The Cross Section of Household Preferences and the MPC: Evidence from high frequency data

Elena Andreou Maria Demetriadou Andreas Tryphonides

University of Cyprus

18 July 2023

NBER SI, Micro Data and Macro Models

Motivation

- Renewed attention on *ex-ante* heterogeneity in heterogeneous agent macroeconomics.
- Matters for wealth heterogeneity and portfolio choice (Krusell and Smith (1998), Gomes and Michaelides (2005, 2008)) and the marginal propensity to consume (Kaplan and Violante (2022); Carroll, Slacalek, Tokuoka, and White (2017)).
- Recent evidence suggesting that observed consumption behavior is inconsistent with the benchmark incomplete markets model with homogeneous preferences (Aguiar, Bils, and Boar, 2023).

$\Longrightarrow$ Need to identify the distribution of preferences.

In this paper

- We show that assuming away liquidity constraints when identifying the distribution of preferences is not supported by the data (both in the PSID (1968-2019) and the KNCP (2006-2019) (Kilts-Nielsen Consumer Panel).
- We hence allow for:
  - **Inequality** in the Euler Equation $\Longrightarrow$ Consistent with arbitrary forms of liquidity constraints (and risk).
  - **Unrestricted heterogeneity** in preferences: High frequency observations of individual household consumption expenditures in the KNCP $\Longrightarrow$ Can estimate preferences *per household.*
- We provide **distribution estimates** for risk aversion, EIS and discount factor
  - Useful for calibrating ex-ante heterogeneity.
  - Less degrees of freedom for the HA model to match the cross sectional distributions of consumption and wealth.

In this paper

We employ the estimated preferences to answer two questions:

1. How do preferences effect the MPC?
2. How do preferences relate to HtM status?

In this paper

How do preferences effect the MPC?

- Using the supplementary survey on the 2008 tax rebate related to the Economic Stimulus Act, we match participating households to their estimated preferences.
- Controlling for other characteristics, we employ non-parametric techniques to estimate how preferences affect the reported MPC.
  - Evidence of non-monotonicity.
- We rationalize these patterns by calibrating an analytic approximation of the MPC to the estimated preferences.
  - Preferences enter non-linearly in the MPC.
  - The cross-sectional correlation of preferences is as important as the functional form of the MPC.

In this paper

Preferences and Hand-to-Mouth (HtM) status:

- No clear association of preferences with reported HtM status in the 2008 supplemental survey.
- But: We construct a test for who is HtM on average (during 2006-2019).
  - Strong association of preferences with average HtM status, consistent with standard theory.

Plan for rest of the talk

- Empirical methodology.
- Distribution estimates and the MPC.
- Preferences and HtM status.

Preference Identification and Liquidity Constraints

- How can we identify the distribution of preferences?
  - Popular identification approaches are model based, or utilize the Euler equation.
- A common assumption is the absence of liquidity constraints.

▸ Existing approaches ▸ Liquidity Matters

- We show that even in a model that allows for substantial heterogeneity in consumption, income and their co-dependence (Alan, Browning and Ejrnaes (2018)), extending the sample beyond 2009 leads to rejection of no excess sensitivity.

▸ Replication Details

- Assuming away constraints is particularly problematic if we want to identify the distribution of preferences.
  - We do not know a priori who is constrained or not.
  - Being constrained has to do with preferences as well.

Preference Identification and Liquidity Constraints

- The key idea is to allow for such constraints without taking a stance on the exact mechanism.
- Under general restrictions on trading across assets, the Euler condition holds with an inequality. Using Epstein-Zin-Weil preferences,

$$\mathbb{E}_{i,t}\left(\beta_i^{\theta_i}\left(\frac{c_{i,t+1}}{c_{i,t}}\right)^{-\frac{\theta_i}{\psi_i}} R_{w,t+1}^{\theta_i-1} R_{j,t+1}\right) \quad = \quad 1-\lambda_{i,t} \leq 1$$

where $\theta_i = \frac{1-\gamma_i}{1-\frac{1}{\psi_i}}$ and $R_{w,t+1}$ the return to total wealth.

Methodological Contribution

- We allow for moment inequalities to hold at the household level:

$$\mathbb{E}_i \left[\left(1 - \beta_i^{\theta_i} \left(\frac{c_{i,t+1}}{c_{i,t}} \right)^{-\frac{\theta_i}{\psi_i}} R_{w,t+1}^{\theta_i - 1} R_{j,t+1} \right) z_{i,t} \right] \geq 0$$

- Preferences are no longer point identified i.e. there is no unique value for $(\beta_i, \gamma_i, \psi_i)$ that satisfies the moment restrictions, even if sample size is infinite.
- The distribution of preferences is non-parametrically (set) identified.

Data

- We employ the consumption expenditures in the KNCP dataset: Monthly aggregation (Bansal, Kiku and Yaron (2016) find empirical support for monthly decision interval using aggregate data.)
- Time series dimension is important to identify the preferences of each household in the panel.
- Need as much time series dimension as possible: Restrict the cross section to households that are present in the survey for at least 10 consecutive years ($\sim$ 12,500 households).

