Heterogeneity in MPCs and Unexplained Variation in Consumption Expenditures*

Filip Rozsypal† Kathrin Schlafmann‡

December 4, 2025

Abstract

Quantile regression is a popular method to estimate the dispersion in MPCs in the population. We discuss the challenges that this method faces given the large unexplained variation in consumption expenditures in survey data. We highlight that quantile regression estimates do not recover the distribution of MPCs if either unexplained variation is due to measurement error or if differences in MPCs are partly driven by ex ante heterogeneity across households. To quantify the likely extent of the bias, we propose a simulation-based approach where we back out the underlying distribution of MPCs for a range of calibrations that attribute the unexplained variation to a split between unobserved factors and measurement noise. All results point in the same direction: the true distribution of MPCs is significantly more dispersed than what is estimated by quantile regression.

JEL codes: D12, E21
KEYWORDS: Marginal Propensity to Consume, heterogeneity, measurement error

*Kathrin Schlafmann is affiliated with the Danish Finance Institute (DFI) which financially supports researchers in Finance.
†Danmarks Nationalbank, email: firo@nationalbanken.dk
‡Copenhagen Business School, Danish Finance Institute (DFI) and Centre for Economic Policy Research (CEPR), email: ksc.fi@cbs.dk

1 Introduction

The Marginal propensity to consume (MPC) is a key characteristic that captures households' responsiveness to shocks and hence affects the effectiveness of both monetary (Kaplan et al., 2018; Auclert, 2019) and fiscal (Jappelli and Pistaferri, 2014; Kaplan and Violante, 2014; Rozsypal and Schlafmann, 2021; Auclert et al., 2024) policy. Because MPCs matter beyond their average value in the population, empirical studies have been trying to uncover the distribution of MPCs using competing approaches, one of which is quantile regression. We contribute to this effort by highlighting the challenges that quantile regression faces when used to uncover the heterogeneity in MPCs based on survey data that exhibits large amounts of unexplained variation in expenditures. Using these insights we propose a simulation-based approach to back out the underlying distribution of MPCs that is necessary to generate the QR estimates estimated from consumption surveys. Applied to expenditure data from the Consumer Expenditure Survey (CEX) for the 2008 stimulus payment, we find that the distribution of MPCs is much more bimodal than what the QR estimates suggest.

One way of uncovering the heterogeneity in MPCs is to compute the average MPC within different groups and then show how these averages differ across groups.[1] A disadvantage of this approach is that within every bin, the heterogeneity in MPCs is compressed and instead it is represented by its bin average. A most recent refinement of this approach is Lewis et al. (2025), who use Gaussian mixture linear regression to jointly characterize the groupings and estimate the MPC, which in principle allows the groups to depend on unobserved factors, thereby reducing the within bin heterogeneity. As an alternative, Misra and Surico (2014) propose to use quantile regression (QR) to uncover the heterogeneity in MPCs.[2] This approach relies on identifying the distribution of MPC from the distribution of the marginal effect of the transfer at different quantiles of the conditional distribution of consumption change.

In this paper, we revisit the use of QR in estimating MPCs in the context of the large unexplained variation that is present when empirically estimating drivers of consumption expenditures. For example, Souleles et al. (2006) report an R^2 of less than 1 percent in their analysis of the 2001 stimulus payment, meaning more than 99% of the variation in consumption expenditures is not explained by observables. We highlight that this unexplained variation causes two challenges for using QR to estimate the distribution of MPCs: First, if part of this variation is due to measurement error, we know from recent theoretical advances that QR estimates are biased (Hausman et al., 2021). Second, we show that even if the variation is not due to measurement error, QR estimates do not map into the distribution of MPCs as soon as there is ex ante heterogeneity in MPCs, i.e. heterogeneity that is not driven by the period-specific consumption shifter.

The result of both types of challenges is that the dispersion of the QR estimates is likely a lower bound for the true dispersion in MPCs. To gauge the degree of this bias in practice, we implement a simulation-based procedure that uses the unbiased OLS estimate of the average MPC as an anchor to back out the underlying distribution of MPCs that, if estimated following the standard QR approach, would deliver the distribution that we estimate on data from the CEX. The resulting MPC distribution is much more extreme with a significant number of households consuming either nothing or 100% out of the transfer. Having said that, the purpose of this paper is not the take a strong stand on how to attribute the unexplained variation in consumption growth. Instead, we invite the reader to pick their preferred split and we provide the resulting corrected MPC distribution. Regardless of the exact split, however, all our results point in the

[1] For example, see Fagereng et al. (2021); Patterson (2023)

[2] Estimation of MPCs by QR can also be combined with separating households into groups (see Miranda-Pinto et al. (2025)). Yet another approach is to use structural estimation, such as Carroll et al. (2017) or using statistical deconvolution techniques as in Boehm et al. (2025).

same direction: The distribution in MPCs is more extreme than what QR estimates indicate. In fact, the resulting distributions are more in line with evidence of self-reported MPCs (Jappelli and Pistaferri, 2014; Bunn et al., 2018), who document a significant fraction of households having a MPC of either zero or one.

