studied in the previous subsection, describing wealth, income, and demographics.

We find only a weak relationship between household observables and quick-fixing types, with R^2 values ranging from 0.02 to 0.06 (Columns 1–4 of Table B.3). To mention just two positive examples of predictive relationships, liquidity-constrained households are 5 percentage points more likely to be savings fixers (MPC = 1), while high-liquidity households are 7 percentage points more likely to be consumption fixers (MPC = 0). We observe a clearer relationship between households' characteristics and whether we can classify them as 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.

Can we predict for how long households stick to their quick-fix? The theoretical model has ambiguous predictions for how household characteristics should affect the losses from quick-fixing (and thereby the propensity to optimize). Nevertheless, this is an interesting question to investigate empirically. We find that households' switching thresholds—the smallest shock for which they switch from an extreme to an interior MPC—vary little across quick-fixing types ($R^2 = 0.02$, Column 6). They are smaller for households of lower age, intermediate liquid wealth levels, debt, and low income, but households' economic characteristics can account only for 6% of the total variation (Column 7).

A corollary of these findings is that we observe a marked transition from extreme MPCs to interior MPCs among households from *all ranges* of the socioeconomic distribution. Appendix Figure B.3 visualizes this result by showing the "bowtie" distribution of MPCs for households across the distributions of wealth (liquid and illiquid) and debt.

We summarize this finding below:

> **Fact 4** (Quick-Fixes Are Unpredictable): Quick-fixing behavior is essentially unpredictable from households' economic and demographic characteristics.

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 think harder about how to respond to large shocks. To test this behavioral mechanism, we investigate how carefully households consider their responses to income shocks of various sizes and signs. We draw on an additional survey, which we conducted with 517 additional US households in August 2024. We measure their 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.[6]

[6] We 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 Appendix Table B.4). Appendix C.2 contains the additional survey instructions.

Figure 5: Self-reported deliberation and extreme MPCs

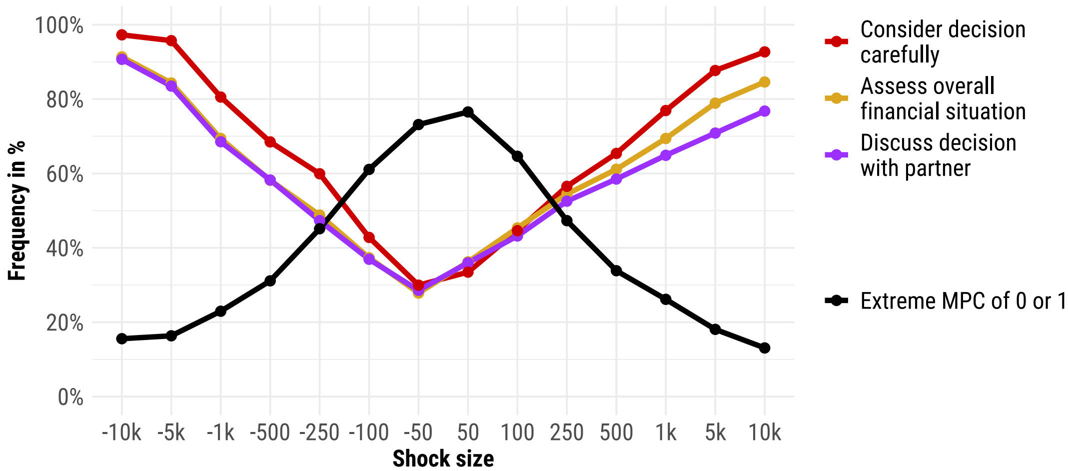

Notes: Data from the additional deliberation study (517 US households). For fourteen different income shocks, the figure shows the average frequency with which households report to consider their response carefully (a score of at least four on a six-point scale, red line), assess their overall financial 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.

Figure 5 shows that deliberation strongly increases with shock size. 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). Likewise, the likelihood that respondents consult other household members increases from 32% to 84% (purple line).

These findings are consistent with the theoretical mechanism that quick-fixes avoid the cost of deliberation and optimization. We find that increases in deliberation mirror the *decreasing* likelihood that households choose an extreme MPC of 0 or 1 (black line). On 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 (Appendix Table B.4). We 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 (Appendix Table B.5).

We summarize these findings below:

> **Fact 5** (Quick-Fixing Requires Less Deliberation): Households think more carefully about their consumption-savings responses when they face 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 households to explain in their own words why they switch from extreme MPCs for a small shock ($100) to interior MPCs for a large shock ($1,000) (see Appendix C.5 for details).

We summarize four patterns in these qualitative data. First, almost all households refer to the difference in shock size to explain the difference in their MPCs. Second, households mention 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. Third, 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. Fourth, many households recognize an income gain as a welcome opportunity to treat themselves or 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. We relegate a more detailed discussion to Appendix C.5. Of course, it seems plausible that further psychological forces are at work, which are harder for households to explicitly articulate. For example, finding a good 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.

3.7 Can Existing Models Rationalize These Data?

In short, we argue the answer is no. 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.

Incomplete Markets Models. The standard incomplete markets model 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.

Incomplete Markets Models with Heterogeneity. This failure of the incomplete markets model to match the data is not remedied by the many extensions of the model that feature *ex ante* heterogeneity in preferences and investment technologies among agents. Such models account for heterogeneous discount factors (Carroll et al., 2017; Aguiar et al., 2024), elasticities of intertemporal substitution (Aguiar et al., 2024), risk aversion (Kaplan and Violante, 2022), and investment opportunities (Kaplan and Violante, 2022). Importantly, our design elicits consumption responses within-subject. The consumption fixers, savings fixers, consumption prioritizers, and savings prioritizers that we detect in the data constitute 68% of all respondents, and none of these respondents have the behavior predicted by an incomplete markets model. Thus, *ex ante* heterogeneity cannot explain our findings.

Incomplete Markets Models with Multiple Assets or Accounts. 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 et al. (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 B.3, we document a pronounced "bowtie" pattern of adjustment for all levels of liquid wealth. Even households who have more than $100,000 of liquid wealth or liquid wealth exceeding ten times their monthly income display the "bowtie" pattern. Thus, liquidity and costly reallocation of assets across accounts cannot account for the behavior of our respondents.

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 et al., 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 as positive shocks get larger. These models are also inconsistent with the behavior of consumption prioritizers and savings fixers, which 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.,* Laibson et al., 2021), temptation (see, *e.g.,* Krusell et al., 2002), rational inattention (Sims, 2003), finite planning horizons (Boutros, 2022), sparsity (Gabaix, 2014), 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. To explain the failure of conventional models to generate a realistic equity risk premium, Gabaix and Laibson (2001) propose a model in which agents optimize and adjust their savings accounts only every D periods. This model does not generate the substantial fraction of agents in our data that report adjusting their savings within the next three months, as Gabaix and Laibson (2001) argue that D is at least one year. Other models generate infrequent optimization through "menu costs" of adjusting behavior (see, *e.g.*, Fuster et al., 2021). In our data, quick-fixing households adjust their behavior even in response to small shocks but differ in whether they adjust along the margin of consumption or spending. This is consistent with our model with optimization costs and heterogeneous quick-fixes and inconsistent with a model with adjustment costs.

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 avoid the costs of figuring out the optimal balance between spending and savings. 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 motivate many of the approaches described above.

4 Quantitative Model

To study the macroeconomic implications of our findings, we incorporate quick-fixing into a quantitative dynamic model of consumption and savings. This model enriches the framework of Section 2 with three features present in benchmark quantitative models: more than two periods, stochastic income, and borrowing constraints. In this section, we set up the model and calibrate it to match our survey evidence.

4.1 Set-up

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 behaviors, as we elaborate on below.

Rational Behavior. We first introduce dynamic optimization of rational households, whom we henceforth refer to as type "R." We let $V^R : [0,\infty) \times \mathcal{Y} \to \mathbb{R}$ denote their value function, defined by

$$\begin{aligned} V^R(a,y) = \max_{a'} &\left\{u(c) + \beta\mathbb{E}\left[V^R(a',y') \mid y\right]\right\} \\ \text{s.t. } & a' = Ra + y - c \\ & a \geq 0 \end{aligned} \tag{12}$$

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

Quick-Fixing Households. 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 *reoptimization 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 $q \in \{\mathrm{CF}, \mathrm{SF}, \mathrm{CP}, \mathrm{SP}\}$, that depend on reference consumption $\bar{c}$ and an income deviation $y - \bar{y}$:

$$\begin{aligned} c^{\mathrm{CF}}(\bar{c}, y-\bar{y}) &= \bar{c} & c^{\mathrm{CP}}(\bar{c}, y-\bar{y}) &= \bar{c} + \max\{y-\bar{y}, 0\} \\ c^{\mathrm{SF}}(\bar{c}, y-\bar{y}) &= \bar{c} + (y-\bar{y}) & c^{\mathrm{SP}}(\bar{c}, y-\bar{y}) &= \bar{c} + \min\{y-\bar{y}, 0\} \end{aligned} \tag{13}$$

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 as reduced savings. *Savings fixers* (SF) save a fixed amount, $s^{\mathrm{SF}} = y - c^{\mathrm{SF}}(\bar{c}, y-\bar{y}) = \bar{y} - \bar{c}$. *Savings prioritizers* (SP) save a fixed amount plus the *negative* component of income shocks, absorbing the positive component as increased savings.[7]