Estimated Distributions

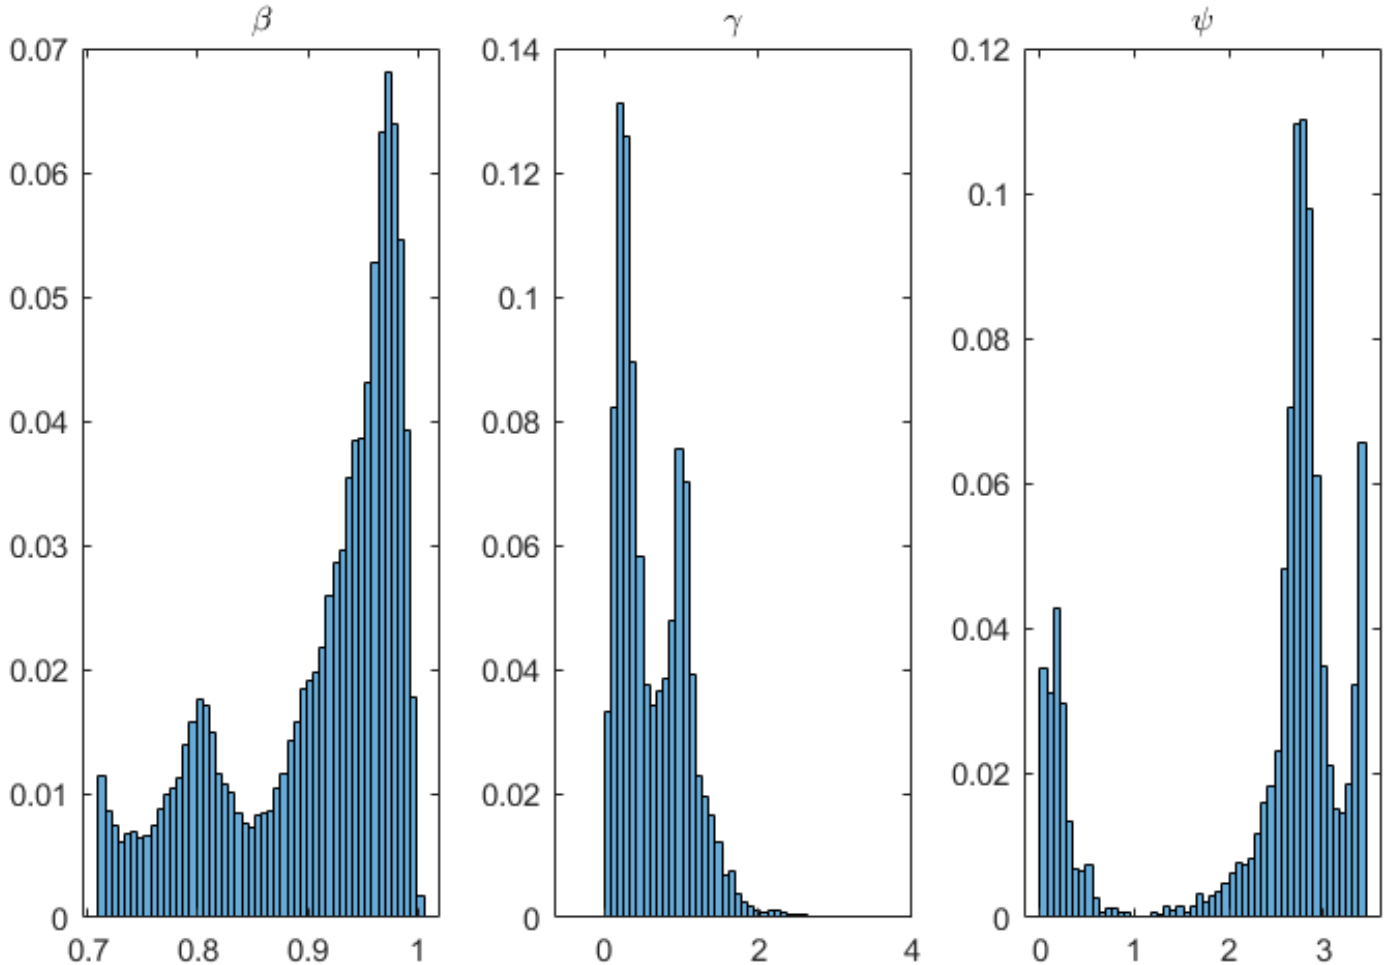

Figure: Histogram of the distribution of estimated preferences (pooled sets, re-sampled using US population weights.)

Correlations

	Discount Factor	Risk Aversion	EIS
Discount Factor	1		
Risk Aversion	0.090	1	
EIS	0.182	-0.545	1

Table: Correlation matrix of estimated EZ preferences.

▸ Scatter Plots

The 2008 Tax Rebate Survey

- Respondents were asked about the dollar amount of their tax rebate (stimulus payment) and then to report their MPC (the extra amount they are spending on consumer goods) in response to this rebate.
- We match the households that reported receiving or going to receive the tax rebate to their estimated preferences (2,159 households).
- We focus on modelling the positive MPC, and discard households with MPC=1 if they report that they are mostly saving/repaying debt.
- The decision to consume versus save/repay debt deserves separate analysis.