2 Unexplained Variance in Consumption and MPC Estimation

Kaplan and Violante (2022) define the MPC as "*the fraction of a small, unanticipated one-time windfall that a household spends within a given time period.*". The traditional way to estimate MPCs (Souleles et al., 2006; Parker et al., 2013) is by using ordinary least squares (OLS) regressions

$$\Delta C_{i,t} = \boldsymbol{\gamma}_0 \boldsymbol{X}_{it} + \gamma_1 R_{it} + u_{it}, \tag{1}$$

where $\boldsymbol{X}_{it}$ includes an intercept, household controls as well as time fixed effects and R_{it} is a transfer, e.g. a tax rebate. Under reasonable assumptions about the exogeneity of the transfer, the OLS estimate of γ_1 correctly identifies the average treatment effect, which reflects the average MPC. This is true even though equation (1) in practice only explains a relatively small portion of the variation in $\Delta C_{i,t}$. For example, Souleles et al. (2006) report an R^2 of less than 1 percent, meaning there is a large amount of unexplained variation in consumption expenditures.

Due to the focus on heterogeneity in macroeconomics over the last two decades, there has been increasing interest in the whole distribution of MPCs. Misra and Surico (2014) proposed to use QR as means to go beyond estimating the average effect and to recover the full distribution of MPCs. In this article, we argue that the large unexplained variation in consumption expenditures poses a challenge for the interpretation of QR estimates as the distribution of MPCs and propose a procedure to infer the latter from QR estimates.

Koenker (2005) shows that QR recovers the marginal effect at a given quantile τ of the conditional distribution of the dependent variable, in our case of consumption growth:

$$Q_\tau(\Delta C_{it} | R_{it}, \boldsymbol{X}_{it}) \ = \beta_0(\tau)\boldsymbol{X}_{it} + \beta_1(\tau) R_{it}. \tag{2}$$

The QR estimate of $\beta_1(\tau)$ is the effect of the rebate at the τ-th quantile of the dependent variable, here $\Delta C_{it} | R_{it}, \boldsymbol{X}_{it}$. To see under which conditions $\beta_1(\tau)$ maps into the quantile of the MPC distribution, we rewrite the QR model as a special case of a random-coefficients model to estimate heterogeneous treatment effects (Koenker, 2005, section 2.6):

$$\Delta C_{it} \ = \ \beta_0(u_{it})\boldsymbol{X}_{it} + \beta_1(u_{it}) R_{it}, \tag{3}$$

where $u_{it} \sim U[0,1]$ represents the unobserved quantile of $\Delta C_{it} | R_{it}, \boldsymbol{X}_{it}$. This formulation emphasizes that QR assumes that all unobserved heterogeneity enters through the treatment effects as one-dimensional heterogeneity in form of the scalar u_{it}. We now describe two reasons why the presence of unexplained variation in consumption expenditures implies that the estimates from this model in practice might not map into the distribution of MPCs, our object of interest.

Challenge 1: Measurement Error in Consumption Expenditures Unexplained variation in expenditures can be due to measurement noise in consumption surveys. While not problematic in the context of OLS and hence for the estimates of *the average* MPC, Hausman et al. (2021) show that quantile regressions—unlike OLS—are biased in the presence of measurement error in the dependent variable. Specifically, Hausman et al. (2021) show that measurement error in the dependent variable leads to a "compression bias", i.e. that all estimated quantiles are strictly between the true minimum and maximum of the marginal effects.

QR's compression bias is present in any setting with measurement error, but is particularly important in the context of estimating MPCs since there is evidence that a large fraction of the unexplained variation in consumption expenditures is due to measurement error. For example, comparisons of the variance and autocovariance of consumption growth in Blundell et al. (2008) suggest that noise explains about one third to one half of the variance in consumption growth. Carroll (1997, Table 4, page 19) uses 1/2 and 1/3 as plausible guesses for signal to noise ratio in CEX data. Using Canadian Food Expenditure Survey, Ahmed et al. (2006) document that measurement error can explain up to 70% of the variance in log food expenditures.[3] All these studies point to measurement error of substantial magnitude in the range of 30-70% of the overall variance of consumption growth.