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 reoptimize by adopting the unconstrained, rational choice, in which case $D = 1$ (D for deliberation). If the household reoptimizes, then they also reset their reference consumption and income states to $c^*(a, y)$ and y, respectively. For each type q, this behavior is described by the dynamic program

$$\begin{aligned} V^q(a, y, \bar{c}, \bar{y}) = \max_{D\in\{0,1\}} &\left\{D\left(u(c^*(a,y)) + \beta\mathbb{E}\left[V^q(a', y', c^*(a,y), y) \mid y\right] - \kappa_q\right)\right. \\ &\left. +(1-D)\left(u(c^q(a,y,\bar{c},\bar{y})) + \beta\mathbb{E}\left[V^q(a', y', \bar{c}, \bar{y}) \mid y\right]\right)\right\} \\ \text{s.t. } a' &= Ra + y - (D(c^*(a,y)) + (1-D)(c^q(a,y,\bar{c},\bar{y}))) \\ a &\geq 0 \end{aligned} \tag{14}$$

Two modeling choices that are necessary in the dynamic model are the treatment of the persistence of types and 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 finding that household wealth and financial status are poor predictors of quick-fixing types (Fact 4). Second, households reoptimize by adopting the choice of rational households. That is, households are not sophisticated in the sense that their optimizations embed the best way to manipulate their future quick-fixing choices. We argue that this approach is most consistent with

[7] If c^q is negative, we adopt the convention that the household automatically reoptimizes.

the idea of near-rationality. Indeed, the costs of such naivete are bounded above by the costs of near-rationality, which we will shortly find are very small, making the distinction quantitatively irrelevant for households' welfare.

One-Time Income Shocks. We next describe how the household responds to one-time shocks like the hypothetical scenarios of our survey. 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 an unanticipated shock x, is

$$\begin{aligned} \max_{D_x \in \{0,1\}} & \{D_x \left(u(c^*(a, y+x)) + \beta \mathbb{E}\left[V^q(a', y', c^*(a, y+x), y) \mid y\right] - \kappa_q\right) \\ & \quad + (1 - D_x) \left(u(c^q(c^{q^*}, x)) + \beta \mathbb{E}\left[V^q(a', y', \bar{c}^*, \bar{y}^*) \mid y\right]\right) \} \\ \text{s.t. } & a' = Ra + y - (D_x(c^*(a, y+x)) + (1 - D_x)(c^q(c^{q*}, x))) \\ & a \geq 0 \end{aligned} \tag{15}$$

If a household reoptimizes, then it follows the rational, forward-looking behavior embodied in c^*. This entails an optimization cost. If the household quick-fixes, then it treats c^{q*} as its reference consumption and the shock x as the income shock (Equation 13).

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.[8] We scale income such that one unit coincides with the median quarterly income reported in our survey, \$15,625. We set the quarterly interest

[8] This 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.

rate to $R = 1.01$.

Second, we calibrate the discount rate to match the spending behavior of households who do *not* quick-fix in the survey (*i.e.,* are "unclassified"). For the "rational" or non-quick-fixing type in the model, we calculate the average MPC out of transfer shocks of size x as $\text{MPC}_x^R = \frac{1}{x}\int (c^*(a, y + x) - c^*(a, y))\, \mathrm{d}\Phi^R(a, y)$, where $\Phi^R(a, y)$ is the stationary distribution over assets and income for rational agents. 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 $\beta = 0.92$.

Third, we calibrate the fraction of agents of each quick-fixing type to match the categorization in Figure 2. The fraction of rational agents 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 rational households and is therefore conservative for the near-rational theory.

Fourth, we calibrate the four type-specific reoptimization costs, $(\kappa_{\text{CF}}, \kappa_{\text{SF}}, \kappa_{\text{CP}}, \kappa_{\text{SP}})$, to match our main findings about quick-fixing behavior in the survey (Figures 1 and 2). 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 reoptimize by reporting a propensity to consume that does not coincide with the quick-fix. In the model, we calculate, for each shock x,

$$\text{ReoptFraction}_x^q = \int D_x^{q*}(a, y, \bar{c}, \bar{y})\, \mathrm{d}\Phi^q(a, y, \bar{c}, \bar{y}) \tag{16}$$

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. The optimization 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 model versus data:

$$\kappa_q^* = \arg\min_{\kappa_q > 0} \left\{ \sum_{i=1}^{12} \left(\text{ReoptFraction}_{x_i}^q - \widehat{\text{ReoptFraction}}_{x_i}^q \right)^2 \right\} \tag{17}$$

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 2 in Section 5.1.

4.3 Model Fit

Reoptimization. Figure 6(i) compares the model prediction and data for the key moments that discipline the costs of reoptimization, namely the propensities to reoptimize for different shock sizes. The model fits the overall adjustment pattern quite well: as shocks get larger, considerably more individuals reoptimize. Because the model is overidentified, with 48 moments and 4 parameters for the final step of calibration, 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 under-estimates adjustment for low shock sizes and over-estimates adjustments for large shock sizes. To rationalize optimization in response to relatively small shocks like gaining or losing $250, the model struggles to rationalize *not* reoptimizing after gaining or losing $10,000. This is a limitation for the model's fit to the data. But, as large shocks of $10,000 are potentially unusual and hard to think about, this is exactly where we might expect the greatest error in the survey responses.

Average MPCs. Figure 6(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 6(i)) and the MPC profiles of unclassified agents (see Figure B.4), it does not target the average MPCs of quick-fixing agents. The model deviation in average MPCs is small (Panel A). The maximum difference is attained for small 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 8). This generates the correct qualitative pattern that residual variation is largest for small negative shocks but slightly overstates the quantitative magnitude.

The MPC Distribution. We finally re-create our key empirical finding about the shape of the MPC distribution (Appendix Figure B.5). 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)

Figure 6: Quantitative model fit

(i) Reoptimization in Response to Shocks

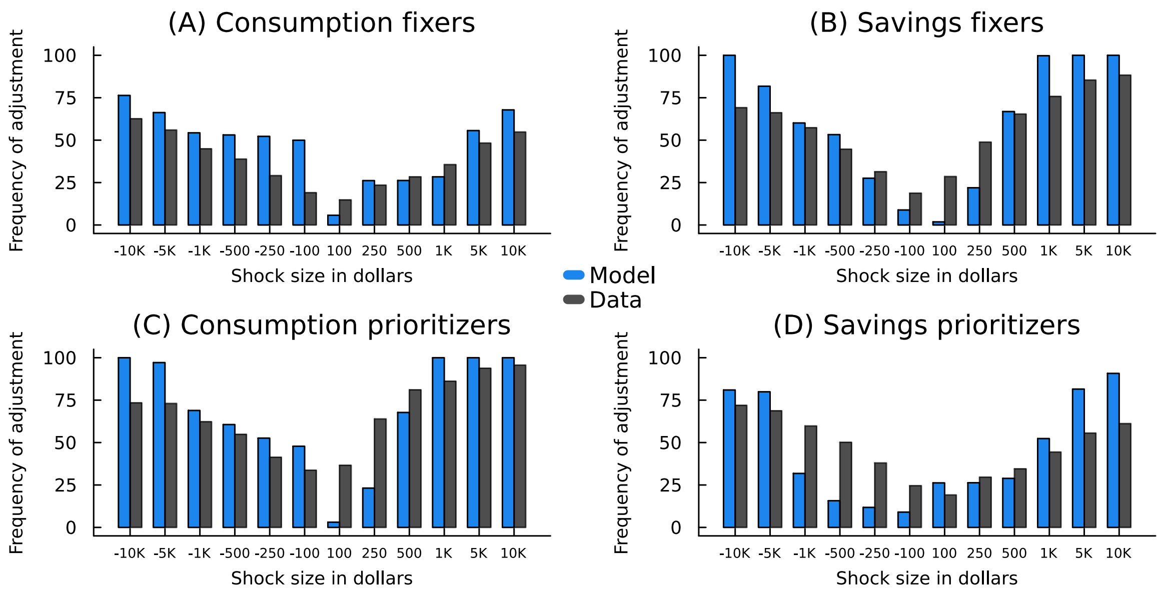

Notes: The bar graphs compare empirical measurements and model predictions for the propensity of agents to reoptimize following unanticipated shocks. Each panel corresponds to one of the four types and therefore to the calibration of the corresponding optimization cost parameter. The blue bars denote model predictions, described in Equation 16, and the grey bars are empirical measurements, as reported in Figure 2.

(ii) Average Marginal Propensities to Consume (MPCs)

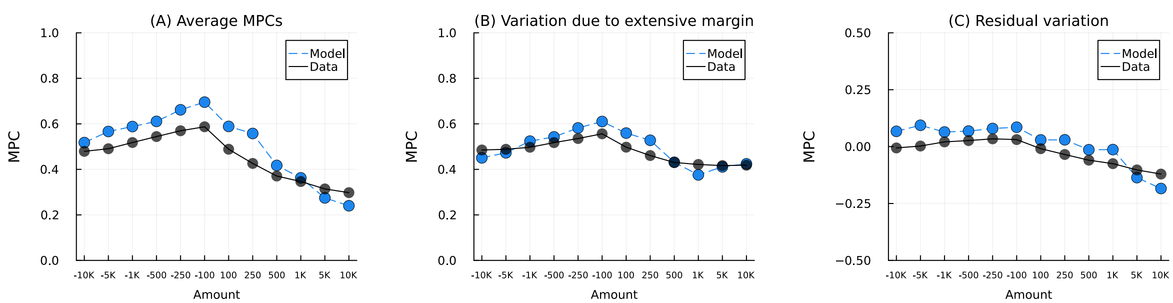

Notes: The line graphs compare empirical measurements and model predictions for the average marginal propensity to consume out of shocks of different amounts. The decomposition corresponds exactly to that of Figure 4. Panel A shows average MPCs. Panel B shows variation due to the extensive margin, calculated 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 4).