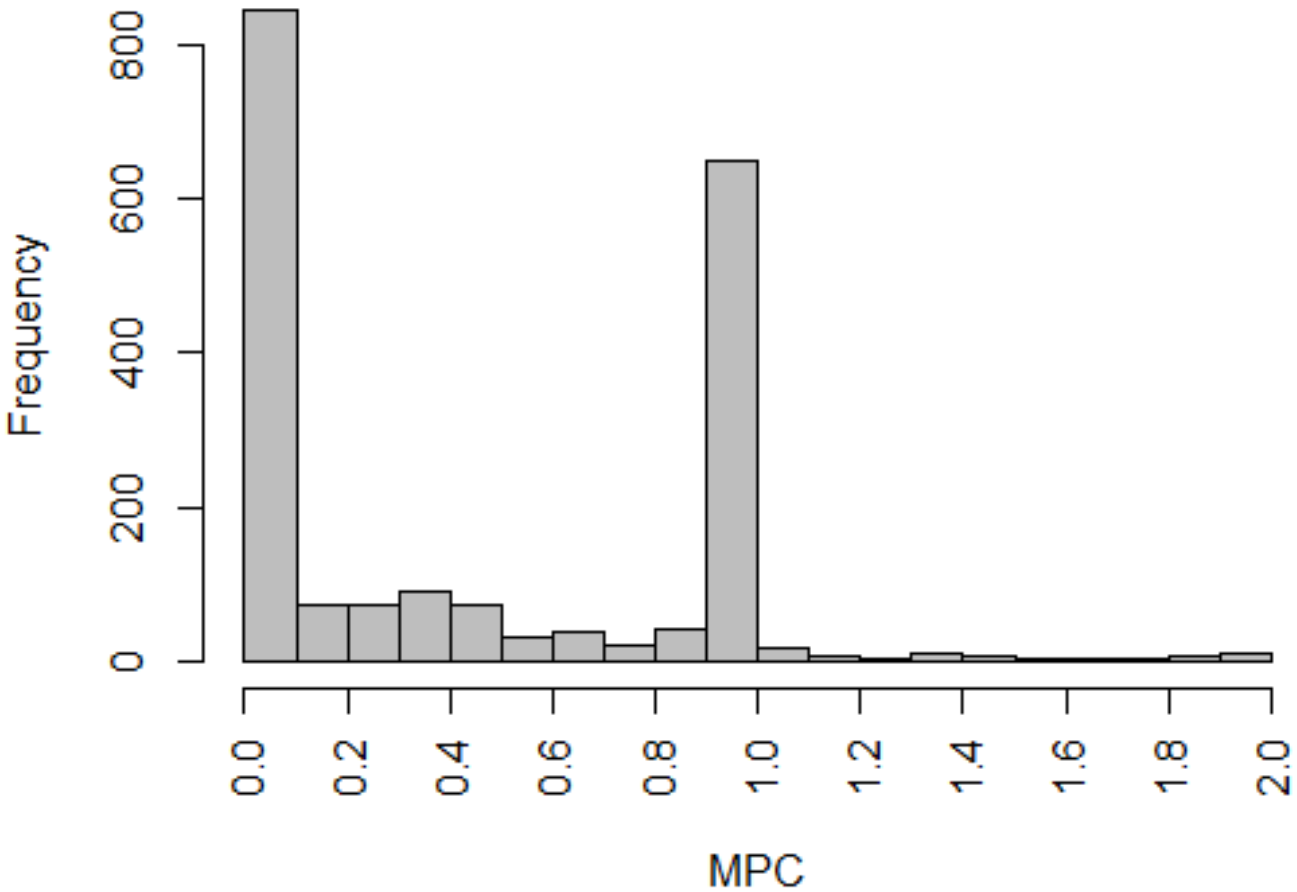

Figure: Histogram of the distribution of reported MPC, 2008

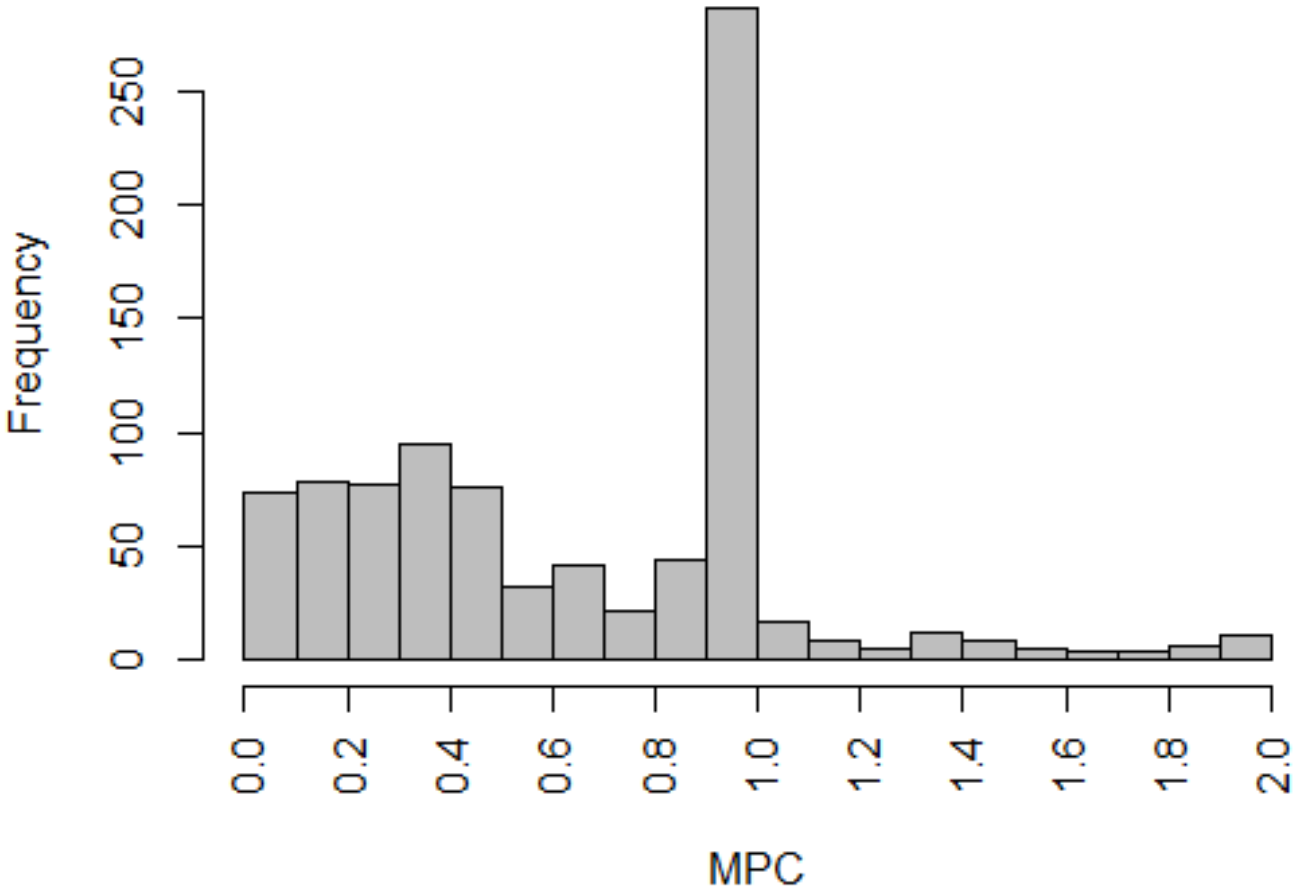

Figure: Histogram of the distribution of reported MPC positive values, 2008, without MPC=1 if reporting that they are mostly saving/paying off debt

Bayesian Non-Parametric model

We consider a Bayesian Additive Regression Tree (BART) model:

$$MPC_i^* = f(\hat{\vartheta}_i^*) + e_i, \qquad e_i \overset{\text{iid}}{\sim} N(0, \sigma^2) \tag{1}$$