Challenge 2: Sources of Heterogeneity in MPCs For the part of unexplained variation that is not due to measurement error, the implications for QR are more nuanced. The reason is that unexplained drivers of the change in consumption might also affect the response of a household to receiving a transfer. In other words, the MPC might be a function of these unobserved drivers. To see how this impacts the interpretation of the QR estimates, consider the following reduced-form setting. Let the change in consumption expenditures in the absence of the transfer (conditional on all observables) be described by a "consumption shifter" z_{it} (e.g. an unexpected health shock leading to a large medical bill):

$$\Delta C^*_{it} \;=\; z_{it}. \tag{4}$$

Turning to the spending out of the transfer, let the MPC be a function of two components: unobserved ex ante heterogeneity across households x_i and a state-dependent component that is reflected in the consumption shifter z_{it}. The former can be understood as underlying differences across households in, e.g., time preferences or risk aversion, as well as persistent states that do not vary during the relatively short amount of time that a household remains in a consumption survey. For example, liquid assets are only reported once per household in the CEX and hence do not vary in each period. Heterogeneity in MPCs related to liquid assets will therefore be part of ex ante heterogeneity. The latter captures the period-specific situation a household is in: In periods where households for reasons unexplained by observables have unusually large or small consumption expenditures, it is plausible that they will also be likely to spend a larger or smaller fraction of the rebate. For instance, in the medical bill example, a household who has to use a lot of their liquid assets for the medical expenditures might be likely to spend more of the transfer to keep their regular consumption expenses at their usual level.[4] In this setup, (conditional) consumption growth in the presence of the transfer is

$$\Delta C_{it} \;= z_{it} + MPC(x_i, z_{it}) R_{it}. \tag{5}$$

Equation (5) simplifies to equation (3) only if either, there are no changes in consumption in the absence of a rebate (i.e. no consumption shifters z_{it}), or there is no ex ante heterogeneity in MPCs across households (no heterogeneity in x_i). The former is unlikely given the large unexplained variation in consumption

[3]See Crossley and Winter (2013); Brzozowski et al. (2017) for more details and results about measurement errors in consumption surveys.

[4]For a structural model where the presence of consumption shocks has implications for MPCs see Miranda-Pinto et al. (2025).

expenditures. In the latter case, u_{it} in equation (3) reflects the quantile in the consumption shifter distribution (which is the only remaining driver of MPC). The quantile-specific intercept will then trace out the distribution of the consumption shifter while $\hat{\beta}_1(\tau)$ traces out the distribution of MPCs.

However, as soon as the MPC depends on *both* ex ante heterogeneity across households *and* situational consumption shifters, QR estimates no longer recover the underlying MPC distribution. Formally, that is because there are two sources of heterogeneity in MPCs while the QR regression in equation (3) assumes that there is only one. Intuitively, QR estimates reflect the marginal effect of the transfer at different quantiles of the conditional distribution of consumption change, but consumption change is not only driven by the MPC. Take observations with high (ex ante) MPCs. These observations can still end up with low consumption growth due to a low realization of the consumption shifter. In this case these observations will increase the estimated marginal effect at lower quantiles of consumption growth and hence the QR estimate at these lower quantiles. The reverse is true for observations with low (ex ante) MPC. Hence—similarly to the compression bias due to measurement error—QR estimates will be less heterogeneous than the distribution of MPCs if there is ex ante heterogeneity in MPCs.

3 Inferring the Distribution of MPCs

We now use the insights from the previous section to back out the degree of heterogeneity in MPCs that must be present such that the QR estimates in equation (2) match the empirically estimated quantiles from the CEX for the 2008 tax rebate. The fundamental idea is that unexplained variation in consumption data make the estimated QR quantiles flatter, i.e., less heterogeneous, compared to the underlying distribution of MPCs. We therefore use Monte Carlo simulations to back out how steep the underlying distribution of MPCs needs to be in order for it to generate the empirically observed QR estimates. Additionally, we also use the unbiased OLS estimate for the average MPC as an additional moment that restricts the set of possible MPC distributions.