Table 2: The small costs of optimization and near-rationality

Panel A: Optimization costs κ_q		
Household type	% reduction in consumption	Average dollar cost
Consumption fixer	1.10	\$176.67
Savings fixer	0.007	\$1.47
Consumption prioritizer	0.006	\$1.44
Savings prioritizer	0.11	\$18.56

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	\$71.67
Savings fixer	0.004	\$0.58
Consumption prioritizer	0.003	\$0.54
Savings prioritizer	0.06	\$8.68

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

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.

5 Quantitative Results and Macro Implications

Having estimated the model, we now explore its economic properties. We first show quantitatively that the near-rational model predicts large differences in behavior compared to the nested rational model, but with economically very small losses from near-rational behavior. We next use the model to decompose how quick-fixing explains latent heterogeneity in the MPC, consistent with our empirical findings. We finally use the model to gauge the model's implications for the efficacy of government transfers to stimulate aggregate consumption.

5.1 The Losses from Near-Rationality Are Very Small

In Panel A of Table 2, we show that the calibrated costs of optimization 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.[9] The costs are by some margin highest for consumption fixers, as this strategy features relatively little reoptimization in the survey. 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 rational agents if they were to adopt the behaviors and bear the decision costs of quick-fixers:

$$\Delta V^q = \int (V^R(a,y) - V^q(a,y,c^*(a,y),y))\, \mathrm{d}\Phi^R(a,y) \tag{18}$$

We express this in units of an equivalent per-period reduction in consumption.[10] One expects this to be lower than the optimization costs because it (i) averages over a lifetime in which reoptimizations do not always occur and (ii) accounts for the benefits from reoptimizing.

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 2). On average, among quick-fixers, the loss is \$17 per quarter. 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 Households Tolerate Large Deviations from Rationality

We now show that, despite the very small opportunity costs of near-rational behavior, many households nevertheless tolerate and make large consumption mistakes. Figure 7 visualizes the states in which households quick-fix rather than reoptimize as shaded "inaction regions" around the rational consumption function, holding fixed income at its median state. The width of the inaction region can be interpreted in units of mean quarterly income. In dollars, at the asset state $a = 1$ (*i.e.*, assets equal to one quarter's income), the inaction regions correspond to tolerated deviations of \$3,466 (consumption fixers), \$168

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

[10]The equivalent percentage reduction in any consumption stream $\{c_t\}_{t=0}^\infty$ solves $\Delta V^q = \sum_{t=0}^\infty \beta^t \log(c_t) - \sum_{t=0}^\infty \beta^t \log\big((1-\delta^d)c_t\big)$ and is therefore $\delta^d := 1 - e^{-(1-\beta)\Delta V^q}$.

Figure 7: Inaction regions in the calibrated quick-fixing model

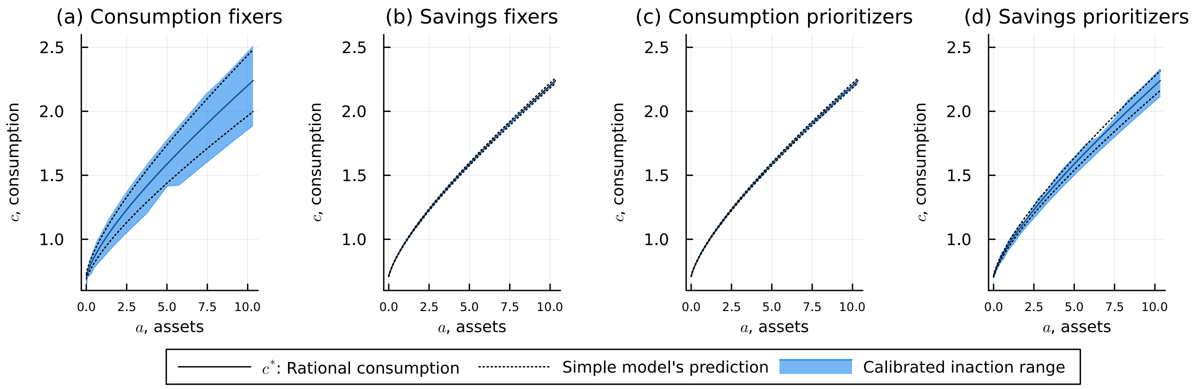

Notes: Each panel corresponds to one of the four behavioral optimization types and shows that type's consumption inaction behavior as a function of assets in the median income state. If consumption and assets lie in the blue shaded area, the agent does not reoptimize their consumption level. The black line is the consumption function for rational, always-optimizing agents. The dotted lines show the boundaries of the predicted inaction region from the simple model (Proposition 1).

(savings fixers), \$10 (consumption prioritizers), and \$1,308 (savings prioritizers). These are all an order of magnitude larger than the equivalent dollar cost of reoptimizing (Table 2), as one would predict from the logic that losses from misoptimization are second order in the size of consumption deviations.

We also observe that the very different widths of inaction bands across types are consistent with similar marginal probabilities of quick-fixing in the stationary distribution: 68% for consumption fixers, 53% for savings fixers, 52% for consumption prioritizers, and 63% for savings prioritizers. This is because quick-fixes that fix savings *do* allow consumption to respond (one-to-one) to shocks and, moreover, achieve close to the rational response to a highly persistent income shock when away from the borrowing constraint. They *exactly* coincide with the rational response when the borrowing constraint binds strictly.

We finally note that the approximate calculation of inaction regions from the two-period model (Proposition 1), denoted with dotted lines in Figure 7, is close to correct in this much richer setting. Intuitively, with a persistent income state, the continuation ($t+1$) costs of quick-fixing—a force necessarily ignored in the two-period model—scale very closely with the static (t) costs. Encouragingly, this implies that the simple economics of the two-period model largely apply in this richer setting. The approximation is relatively worse for lower asset households because of two forces missing in the simple model: the changing second derivative of the utility function (prudence) and the borrowing constraint.

Figure 8: Wealth distributions by type

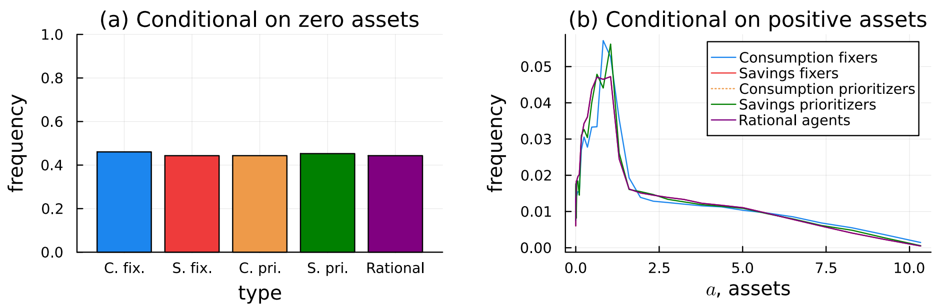

Notes: This figure shows the distribution of assets for each type. Panel (a) shows the frequency of households at the borrowing constraint. Panel (b) shows the distribution of wealth conditional on $a > 0$.

5.3 Quick-Fixing is Unrelated to Wealth, But Generates MPC Heterogeneity

One moment that we did not use for calibration is the limited relationship between quick-fixing types and wealth (Fact 4). 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 who default to increasing saving in response to shocks (savings prioritizers) accumulate considerably more wealth than those who default to increasing consumption in response to shocks (consumption prioritizers).

We find instead that the wealth distributions of each of the four quick-fixing types as well as rational agents are essentially identical (Figure 8). 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. 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 findings also contrast with those obtained in the class of models of agents with heterogenous discount factors and/or present bias, which naturally lead to very different long-run rates of saving.

We next explore how much quick-fixing contributes to heterogeneity in the marginal propensity to consume. Table 3 reports the percent of variance in this object that can be explained by assets and income in the quick-fixing model and the nested rational model.

Table 3: Variance in MPCs unexplained by assets and income

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

Notes: In each model, we calculate $\text{Var}[MPC_i|a_i, y_i]$, 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. We report $100 \cdot \mathbb{E}[\text{Var}[MPC_i|a_i, y_i]]/\text{Var}[MPC_i]$, or the fraction of variance unexplained by wealth and income. Since (a, y) is the state variable in the rational problem, this is 0 in that model.

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 rational model, 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. Thus, quick-fixing helps break the tight connection in incomplete markets models between financial observables and MPCs.[11]

5.4 Implications for the Aggregate Response to Macro Shocks

We finally use the model to study how quick-fixing shapes the response of aggregate consumption to income shocks. To do this, we use the quantitative model to calculate a sequence of *intertemporal marginal propensities to consume* ("iMPCs"), or consumption responses at various future horizons in response to contemporaneous transfer shocks. A leading application of this analysis is to study how the economy responds to an aggregate transfer shock (*i.e.,* "stimulus checks"). In this context, iMPCs measure how quickly consumers spend their checks and therefore how "front-loaded" the stimulus is. More generally, iMPCs are sufficient statistics for calculating the first-order effects of shocks on aggregate consumption in the "household block" of (New) Keynesian models (Auclert et al., 2024). Under standard assumptions in the modern HANK literature, this also

[11] This holds even while we enforce a strict borrowing constraint. More radically, one could argue that quick-fixing could deliver a high average MPCs and heterogenous MPCs even without a strict borrowing constraint. In such a model, almost all of the variation in MPCs would be attributed to quick-fixing.

corresponds to the general equilibrium effect of a broader class of shocks (Auclert et al., 2024; Angeletos et al., 2024).