- $*$ are the residuals after projecting out observed characteristics. ▸ Results
- f is represented as the sum of many regression trees.
- $\hat{\vartheta}_i = (\hat{\beta}_i, \hat{\gamma}_i, \hat{\psi}_i)$ are the medians of the estimated preferences that correspond to the i^{th} household. ▸ Bayesian algorithm that allows for sets

- We control for the effect of wealth on the MPC, by focusing on the (liquidity) unconstrained households. ▸ MPC example

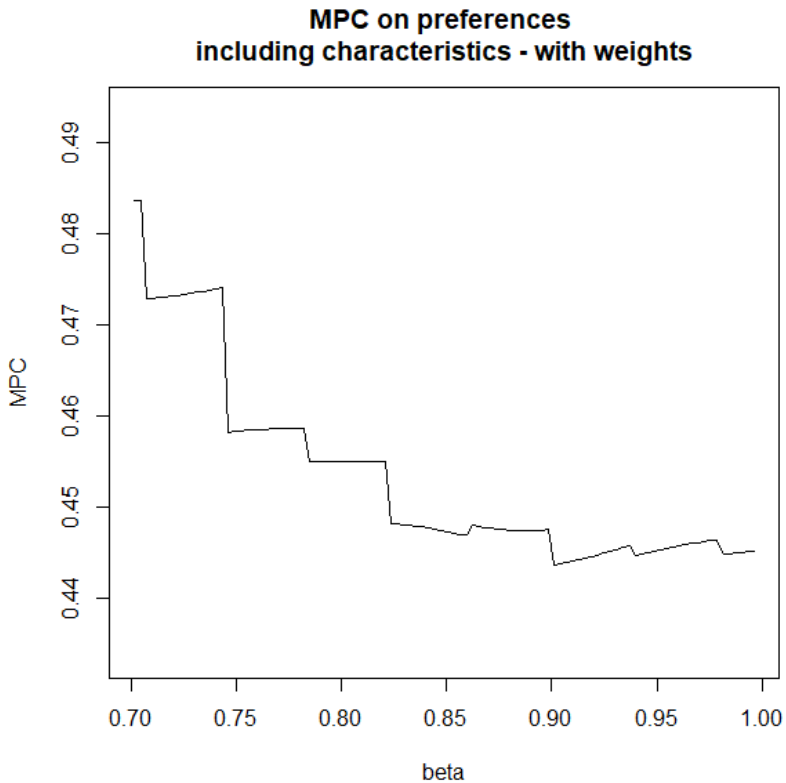

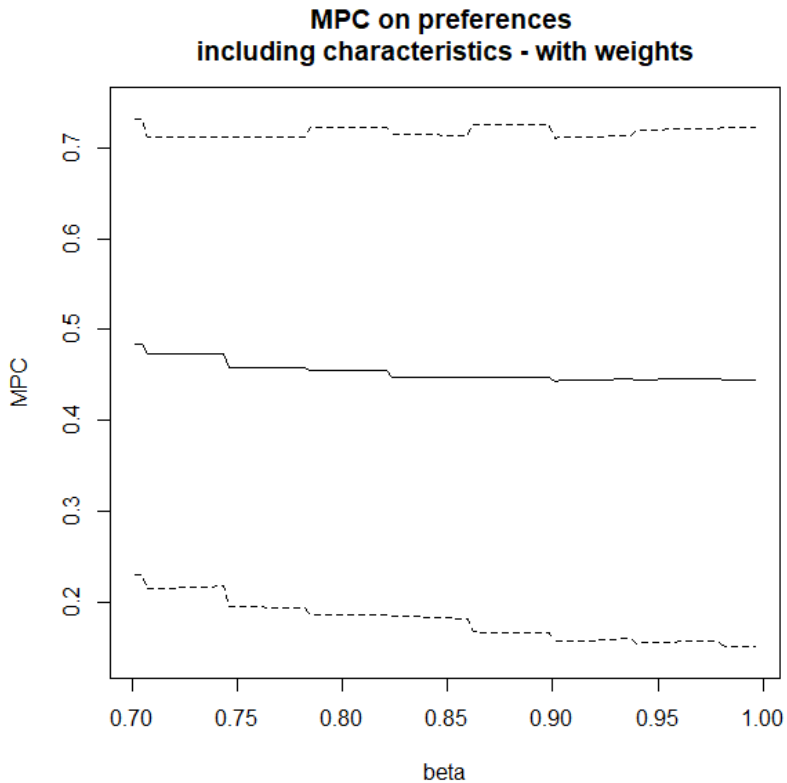

Figure: The marginal effect of discount factor on MPC while aggregating over the other preferences.