Simulation Approach to Recover the Distribution of MPCs We extend equation (4) to allow for the presence of measurement error in consumption expenditures

$$\Delta C^*_{it} \;=\; z_{it} + \epsilon_{it}, \tag{6}$$

where $z_{it} \sim N(0, \sigma_z)$, and $\epsilon_{it} \sim N(0, \sigma_\epsilon)$. For the MPC, we assume a linear functional form within the bounds of $[0, 1]$ for the MPC[5]:

$$MPC(x_i, z_{it}) \;=\; \begin{cases} 0 & \text{if } \; \alpha_0 + \alpha_1 x_i + \alpha_2 z_{it} < 0 \\ 1 & \text{if } \; \alpha_0 + \alpha_1 x_i + \alpha_2 z_{it} > 1 \\ \alpha_0 + \alpha_1 x_i + \alpha_2 z_{it} & \text{otherwise} \end{cases}, \tag{7}$$

where $z_{it} \sim N(0, \sigma_z)$, $x_i \sim N(0, 1)$, and $\epsilon_{it} \sim N(0, \sigma_\epsilon)$. We also use the notation $MPC(x_i, 0)$ as describing the ex ante heterogeneity, i.e. the part of the MPC that is not driven by the consumption shifters. The data generating process for observed consumption changes is therefore described by

$$\Delta C_{it} \;=\; z_{it} + MPC(x_i, z_{it}) R_{it} + \epsilon_{it}. \tag{8}$$

[5]We bound the MPC between 0 and 1 in line with the definition from Kaplan and Violante (2022) quoted in the beginning of Section 2.

The two challenges highlighted in the previous section directly map into parametric restrictions for this data generating process: First, the extent of measurement error is governed by the share of unexplained variance in consumption growth that is not due to the consumption shifter:

$$\text{Degree of measurement error} \;=\; \frac{\sigma_\epsilon^2}{\sigma_\epsilon^2 + \sigma_z^2}. \tag{9}$$

Second, the importance of ex ante heterogeneity in overall heterogeneity of MPCs is governed by the share of the variance in MPCs that is driven by ex ante heterogeneity:

$$\text{Degree of ex ante heterogeneity} \;=\; \frac{\mathbb{V}\left[MPC(x_i, 0)\right]}{\mathbb{V}\left[MPC(x_i, z_i)\right]}. \tag{10}$$

We apply the following approach: In a first step, we empirically estimate the OLS regression in equation (1) on the CEX data for the 2008 stimulus payment.[6] This gives us an estimate of the amount of overall unexplained variation in consumption change as $\sigma_\epsilon^2 + \sigma_z^2 = \hat{u}_{it}^2$. Moreover, we obtain the estimate of the average MPC as $\hat{\beta}_1$. Recall that OLS does not suffer from either of the two challenges. Any underlying distribution of MPCs must therefore be consistent with an average MPC as estimated by the OLS regression.[7] In a second step, we fix a given set of assumptions for the degree of ex ante heterogeneity and measurement error. For a given degree of measurement error and our estimate for overall unexplained variation, equation (9) determines the variation in consumption shifters σ_z. Any candidate distribution, characterized by the parameters α_0, α_1, and α_2, therefore needs to satisfy: (1) the average MPC is equal to the OLS estimate of the average MPC, and (2) the degree of ex ante heterogeneity according to equation (10) is equal to the assumed degree. Finally, in the last step, we run a Monte Carlo simulation using each candidate distribution, where in each repetition we run the QR estimation. We then determine the underlying MPC distribution as the candidate distribution where the average of QR estimates across all repetitions is closest to the QR estimate from the CEX.[8]

Implied Distribution of MPCs What are plausible degrees of measurement error and ex ante heterogeneity? For the degree of measurement error, we have argued in Section 2 that previous studies found 30-70% of pure measurement error. However, in our context, true measurement error is indistinguishable from unexplained variation in consumption change that is uncorrelated with the MPCs. The relevant amount of measurement error—including uncorrelated unexplained variation—is hence likely even higher. We therefore focus our discussion on degrees of measurement error of at least 50%. As for the degree of ex ante heterogeneity, it is important to take into account that this source of heterogeneity not only comprises differences across people in their preferences. Instead, it also contains differences in persistent states households are in (e.g. the level of liquid assets a household has). This is because in consumption surveys such as the CEX we typically observe a household only for a few quarters. The estimation is thus unable to distinguish between heterogeneity due to preferences and heterogeneity due to a household's situation as long as this situation does not change during the observation period. Previous studies of average MPCs by subgroups

[6] We use the CEX data provided by Misra and Surico (2014). The analysis in this paper focuses on consumption expenditures according to their definition of "all nondurables" and as in their paper the control variables $\boldsymbol{X}_{it}$ include a quadratic function of age, of the change in the number of adults, and of the change in the number of children, as well as time fixed effects.