Concretely, for individual i with current idiosyncratic state z_i, we can define their iMPC at horizon h out of an income change x occuring at time t:

$$\text{MPC}_{h,x}(z_i) = \frac{\mathbb{E}[c_{i,t+h}|z_i, x] - \mathbb{E}[c_{i,t+h}|z_i, 0]}{x} \tag{19}$$

That is, averaged over realizations of income and expressed as a fraction of the shock, how much more does that household consume at $t + h$ because of an unanticipated shock of size x at t? Due to the budget constraint, intertemporal MPCs must add up to unity in present value: $\sum_{h=0}^{\infty} \frac{1}{(1+r)^h} \text{MPC}_{h,x}(z_i) = 1$ for all z_i. While our survey directly disciplines the contemporaneous MPC, the quantitative model is necessary to understand how quick-fixing households spend their savings over subsequent quarters.

iMPC profiles vary considerably across types, shock sizes, and horizons (Figure 9(i)). For a \$100 shock, savings fixers and consumption prioritizers consume almost everything in one quarter, consumption fixers and savings prioritizers keep about 40% saved after one year, and rational consumers behave somewhere in between. For a \$1,000 shock, by contrast, iMPC profiles are more uniform across groups due to most types' reoptimizing—except for a significant fraction of consumption fixers who, consistent with our survey evidence, fully save even these large amounts.

To better illustrate how quick-fixing and reoptimization drive these patterns, Figure 9(ii) shows the impulse response of the fraction of agents who reoptimize. In response to small shocks (Panel (a)), very few households reoptimize on impact, and essentially no one does after three quarters. As shocks get larger (Panels (b) and (c)), households reoptimize both on impact and after many quarters.

Comparing the Near-Rational and Rational Models. Using the model, we can formally evaluate how the presence of quick-fixing affects aggregate consumption by comparing it to a world with only rational households. To isolate this, Figure 9(iii) plots iMPCs at specific horizons (1, 2, 4, and 6 horizons) for each gain scenario under the calibrated quick-fixing model and a restriction to only rational agents. For large shock sizes (\$1,000 and above), the models are much closer to one another at all horizons. This is natural because, as uncovered in the survey, most households do not quick-fix when confronted with large shocks. For small shocks (less than \$500), the rational model under-states contemporaneous responses and over-states long-run iMPCs.

Figure 9: Dynamic consumption responses to transfer shocks

(i) Profiles of intertemporal marginal propensities to consume

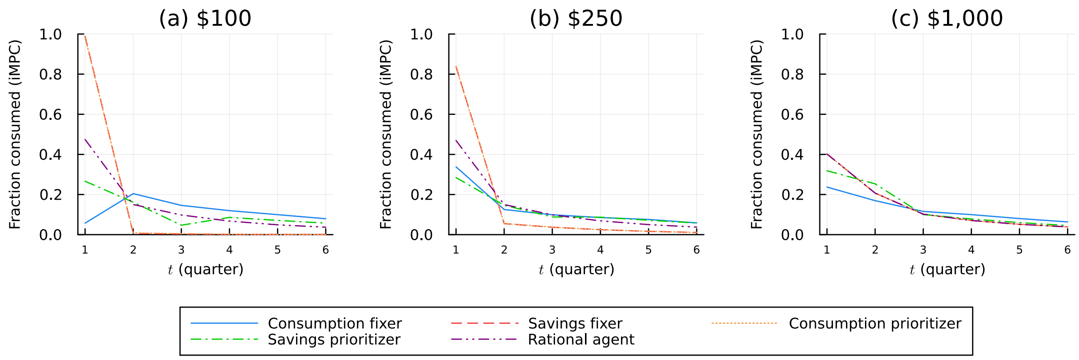

Notes: Each panel shows the intertemporal marginal propensity to consume (iMPC) in response to a different stimulus experiment. Each line corresponds to a different household type.

(ii) Profiles of delayed reoptimization

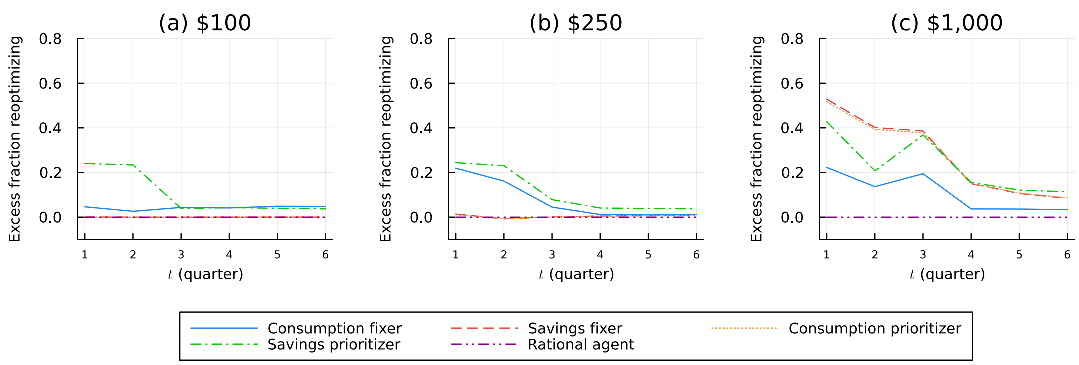

Notes: Each panel shows the fraction of additional households that reoptimize compared to the steady state in response to a different stimulus experiment. The fraction at $t = 1$ is calculated based on whether households reoptimize in response to an income shock *or* the unanticipated stimulus. Each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

(iii) Comparing iMPC profiles with the rational model

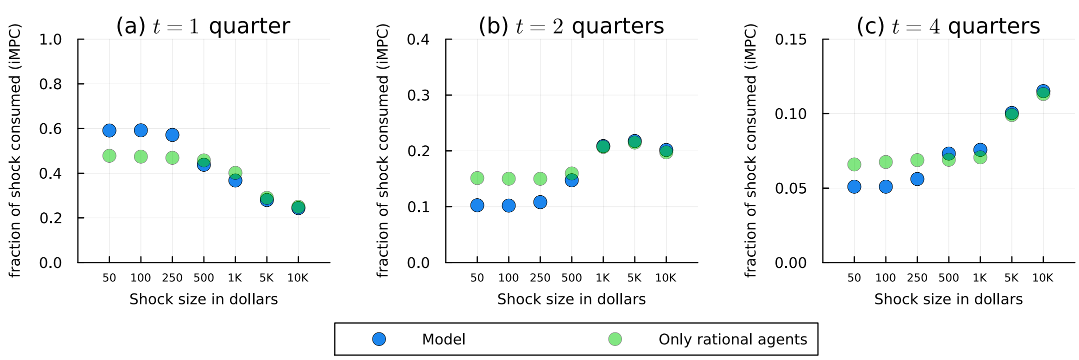

Notes: Each panel shows the intertemporal marginal propensity to consume (iMPC) averaged over households of all types in response to the specified shocks (x-axis) at the specified horizon. Blue dots correspond to the model prediction. Green dots correspond to the subset of rational agents.

Table 4: Impact MPCs depend significantly on shock size

Amount	**Context**	MPC in **Full Model**	MPC in **Rational Model**
\$100	Survey	0.59	0.47
\$600	2001 Tax Rebate	0.42	0.44
\$1200	2020-21, Round 2	0.35	0.39
\$2400	2020-21, Round 1	0.31	0.34
\$2800	2020-21, Round 3	0.30	0.33

Notes: This table reports contemporaneous (one-quarter) MPCs in the full model and the nested rational model for positive transfers of different amounts. Rows 2-5 correspond to baseline transfer amounts for two-taxpayer households in the indicated historical US stimulus policies.

Size-Dependence in Transfer Responses. We also can revisit the size-dependent response to transfers uncovered by the survey and formally relate this to the presence of near-rationality. Table 4 calculates average MPCs for four benchmark amounts that have been used in US stimulus design as well as the \$100 experiment in our survey. For each historical example, the exact size of transfer payments depends on other features (*e.g.*, the number of taxpayers and children in each household), so we benchmark to the payment for a two-taxpayer household. Transfer amounts have varied significantly over time. Our near-rational model suggests that if these amounts were transferred to a population that approximates current US households, their stimulus effects would widely differ. Moreover, much smaller stimulus checks, for instance of \$100, could potentially have a significantly larger per-dollar effect on aggregate consumption. This prediction of sharp size-dependence is *not* shared with the rational model. Concretely, the difference in MPCs between a \$100 check and a \$1,200 check is 0.24 in the quick-fixing model versus 0.09 in the rational model, an almost three-fold difference between models.

A Lucas Critique for (i)MPCs. An additional takeaway is that measurements of MPCs based on transfers of a particular amount may not be informative about MPCs out of different amounts. In particular, those based on large income shocks may understate the MPC in counterfactuals with small income shocks, and *vice versa*. This is a particular articulation of the Lucas (1976) critique of consumption functions, which we show is quantitatively relevant for measuring (i)MPCs.