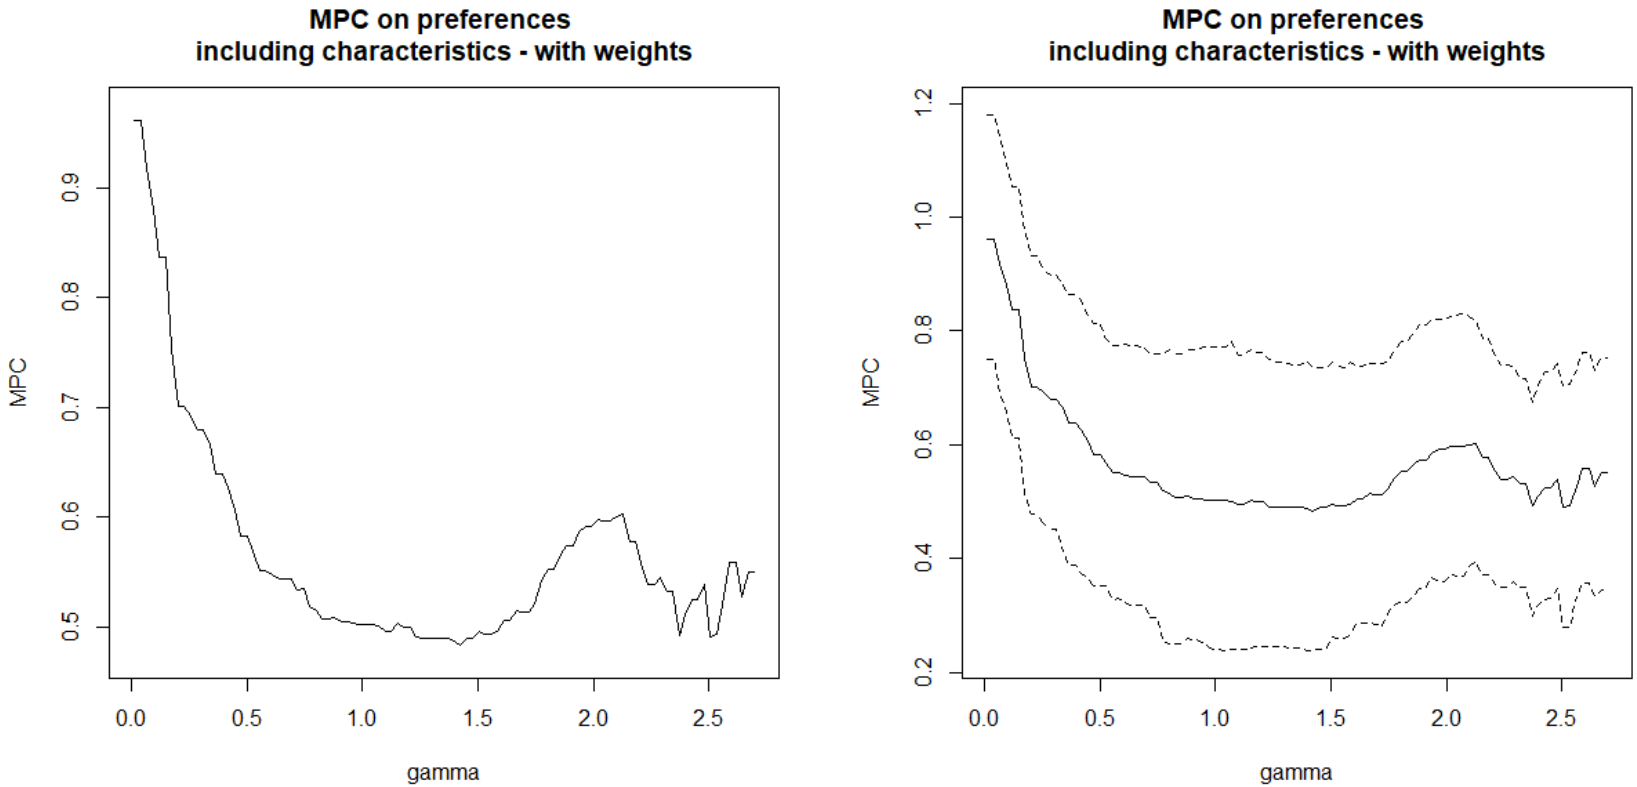

Figure: The marginal effect of risk aversion on MPC while aggregating over the other preferences.