[7] The OLS estimate of the average MPC is equal to 13 percent, consistent with the findings in (Parker et al., 2013) and Misra and Surico (2014).

[8] For each Monte Carlo repetition we use a sample size equal to the sample size in the CEX (17718 observation, 3076 obtain a transfer) and a transfer size R of 500 US\$. We run 100 Monte Carlo repetitions and compute the average over these repetitions for each estimated quantile. The criterion to determine the best fit is the sum of squared deviations of these averages from the corresponding CEX QR estimate in the range of quantiles 0.1 to 0.9.

Figure 1: Implied distribution of MPCs

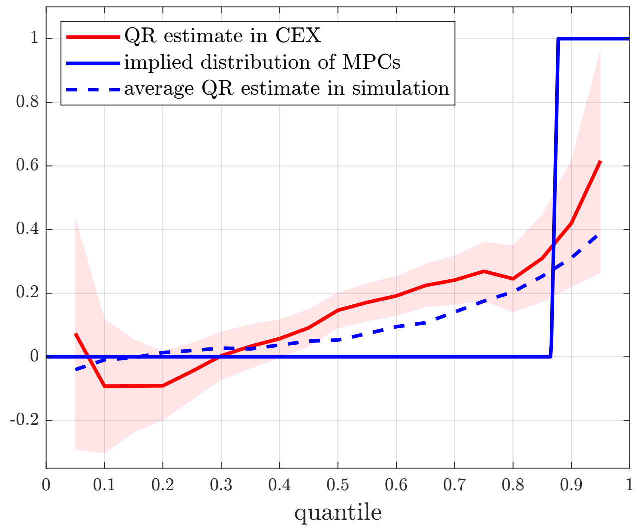

Note: The figure shows the full implied distribution of MPCs under the assumption of 30% of variance in MPCs is due to ex ante heterogeneity and 50% measurement error in consumption expenditures. It depicts the QR estimates based on CEX data of the 2008 tax rebate (red line, shaded areas refer to 90% confidence interval), the implied true distribution of MPCs (solid blue line), and the QR estimate in this setup (average over 100 Monte Carlo repetitions, dashed blue line).

have found substantially heterogeneity in average MPCs (Fagereng et al., 2021; Patterson, 2023). In our preferred specification we therefore analyze the effects of a degree of ex ante heterogeneity of 30%.

Figure 1 shows the full implied distributions of MPCs for our preferred specification. The figure plots the CEX QR estimates in red (with 90 percent confidence interval), as well as the true underlying distribution of MPCs as the solid blue line. The dashed blue line depicts the average QR estimates from the Monte Carlo simulations. As expected from our discussion of the challenges, ex ante heterogeneity in MPCs and measurement error together lead to a compression bias that substantially flatten the QR estimates relative to the true underlying distribution of MPCs. The underlying distribution of MPCs has to be very dispersed to generate the observed QR estimates: underlying MPCs are almost bimodal where almost all households either have an MPC of zero (86 percent) or an MPC of one (11 percent).

Figure 2 extends this analysis to show a systematic overview over the implied heterogeneity across a wide range of assumptions regarding the degree of ex ante heterogeneity and measurement error.[9] The three panels depicts the fraction of households with MPC $= 0$, with $0 <$ MPC < 1, and MPC $= 1$, respectively. Each panel plots the implied fraction against the share of assumed measurement error, and the three lines correspond to increasing degrees of ex ante heterogeneity in MPCs. The figure highlights that both the fraction of households that spend nothing out of the transfer and the fraction that spends the whole transfer increases with both challenges. This is because both challenges lead to flatter QR estimates relative to the underlying distribution of MPCs. So in order for the QR estimates to be close to the ones observed in the CEX, the underlying heterogeneity in MPCs needs to increase the more measurement error there is and/or the more ex ante heterogeneity there exists in MPCs.

The quantitative results in the figure show that for degrees of measurement error of more than 50%, almost all households either do not spend anything out of the transfer or spend all of it. If in addition to measurement error we think that 15-30% of the heterogeneity in MPCs is driven by ex ante differences across

[9] In the appendix, we also provide plots of the full distributions the moments displayed in figure 2 are based on.