6 Conclusion

This paper starts from the idea that, when optimizing consumption-savings decisions is costly, people may instead rely on simple near-rational policy functions: *quick-fixes* that avoid the cost of optimization. We develop a model of quick-fixing to capture this idea. By asking almost 5,000 US households about their consumption response to a large number of hypothetical scenarios, 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.

Quick-fixing has implications for consumption and savings behavior at both microeconomic and macroeconomic levels. First, quick-fixing is near-rational. Even very small costs of optimization can account for the economically large deviations from the rational benchmark in households' responses to income shocks. 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-dependence in the aggregate response to income shocks, quick-fixing has potentially important implications for business cycle dynamics and policy design. Fourth, given empirical quick-fixing behavior, care should be taken in extrapolating estimates of iMPCs into counterfactuals: if the distribution of income shocks under counterfactuals differs from that in the estimation sample, then these iMPCs may be substantially incorrect.

Finally, our findings convey a broader lesson about how to "put the near-rationality hypothesis to work" in other contexts. The premise that optimization 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 be valuable.

References

Aguiar, Mark, Mark Bils, and Corina Boar, "Who are the Hand-to-Mouth?," *Review of Economic Studies*, 2024. Forthcoming.

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," Working Paper 2024.

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.

Bordalo, Pedro, John Conlon, Nicola Gennaioli, Spencer Kwon, and Andrei Shleifer, "How People Use Statistics," *Working Paper*, 2024.

Boutros, Michael, "Windfall income shocks with finite planning horizons," Paper 2022.

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.

Colarieti, Roberto, Pierfrancesco Mei, and Stefanie Stantcheva, "The How and Why of Household Reactions to Income Shocks," Working Paper 2024.

Deaton, Angus, "Saving and liquidity constraints," *Econometrica*, 1991, *59* (5), 1221–1248.

Debortoli, Davide and Jordi Galí, "Heterogeneity and aggregate fluctuations: insights from TANK models," Technical Report, National Bureau of Economic Research 2024.

Enke, Benjamin, "The Cognitive Turn in Behavioral Economics," Working Paper 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.

— **and David Laibson**, "The 6D bias and the equity-premium puzzle," *NBER Macroeconomics Annual*, 2001, *16*, 257–312.
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.
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," *J. of Monetary Economics*, 2002, *49* (1), 107–135.
Kueng, Lorenz, "Excess Sensitivity of High-Income Consumers," *The Quarterly Journal of Economics*, November 2018, *133* (4), 1693–1751.
Laibson, David, Peter Maxted, and Benjamin Moll, "Present bias amplifies the household balance-sheet channels of macroeconomic policy," Technical Report, NBER 2021.
Lewis, Daniel J., Davide Melcangi, and Laura Pilossoph, "Latent Heterogeneity in the Marginal Propensity to Consume," Working Paper 32523, NBER 2024.
Lian, Chen, "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.
Lucas, Robert E, "Econometric policy evaluation: A critique," in "Carnegie-Rochester Conference Series on Public Policy," Vol. 1 North-Holland 1976, pp. 19–46.
Mackowiak, Bartosz, Filip Matejka, and Mirko Wiederholt, "Rational inattention: A disciplined behavioral model," Technical Report, Mimeo, New York City 2018.
Oprea, Ryan, "Complexity and Its Measurement," Working Paper 2024.
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.
Reis, Ricardo, "Inattentive consumers," *J. of Monetary Economics*, 2006, *53* (8), 1761–1800.
— , "Inattentive producers," *The Review of Economic Studies*, 2006, *73* (3), 793–821.
Sims, Christopher A, "Implications of rational inattention," *J. of Monetary Economics*, 2003, *50* (3), 665–690.
Woodford, Michael, "Modeling Imprecision in Perception, Valuation, and Choice," *Annual Review of Economics*, 2020, *12*, 579–601.

Quick-Fixing: Near-Rationality in Consumption and Savings Behavior
by Peter Andre, Joel P. Flynn, Georgios Nikolakoudis, Karthik A. Sastry

Online Appendix

A Omitted Proofs

A.1 Proof of Proposition 1

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

$$U(c, z_i) = u(c) + \beta u(R(a_i + y_{i,1} - c) + y_{i,2}) \tag{20}$$

and observe that $U^{q_i}(z_i) = U(c^{q_i}(z_i), z_i)$. Toward approximating, we can define the loss function on an enriched domain as $\mathcal{L}^U(c, z_i) = U^*(z_i) - U(c, z_i)$, and we can again observe that

$$\mathcal{L}^{q_i}(z_i) = \mathcal{L}^U(c^{q_i}(z_i), z_i) \tag{21}$$

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

$$\begin{aligned} \mathcal{L}^{q_i}(z_i) &= \mathcal{L}^U(c^*(z_i), z_i) + \mathcal{L}^U_c(c^*(z_i), z_i)(c^{q_i}(z_i) - c^*(z_i)) \\ &\quad + \frac{1}{2}\mathcal{L}^U_{cc}(c^*(z_i), z_i)(c^{q_i}(z_i) - c^*(z_i))^2 + O(|c^*(z_i) - c^{q_i}(z_i)|^3) \end{aligned} \tag{22}$$

We observe however that $\mathcal{L}^U(c^*(z_i), z_i) = 0$ by definition and that $\mathcal{L}^U_c(c^*(z_i), z_i) = 0$ by the optimality of c^*. Thus, we have obtained that:

$$\mathcal{L}^{q_i}(z_i) = \frac{1}{2}\mathcal{L}^U_{cc}(c^*(z_i), z_i)(c^{q_i}(z_i) - c^*(z_i))^2 + O(|c^*(z_i) - c^{q_i}(z_i)|^3) \tag{23}$$

We can moreover compute $\frac{1}{2}\mathcal{L}^U_{cc}(c^*(z_i), z_i)$ as:

$$\mathcal{L}^U_{cc}(c, z_i) = -\beta R^2 u''(R(a_i + y_{i,1} - c) + y_{i,2}) - u''(c) \tag{24}$$

which yields:

$$\mathcal{L}^U_{cc}(c^*(z_i), z_i) = -(1 + \beta R^2)u''(c^*(z_i)) \tag{25}$$

Completing the proof. $\square$

A.2 Calculation for Example 1

Define the consumption mistake as $m = [c^{q_i}(z_i) - c^*(z_i)]/c^*(z_i)$. 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:

$$\begin{aligned} 2\log(c^*(z_i)) - \kappa(m) &= \log((1 - \tau)c^*(z_i)) + \log(c^*(z_i)) \\ \implies \tau &= 1 - e^{-m^2} \end{aligned} \tag{26}$$

B Additional Figures and Tables

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

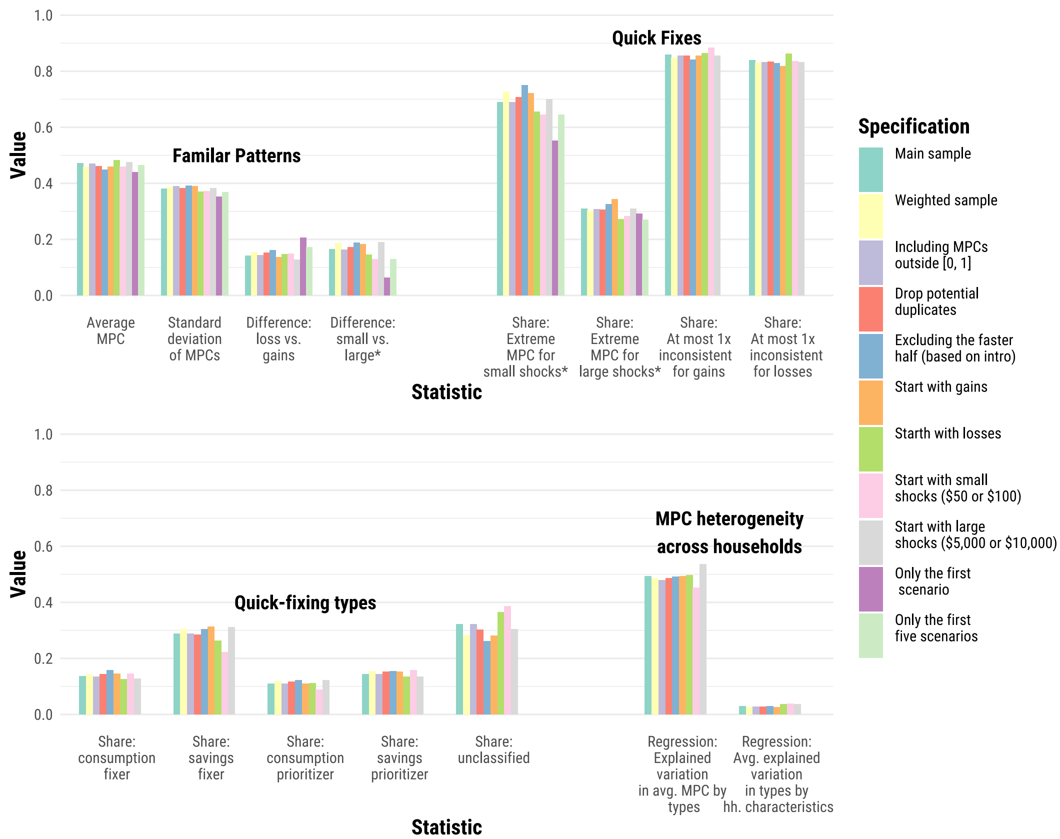

Notes: We recalculate key statistics from Section 3 for a variety of robustness specifications:

- *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.