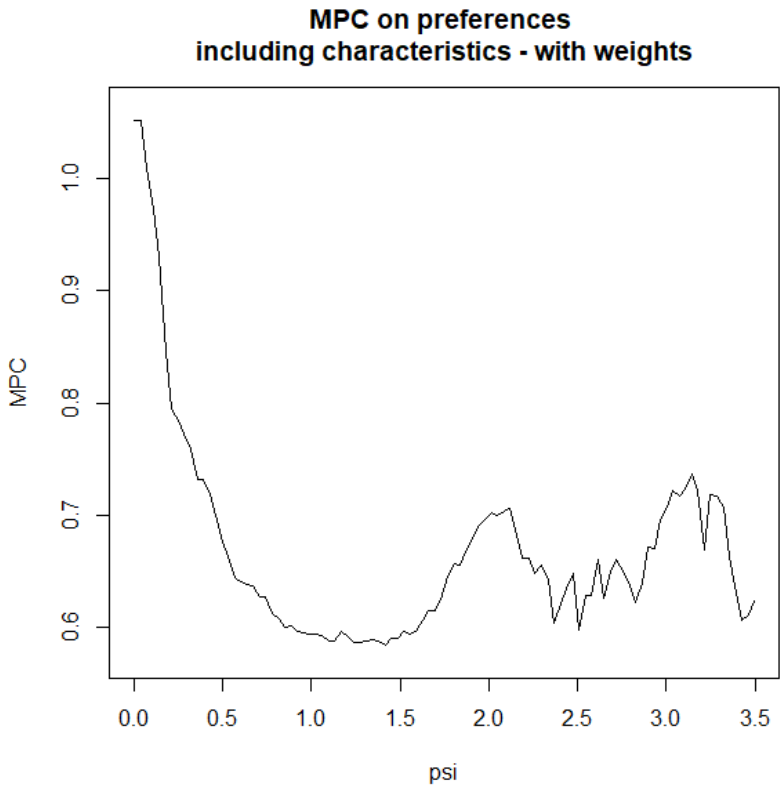

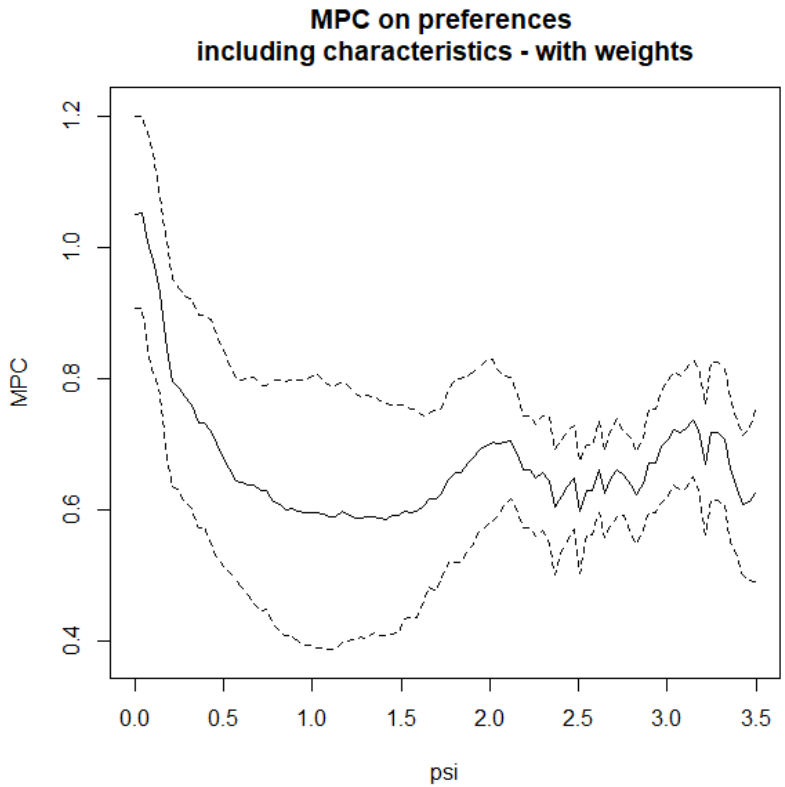

Figure: The marginal effect of intertemporal elasticity of substitution on MPC while aggregating over the other preferences.

Evidence versus theory

- Would theory predict these relationships?
- Assuming jointly log-Normal returns for wealth and asset j, with means (τ_w, τ_j) respectively, the MPC for liquid households is

$$\begin{aligned}\mu_i \quad = \quad & \underbrace{\tau_w - \frac{1}{\gamma_i - 1}(\tau_j - \tau_w)}_{\text{Income effect}} + \underbrace{\psi_i \left(\frac{1}{\beta_i} - 1 - \left(\tau_w - \frac{1}{\gamma_i - 1}(\tau_j - \tau_w) \right) \right)}_{\text{Substitution effect}} \\ & \underbrace{-\gamma_i \frac{1 - \psi_i}{\gamma_i - 1} \left((1 + \gamma_i) \frac{1}{2} \sigma_w^2 - \rho_{w,j} \right)}_{\text{Precautionary savings}} \end{aligned} \tag{2}$$