Figure 2: Moments of the MPC distribution in the presence of ex ante heterogeneity and measurement error

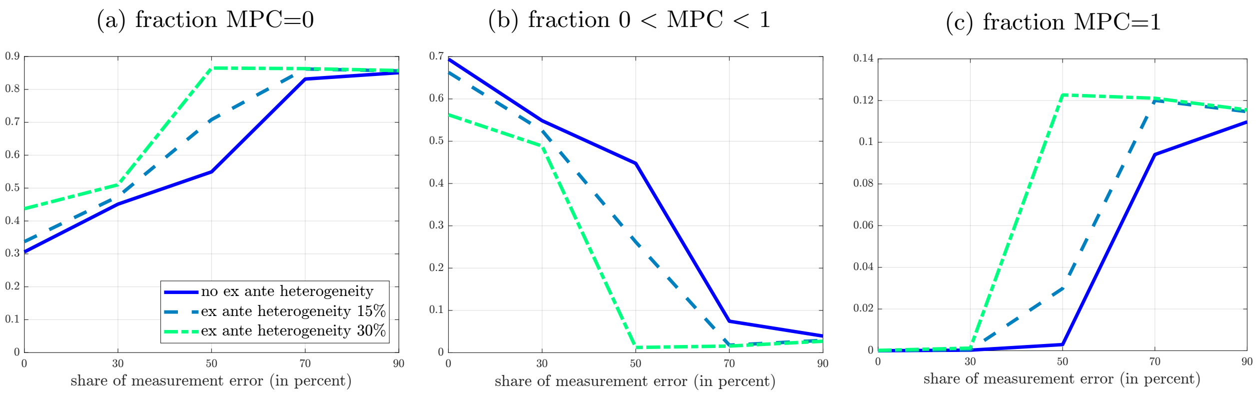

Note: The figure shows the implied fraction of households with MPC = 0 (figure (a)), with 0 < MPC < 1 (figure (b)), and with MPC = 1 (figure (c)). The implied fractions are plotted against the assumed share of measurement error in the total unexplained variance of consumption expenditures. The three lines refer to different assumptions about the share of variance in MPCs driven by ex ante heterogeneity: 0%, i.e. no ex ante heterogeneity (solid dark blue), 15% (dashed light blue), and 30% (dash-dotted green).

households (which includes differences due to persistent states of the household), then this bimodality of MPCs arises for even lower levels of measurement error. We therefore conclude that the implied heterogeneity in the distribution of MPCs that we can infer from the QR estimates of the CEX is much higher than what the QR estimates indicate.

4 Conclusion

Estimating the distribution of MPCs by quantile regression (QR) is challenging in the presence of the large unexplained variation in consumption expenditures found in consumption surveys. QR estimates will underestimate the true heterogeneity of MPCs if this unexplained variation is (at least partly) due to measurement error or if part of the heterogeneity in MPCs is due ex ante heterogeneity across households.

We have used these insights to back out the implied heterogeneity in MPCs that is necessary to generate empirically observed QR estimates. The procedure we propose uses Monte-Carlo simulation while restricting the set of possible MPC distributions using the unbiased OLS estimate of the average MPC. We found that for plausible ranges of ex ante heterogeneity and measurement error, the implied distribution of MPCs is almost bimodal: the majority of households (around 85 percent) does not spend anything of the transfer, while another substantial fraction (around 10 percent) spends the whole transfer. Our results are qualitatively consistent with recent studies on CEX data that find more heterogeneity in MPCs when estimated using alternative estimation approaches (Lewis et al., 2025). Quantitatively, however, our results point to a more bimodal distribution with bunching at 0 and 1which is more in line with evidence from self-reported MPCs (Jappelli and Pistaferri, 2014; Bunn et al., 2018).

References

Ahmed, N., Brzozowski, M., and Crossley, T. F. (2006). Measurement errors in recall food consumption data. Technical report, IFS WP 0621.

Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review*, 109(6):2333–67.

Auclert, A., Rognlie, M., and Straub, L. (2024). The intertemporal keynesian cross. *Journal of Political Economy*, 132(12):4068–4121.

Blundell, R., Pistaferri, L., and Preston, I. (2008). Consumption Inequality and Partial Insurance. *American Economic Review*, 98(5):1887–1921.

Boehm, J., Fize, E., and Jaravel, X. (2025). Five facts about mpcs: Evidence from a randomized experiment. *American Economic Review*, 115(1):1–42.

Brzozowski, M., Crossley, T. F., and Winter, J. K. (2017). A comparison of recall and diary food expenditure data. *Food Policy*, 72:53–61. Food counts. Measuring food consumption and expenditures in household consumption and expenditure surveys (HCES).