Table B.1: Additional significance tests

Statement	***t*-test**	**Randomization test**	
	p-value	Statistic under H_0	p-value
MPCs decline for larger shocks; e.g., for gains, 0.56 for \$50 gain versus 0.30 for \$10,0000 gain. Difference: 0.26.	< 0.001	0.000	< 0.001
The average MPC is larger for losses (0.54) than for gains (0.40). Difference: 0.14.	< 0.001	0.000	< 0.001
From one shock to the next larger shock, a net share of 7% of households transition from an extreme MPC to a more moderate interior MPC.	–	0.000	< 0.001
Most households — namely 68% for gains and 67% for losses — switch to an interior MPC at most once and stick to interior MPCs thereafter. Average: 68%.	–	0.413	< 0.001
Only 14% (for gains) and 16% (for losses) of households deviate from this pattern more than once. Average: 15%.	–	0.398	< 0.001
Extreme MPCs more common for small than large shocks, e.g. 74% for \$50 gain versus 0.26% for \$10,000 gain. Difference: 0.48.	< 0.001	0.000	< 0.001
The average absolute difference between two interior MPCs of adjacent shock sizes is 0.14.	–	0.236	< 0.001
For only 3% of shock size increases, we observe a transition from interior to extreme MPCs.	–	0.126	< 0.001
Conditional on starting from an interior MPC, households stay in the interior for 93% of shock size increases.	–	0.753	< 0.001

Notes: This table reports whether the statistics reported in the statements are significantly different from the patterns we would expect to result from pure chance. The t-test column reports p-values from two-sided t-tests that test for differences in MPCs between shocks. We also use a more flexible randomization test that derives the distribution of the statistics under the null hypothesis that there is no link between MPCs, shock sizes, and valence. To achieve this, we permute the data within each household by (i) reshuffling MPCs within gains, and (ii) reshuffling MPCs within losses, and (iii) randomly replacing all gain MPCs with loss MPCs and vice versa. We draw 10,000 permuted data sets. The column "Statistic under H_0" reports the average value of the statistic in the permuted data sets. The last column reports the p-values of the randomization tests.

Figure B.2: Distribution of MPCs for losses before and after switching to interior

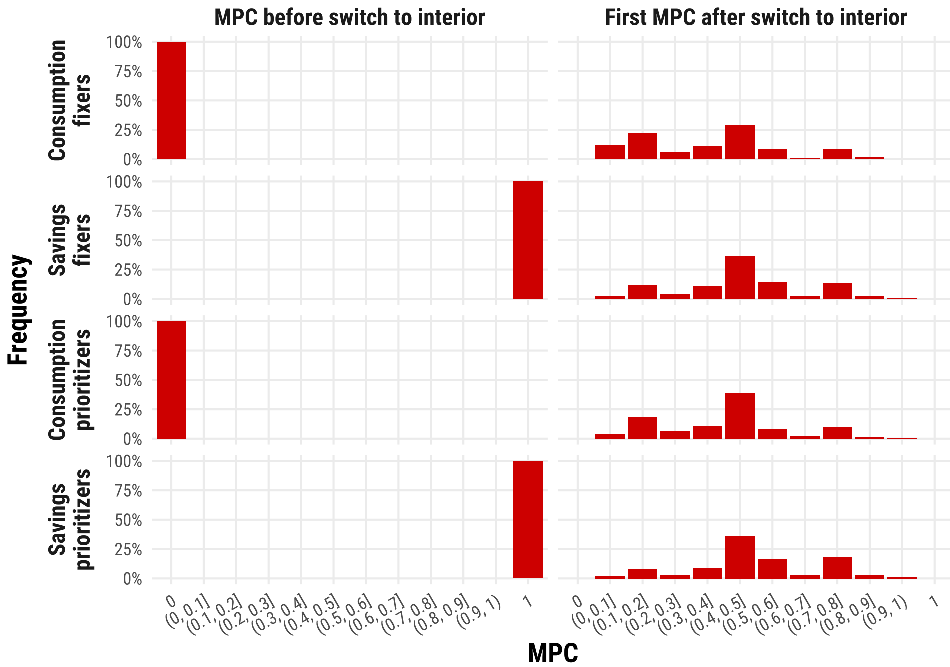

Notes: These histograms show conditional distributions of MPCs for losses. The rows correspond to the four quick-fixing types. The first column shows the distribution of MPCs before households switch to an interior MPC, which by construction puts all mass at either MPC = 0 or MPC = 1. The second column 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 3.

Table B.2: Exploring the variation in MPCs across households

	MPCs					
	Average MPC		**Share MPC = 0**		**Share MPC = 1**	
	(1)	(2)	(3)	(4)	(5)	(6)
Consumption and income						
Monthly spending (log.)		0.012***		−0.014***		0.008**
		(0.003)		(0.004)		(0.003)
Annual income (log.)		0.007		0.003		−0.002
		(0.005)		(0.006)		(0.006)
Income risk (std.)		0.020***		−0.038***		−0.008*
		(0.004)		(0.004)		(0.005)
Liquid wealth: dummies with reference group: [0k, 1k)						
[1k, 10k)		−0.052***		0.003		−0.075***
		(0.009)		(0.010)		(0.010)
[10k, 100k)		−0.069***		0.013		−0.097***
		(0.010)		(0.012)		(0.012)
[100k, more)		−0.072***		0.058***		−0.070***
		(0.013)		(0.015)		(0.014)
Illiquid wealth: dummies with reference group: [0k, 10k)						
[10k, 100k)		−0.018*		0.008		−0.026**
		(0.010)		(0.011)		(0.013)
[100k, 500k)		−0.025***		0.046***		0.008
		(0.010)		(0.011)		(0.012)
[500k, more)		−0.067***		0.102***		0.017
		(0.012)		(0.014)		(0.014)
Debt: dummies with reference group: [0k, 1k)						
[1k, 10k)		0.030***		−0.059***		−0.005
		(0.009)		(0.010)		(0.010)
[10k, 100k)		0.039***		−0.053***		0.027***
		(0.009)		(0.010)		(0.010)
[100k, more)		0.018*		−0.025**		0.018
		(0.010)		(0.012)		(0.011)
Other characteristics						
College		0.004		−0.023***		−0.014*
		(0.007)		(0.008)		(0.008)
Age (in 10y)		−0.015***		0.035***		0.006***
		(0.002)		(0.002)		(0.002)
Female respondent		0.009		−0.014**		0.005
		(0.006)		(0.007)		(0.007)
Household size		0.010***		−0.013***		0.002
		(0.003)		(0.003)		(0.003)
Quick-fixing types						
Consumption fixer	−0.346***		0.625***		−0.006	
	(0.007)		(0.010)		(0.005)	
Savings fixer	0.165***		0.093***		0.457***	
	(0.006)		(0.006)		(0.007)	
Consumption prioritizer	−0.094***		0.321***		0.184***	
	(0.008)		(0.010)		(0.007)	
Savings prioritizer	−0.073***		0.331***		0.236***	
	(0.006)		(0.008)		(0.007)	
Constant	0.492***	0.396***	0.050***	0.185***	0.057***	0.216***
	(0.003)	(0.050)	(0.003)	(0.057)	(0.003)	(0.059)
Obs.	4,981	4,981	4,981	4,981	4,981	4,981
R^2	0.492	0.109	0.563	0.211	0.563	0.028