- We evaluate this expression at the estimated preferences. ▸ One asset

Evidence versus theory

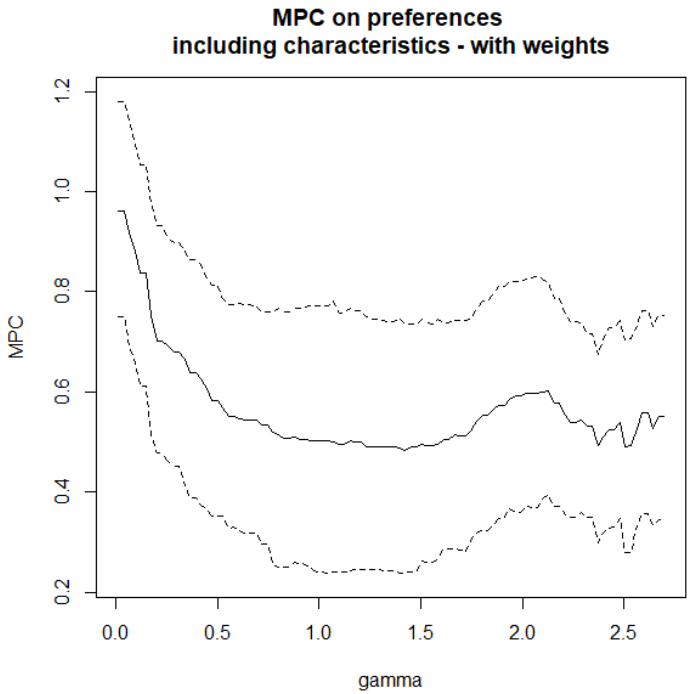

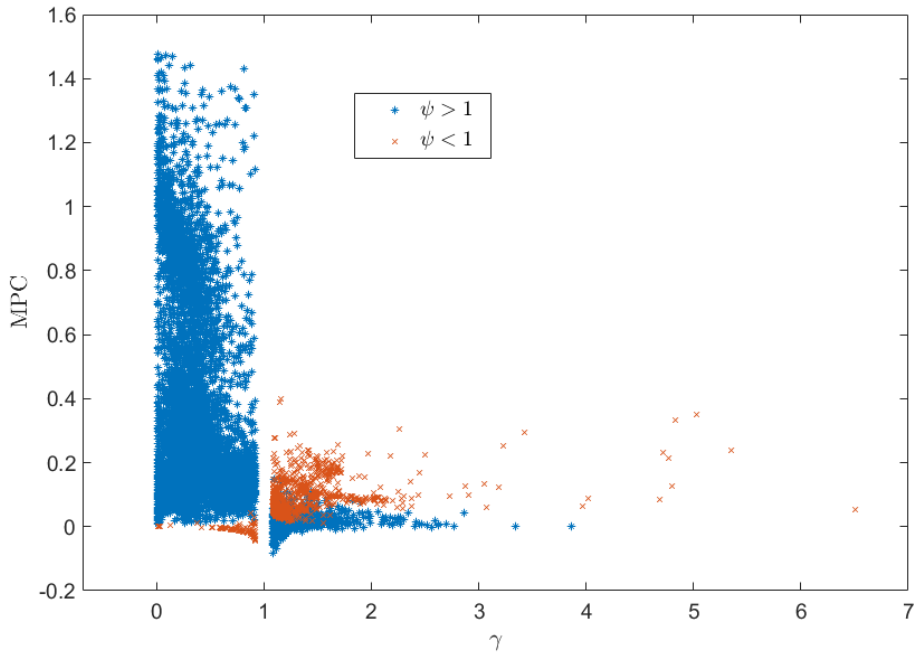

Evidence versus theory

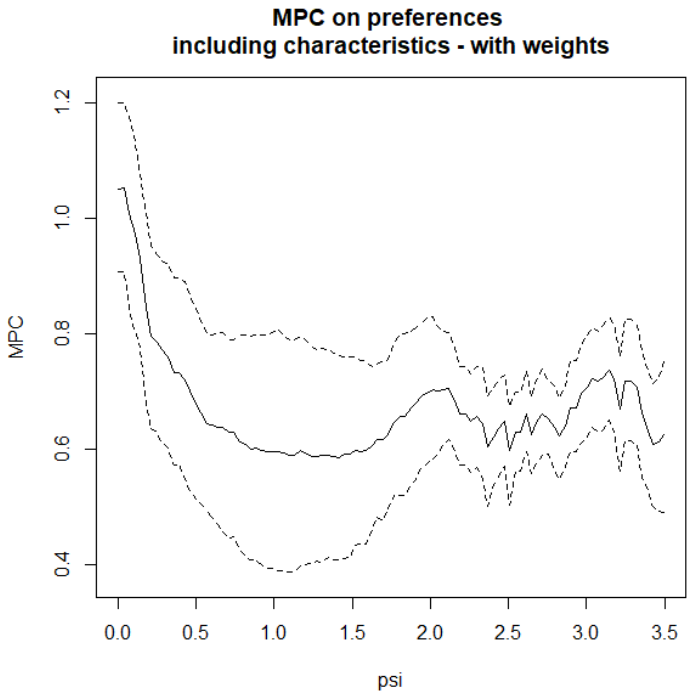

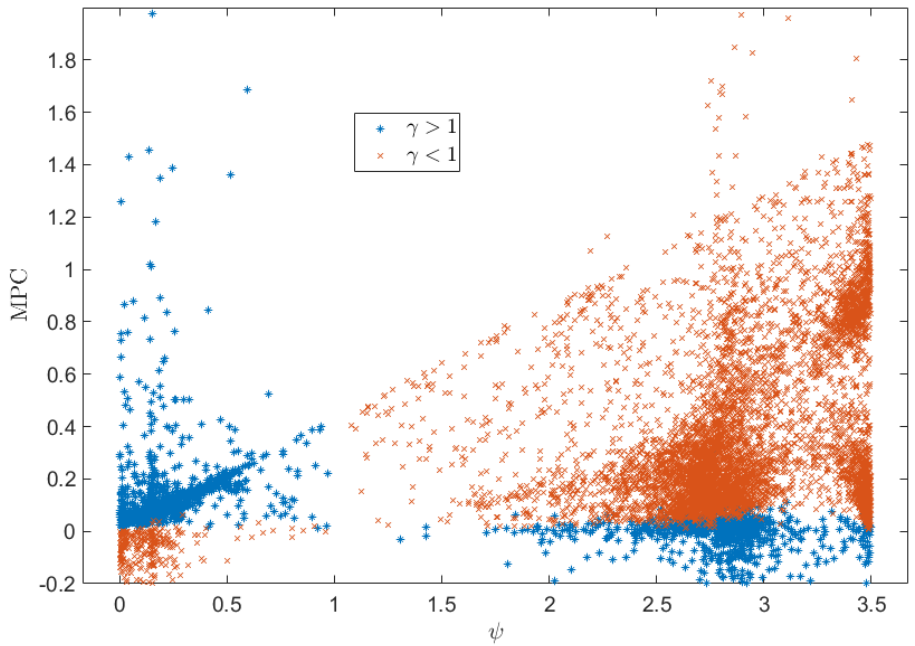

Latent Heterogeneity in the MPC

- In our sample the observed characteristics explain 5% of the variation in the MPC. Preferences add 3pp.
- Much of MPC heterogeneity remains unexplained.
- Same feature exists in almost all relevant studies with $R^2 \sim$ 1%-15%. (e.g. Lewis, Melcangi and Pilossoph (2022)).

Identifying the average Hand-to-Mouth

- By not taking a stance on who is HtM when estimating preferences, we are able to use our estimates to infer who is HtM on average.
- The estimated wedge to the Euler equation is:

$$\bar{\lambda}_i = \frac{1}{T}\sum_{t=1..T}\left(1-\beta_i R^j_{t+1}\left(\frac{c_{i,t+1}}{c_{i,t}}\right)^{-\frac{\theta_i}{\psi_i}} R^{\theta_i-1}_{w,t+1}\right) z_{i,t}$$

- We separately identify the two groups using

$$Q_{i,T}(\vartheta_i^{median}, 0) = \frac{1}{2}\bar{\lambda}_i^T V^{-1}\bar{\lambda}_i$$

- Since $TQ_{i,T} \sim \chi^2(k) \Rightarrow$

$$\text{Constrained if } TQ_{i,T} > \chi^2_{1-\alpha\%}(5)$$

- HtM: 50.7%, higher than Aguiar, Bils and Boar (2022), (40.7%) which uses PSID data and employs liquidity based criteria similar to Kaplan and Violante (2014) and net worth based criteria as in Zeldes (1989).

Preferences are important determinants of hand-to-mouth status.