Bunn, P., Le Roux, J., Reinold, K., and Surico, P. (2018). The consumption response to positive and negative income shocks. *Journal of Monetary Economics*, 96:1–15.

Carroll, C., Slacalek, J., Tokuoka, K., and White, M. N. (2017). The distribution of wealth and the marginal propensity to consume. *Quantitative Economics*, 8(3):977–1020.

Carroll, C. D. (1997). Death to the log-linearized consumption euler equation! (and very poor health to the second-order approximation). Working Paper 6298, National Bureau of Economic Research.

Crossley, T. F. and Winter, J. K. (2013). Asking households about expenditures:what have we learned? Technical report, NBER WP 19543.

Fagereng, A., Holm, M. B., and Natvik, G. J. (2021). Mpc heterogeneity and household balance sheets. *American Economic Journal: Macroeconomics*.

Hausman, J., Liu, H., Luo, Y., and Palmer, C. (2021). Errors in the dependent variable of quantile regression models. *Econometrica*, 89(2):849–873.

Jappelli, T. and Pistaferri, L. (2014). Fiscal Policy and MPC Heterogeneity. *American Economic Journal: Macroeconomics*, 6(4):107–136.

Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary policy according to hank. *American Economic Review*, 108(3):697–743.

Kaplan, G. and Violante, G. L. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239.

Kaplan, G. and Violante, G. L. (2022). The marginal propensity to consume in heterogeneous agent models. *Annual Review of Economics*, 14(Volume 14, 2022):747–775.

Koenker, R. (2005). *Quantile Regression*. Cambridge University Press.

Lewis, D., Melcangi, D., and Pilossoph), L. (2025). Latent heterogeneity in the marginal propensity to consume. Technical report, wp.

Miranda-Pinto, J., Murphy, D., Walsh, K. J., and Young, E. R. (2025). A model of expenditure shocks. *Journal of Monetary Economics*, 154:103807.

Misra, K. and Surico, P. (2014). Consumption, income changes and heterogeneity: Evidence from two fiscal stimulus programmes. *American Economic Journal: Macroeconomics*, 6:84–106.

Parker, J. A., Souleles, N. S., Johnson, D. S., and McClelland, R. (2013). Consumer Spending and the Economic Stimulus Payments of 2008. *American Economic Review*, 103(6):2530–53.

Patterson, C. (2023). The matching multiplier and the amplification of recessions. *American Economic Review*, 113(4):982–1012.

Rozsypal, F. and Schlafmann, K. (2021). Overpersistence bias in individual income expectations and its aggregate implications. *American Economic Journal: Macroeconomics*, 15(4):331–71.

Souleles, N. S., Parker, J. A., and Johnson, D. S. (2006). Household Expenditure and the Income Tax Rebates of 2001. *American Economic Review*, 96(5):1589–1610.

A Appendix

Figure 3: Moments of the MPC distribution in the presence of ex ante heterogeneity and measurement error