Notes: This table reports regressions that explore the heterogeneity of households' MPCs. Columns 1–2 analyze households' mean MPC (averaged across all 14 shocks), Columns 3–4 analyze households' share of MPCs that equal 0 (among all 14 shocks), and Columns 5–6 analyze households' share of MPCs that equal 1 (among all 14 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.3: 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***	0.011**	0.005	0.000	−0.005		−0.091***
	(0.004)	(0.005)	(0.003)	(0.004)	(0.005)		(0.023)
Annual income (log.)	0.005	0.000	−0.009	−0.018**	0.022**		0.126***
	(0.007)	(0.010)	(0.007)	(0.008)	(0.010)		(0.044)
Income risk (std.)	−0.022***	−0.021***	−0.020***	−0.022***	0.084***		0.057*
	(0.005)	(0.007)	(0.004)	(0.005)	(0.007)		(0.031)
Liquid wealth: dummies with reference group: [0k, 1k)							
[1k, 10k)	0.012	−0.051***	0.003	−0.010	0.046**		−0.438***
	(0.013)	(0.019)	(0.013)	(0.015)	(0.018)		(0.082)
[10k, 100k)	0.020	−0.086***	0.005	−0.026	0.088***		−0.205**
	(0.015)	(0.020)	(0.015)	(0.017)	(0.019)		(0.087)
[100k, more)	0.070***	−0.088***	0.037**	−0.045**	0.026		0.052
	(0.020)	(0.026)	(0.019)	(0.020)	(0.023)		(0.104)
Illiquid wealth: dummies with reference group: [0k, 10k)							
[10k, 100k)	−0.002	−0.000	−0.003	0.015	−0.009		−0.053
	(0.015)	(0.022)	(0.014)	(0.017)	(0.021)		(0.100)
[100k, 500k)	0.001	0.050**	0.016	0.020	−0.086***		0.111
	(0.014)	(0.020)	(0.014)	(0.016)	(0.019)		(0.084)
[500k, more)	0.061***	0.082***	0.003	0.043**	−0.189***		0.141
	(0.019)	(0.024)	(0.017)	(0.019)	(0.021)		(0.098)
Debt: dummies with reference group: [0k, 1k)							
[1k, 10k)	−0.057***	−0.002	0.008	0.009	0.042**		−0.364***
	(0.013)	(0.018)	(0.012)	(0.014)	(0.018)		(0.083)
[10k, 100k)	−0.051***	0.046**	−0.000	0.017	−0.012		−0.269***
	(0.013)	(0.018)	(0.012)	(0.014)	(0.017)		(0.076)
[100k, more)	−0.039**	0.045**	0.034**	0.038**	−0.077***		−0.363***
	(0.016)	(0.020)	(0.015)	(0.016)	(0.018)		(0.080)
Other characteristics							
College	−0.011	−0.021	0.007	−0.019*	0.045***		−0.005
	(0.011)	(0.015)	(0.010)	(0.011)	(0.013)		(0.061)
Age (in 10y)	0.016***	0.010**	0.011***	0.007**	−0.045***		0.100***
	(0.003)	(0.004)	(0.003)	(0.003)	(0.004)		(0.018)
Female respondent	−0.009	0.015	−0.017*	−0.003	0.014		−0.011
	(0.010)	(0.013)	(0.009)	(0.010)	(0.012)		(0.056)
Household size	−0.013***	0.005	0.003	−0.007	0.012**		−0.029
	(0.004)	(0.005)	(0.004)	(0.004)	(0.005)		(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*	0.134	0.089	0.324***	0.315***	7.960***	6.987***
	(0.077)	(0.099)	(0.067)	(0.082)	(0.103)	(0.059)	(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. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Figure B.3: MPC profiles across the wealth distribution

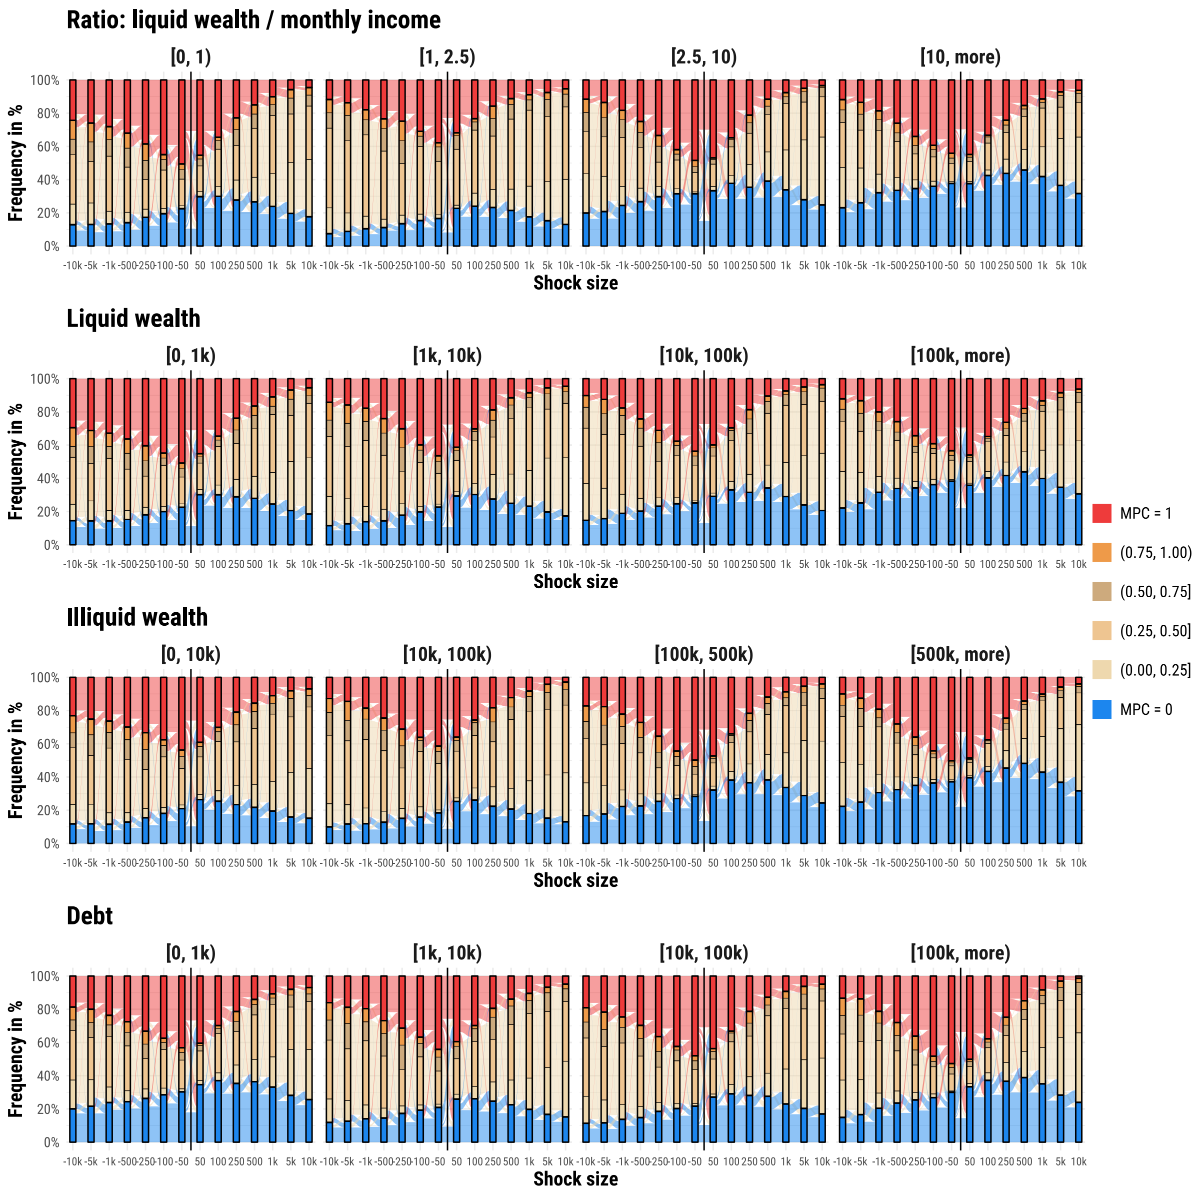

Notes: The alluvial graphs summarize the MPC profiles of households with varying (i) liquid wealth to income ratio (Kaplan, Violante and Weidner, 2014), (ii) liquid wealth, (iii) illiquid wealth, and (iv) debt (see Appendix C.3 for variable definitions). In each panel, each of the 14 columns displays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

Table B.4: 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***	−0.274***	−0.248***	−0.254***	−0.257***	−0.276***
	(0.011)	(0.020)	(0.011)	(0.020)	(0.013)	(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.5: 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***	−0.078***	−0.068***	−0.094***
	(0.001)	(0.002)	(0.002)	(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 repsondents make their first, second, third, ... decision. All regressions use household-level fixed effects. The robust standard errors are robust and clustered at the household level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Figure B.4: Quantitative model fit for rational agents' MPCs

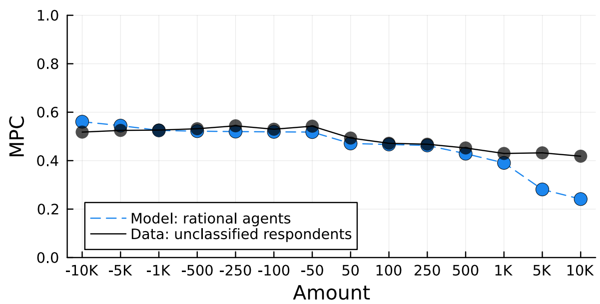

Notes: The line graphs compare empirical measurements and model predictions for the average marginal propensity to consume out of shocks of different amounts, among Unclassified respondents in the survey (black line and dots) and rational agents in the model (blue line and dots). We calibrate the discount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

Figure B.5: Quantitative model fit for the MPC distribution

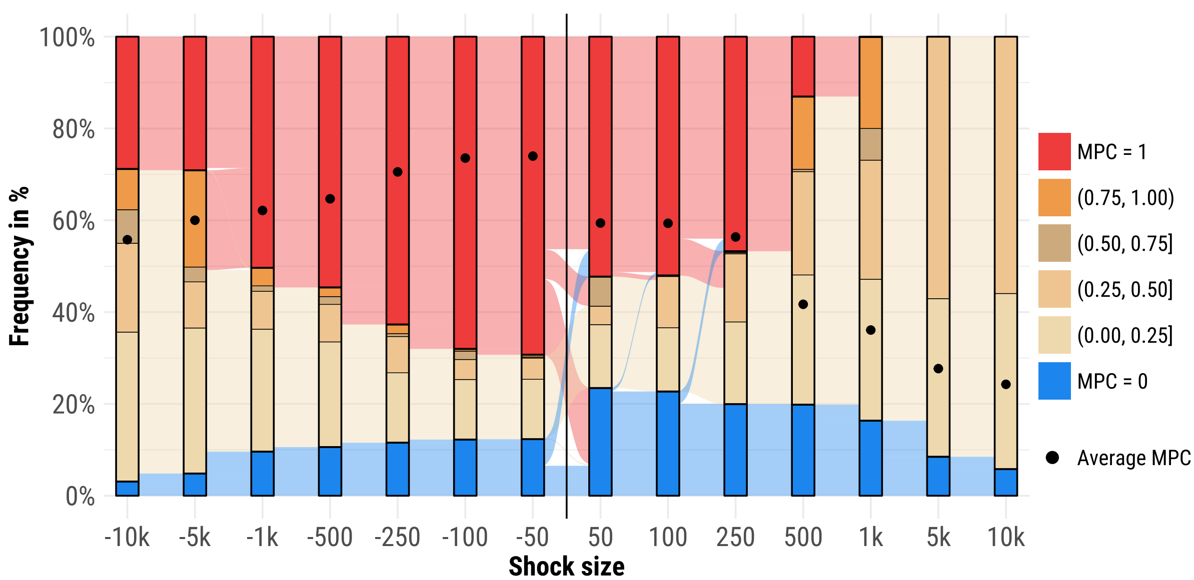

Notes: The alluvial graph summarizes the MPC predictions generated by the quantitative model, following the format of Figure 1. Each of the 14 columns displays the distribution of MPCs for one particular shock size, with colors indicating the size of the MPC. The streams between bars indicate how households' 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.