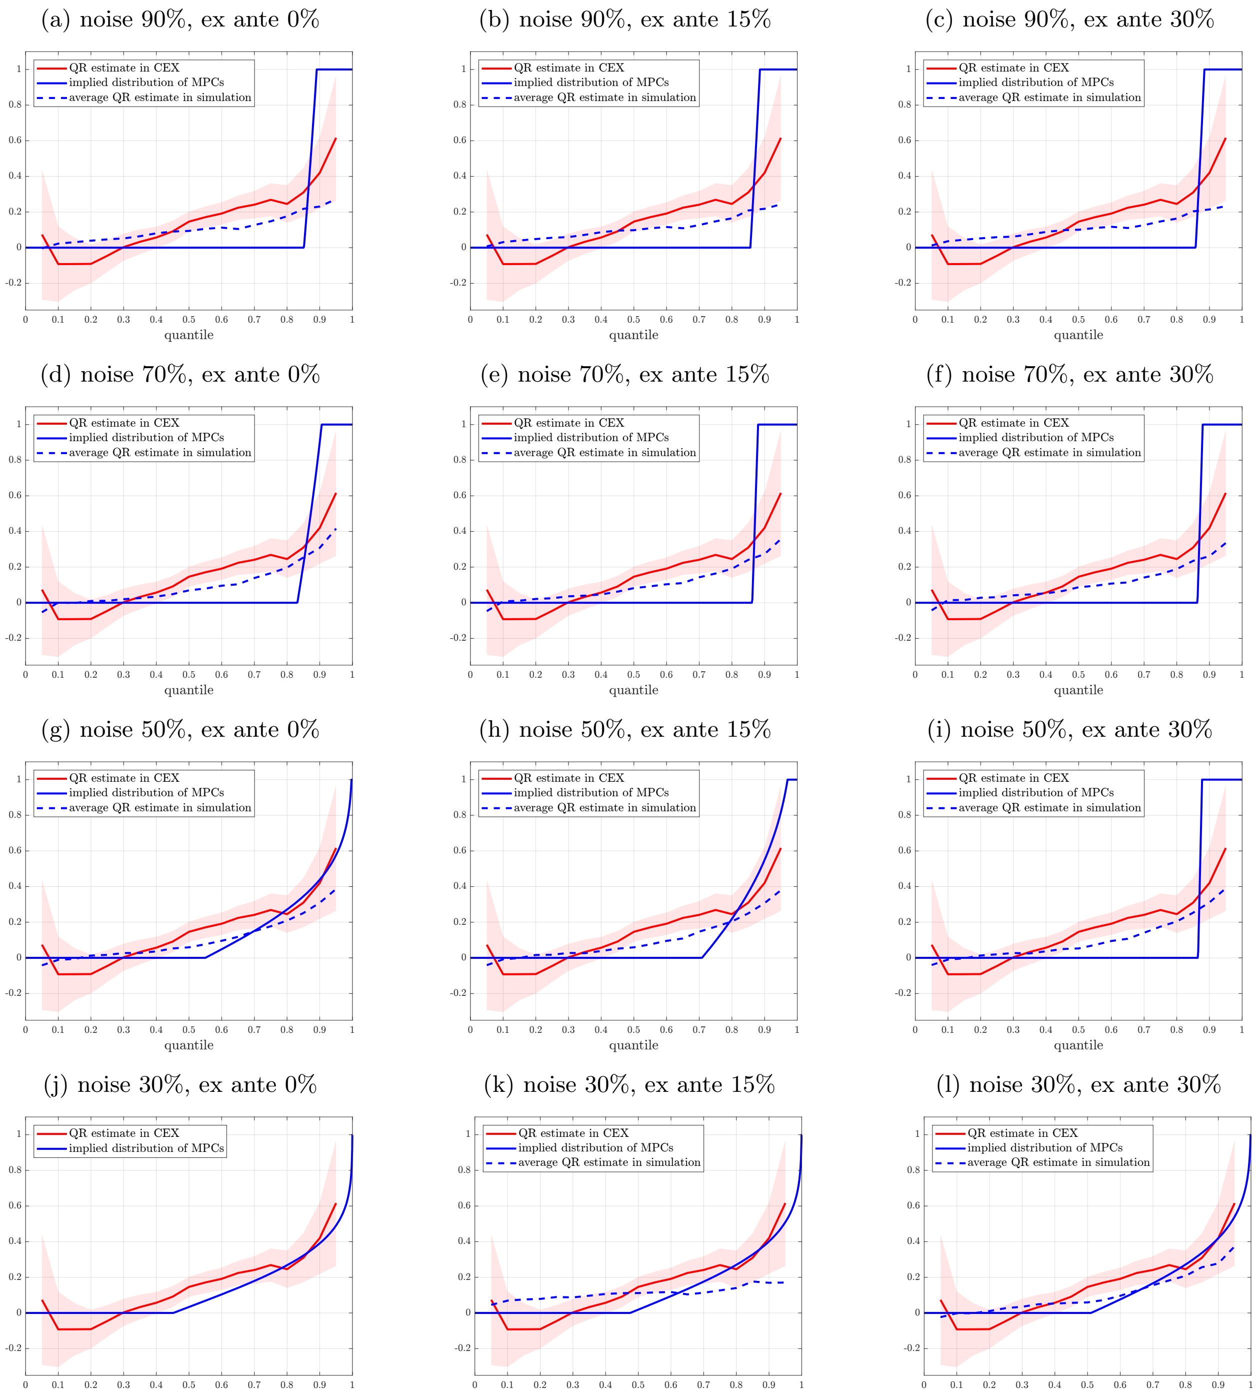

Note: The figure shows the full implied distribution of MPCs under varying assumptions for the degree of ex ante heterogeneity and the degree of measurement error. It depicts the QR estimates based on CEX data of the 2008 tax rebate (red line, shaded areas refer to 90% confidence interval), the implied true distribution of MPCs (solid blue line), and the QR estimate in this setup (average over 100 Monte Carlo repetitions, dashed blue line).