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 MPCs 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 Appendix 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 *de*creases in your household's spending/saving.

Increase in spending (By how much) would your household increase its monthly spending over the next three months?	$ 0
Increase in saving (By how much) would your household increase its monthly saving (which includes increases in debt repayment or <u>de</u>creases in debt-taking) over the next three months?	$ 0
Total	$ 0

Comment: Situation 2–7 are analogous. We randomize the order of shock sizes. Each respondents 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 *in*creases in your household's spending/saving.

Decrease in spending (By how much) would your household decrease its monthly spending over the next three months?	$ 0
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?	$ 0
Total	$ 0

Comment: Situation 2–7 are analogous. We randomize the order of shock sizes. Each respondents 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.

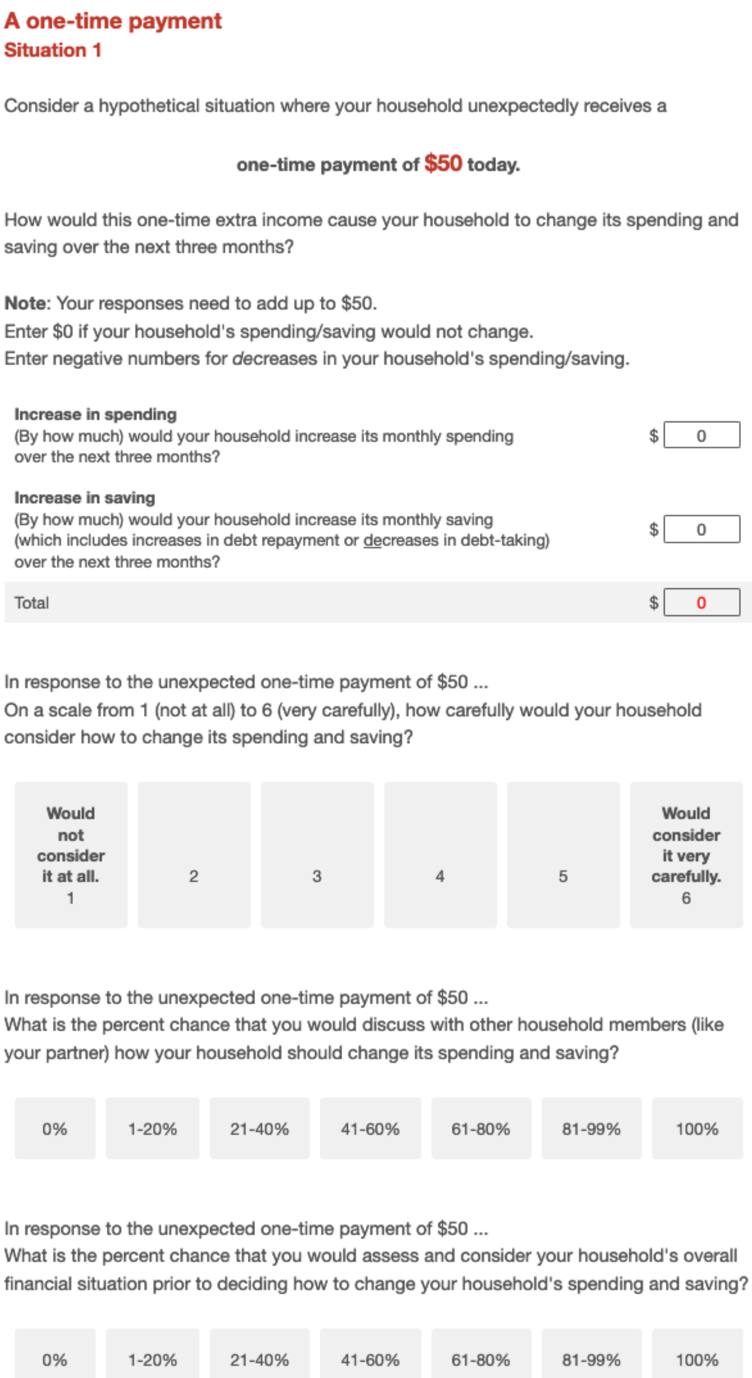

A one-time payment
Situation 1

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?	$ 0
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?	$ 0
Total	$ 0

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 response (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.

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 in Section 3.2 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 *notational 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 on 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 33% to 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, Jaravel and Spiess (2024) and Orchard, Ramey and Wieland (2024) estimate an MPC of 30%, correcting earlier higher estimates by Parker, Souleles, Johnson and McClelland (2013) and Broda and Parker (2014). Estimates for the consumption response to the 2020 Economic Stimulus Payment in the US range from 8–28% (Parker, Schild, Erhard and Johnson, 2022), to 25%-30% (Baker, Farrokhnia, Meyer, Pagel and Yannelis, 2023), or 40% (Coibion, Gorodnichenko and Weber, 2020). In a randomized experiment, Boehm, Fize and Jaravel (2024) 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, Holm and Natvik (2021) estimate a within-year MPC of around 0.50 out of lottery winnings.

MPCs decline for larger shocks. MPCs decline with larger shock size. This has been observed, *e.g.*, by Kueng (2018), Fagereng et al. (2021), and Colarieti et al. (2024). 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, Le Roux, Reinold and Surico (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; Misra and Surico, 2014; Boehm 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 is Fuster et al. (2021) who observe an extremely large share of MPCs of 0 (74% for a $500 gain) but few MPCs close to 1.[12]

Recent research has started to attempt inferring the distribution of MPCs in field settings. Misra and Surico (2014) estimate quantile consumption effects of the 2001 and 2008 US tax rebates and find that many households have an MPC close to 0, while a smaller group of households has MPCs close to 1. Karger and Rajan (2020) estimate consumption responses to the Covid stimulus payments. Their estimated individual-level distribution is noisy but exhibits spikes at 0 and 1 (Figure A11 and A12 in their paper). Lewis et al. (2024) estimate the latent distribution of MPCs in response to the 2008 US tax rebate, taking a parametric approach and assuming that there are three latent MPC types, which they estimate to have MPCs of 0.04, 0.23, and 1.33, respectively. However, their clustering approach is not designed to detect spikes in the distribution. Boehm et al. (2024) use a non-parametric approach to recover the distribution of MPC profiles from receiving an additional 300 Euro debit card in a randomized field. Although they do not detect exact mass points of zero and one, some smoothing of the density is inevitable due to measurement error and the kernel density deconvolution method that the authors employ to non-parametrically recover the CDF.

[12]If we had to speculate, we would attribute this difference to their two-stage response format. Respondents are first asked whether they would increase, not change, or decrease their spending, debt repayment, or savings. Only then can households specify their precise responses. This means response noise overproportionally favors MPCs of 0 and below 0 (8% of respondents select an MPC below 0 for a $500 gain), while it is hard to indicate an MPC of 1.

C.5 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)*

Table C.3: Overview of the coding scheme

Theme (and detected freq.)	**Description**
Gains	
199 cases where respondents choose an extreme MPC for small shock but not for large shock.	
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	
184 cases where respondents choose an extreme MPC for small shock but not for large shock.	
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 they prefer to not reduce spending any further in response to a \$100 loss.

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 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

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," Working Paper 2024.

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess, "Revisiting Event-Study Designs: Robust and Efficient Estimation," *Review of Economic Studies*, 2024. Forthcoming.

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," *J. 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 J.: Macroeconomics*, 2021, *13* (4), 1–54.

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, Diana Farrell, Fiona Greig, and Chris Wheat, "Wealth, Race, and Consumption Smoothing of Typical Labor Income Shocks," 2023.

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, (1), 77–153.

Karger, Ezra and Aastha Rajan, "Heterogeneity in the Marginal Propensity to Consume: Evidence from Covid-19 Stimulus Pay," Working Paper, FRB of Chicago 2020.

Kueng, Lorenz, "Excess Sensitivity of High-Income Consumers," *The Quarterly Journal of Economics*, November 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.

Orchard, Jacob, Valerie A Ramey, and Johannes Wieland, "Micro MPCs and Macro Counterfactuals," *The Quarterly Journal of Economics*, 2024. Forthcoming.

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.