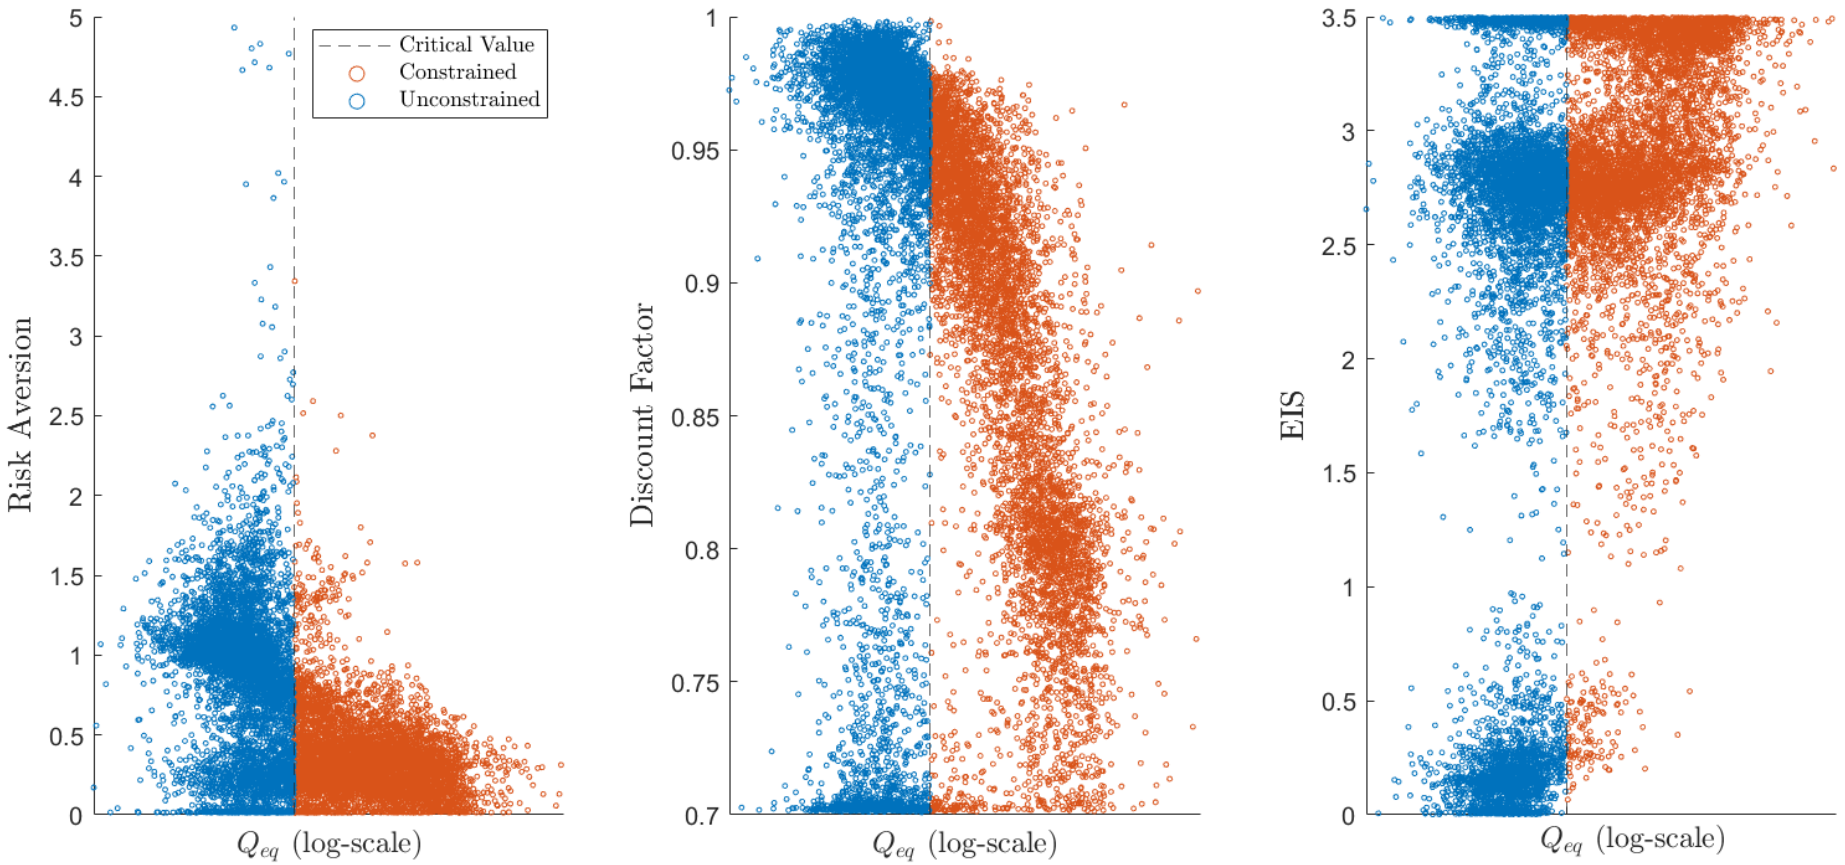

Figure: Liquidity constrained (orange) versus unconstrained (blue)

Conclusions

- This paper makes methodological and empirical contributions on estimating the distribution of preferences and their role in determining the MPC.
- Estimated distributions can be used to calibrate heterogeneous agent models $\Longrightarrow$ makes other mechanisms testable.
- The direct evidence we bring on the MPC provides empirircal support to the standard model, yet, lot's of unexplained variation.
- Empirical support for the role of preferences in determining the HtM status.

Questions?

Thank you for your attention!

Existing approaches

- Blundell, Pistaferri and Preston (2008) identify (homogeneous) Frisch elasticities by employing approximations to the Euler equation and the budget constraint. Theloudis (2021) extends this to identify first and second moments of these elasticities in the cross section using PSID data.
- Alan and Browning (2010) simulate consumption paths using a flexible parametric form for the Euler equation error and preferences which are then matched to moments computed from the PSID. Alan, Browning and Ejrnaes (2018) allows codependence between preference and income parameters.
- Calvet, Campbell, Gomez and Sodini (2022) use Swedish administrative data and estimate the same lifecycle model across household groups.

▸ Back

Liquidity Matters

- Crawley and Kuchler (2020) response of consumption to income shocks differ across households with different levels of liquid wealth.
- Boutros (2019) finds that households with less liquid wealth are more likely to repay debt using fiscal stimulus payments than increasing consumption.
- Jappelli and Pistaferri (2014) and Parker (2017) find that households with low liquidity have a much higher MPC than more liquid households.
- Kaplan, Violante and Weidner (2014) find that hand-to-mouth households, especially wealthy ones, exhibit the largest consumption responses to transitory income shocks. Back

Do liquidity constraints matter?

- Alan, Browning and Ejrnaes (2018) employed PSID data from 1968 to 2009 to estimate the distribution of discount factor and risk aversion, allowing for vast heterogeneity.
- They test their specification by examining the excess sensitivity of consumption to current income using a QLM test. $\Longrightarrow$ No evidence of excess sensitivity.
- We replicated the results of Alan, Browning and Ejrnaes (2018) for 1968-2009 (their sample) and 1968-1997 (before PSID turned biennial)$\Longrightarrow$ Indeed, no evidence of excess sensitivity.
- We repeated the estimation using the updated full sample 1968-2019 and the biennial sample 1999-2019 $\Longrightarrow$ Evidence of excess sensitivity.
- We repeated all of the above tests using the Kilts-Nielsen Consumer Panel (KNCP) dataset, which involves a more comprehensive measure of consumption to PSID, where we aggregated to yearly consumption $\Longrightarrow$ Evidence of excess sensitivity.

Back

Discount Factor

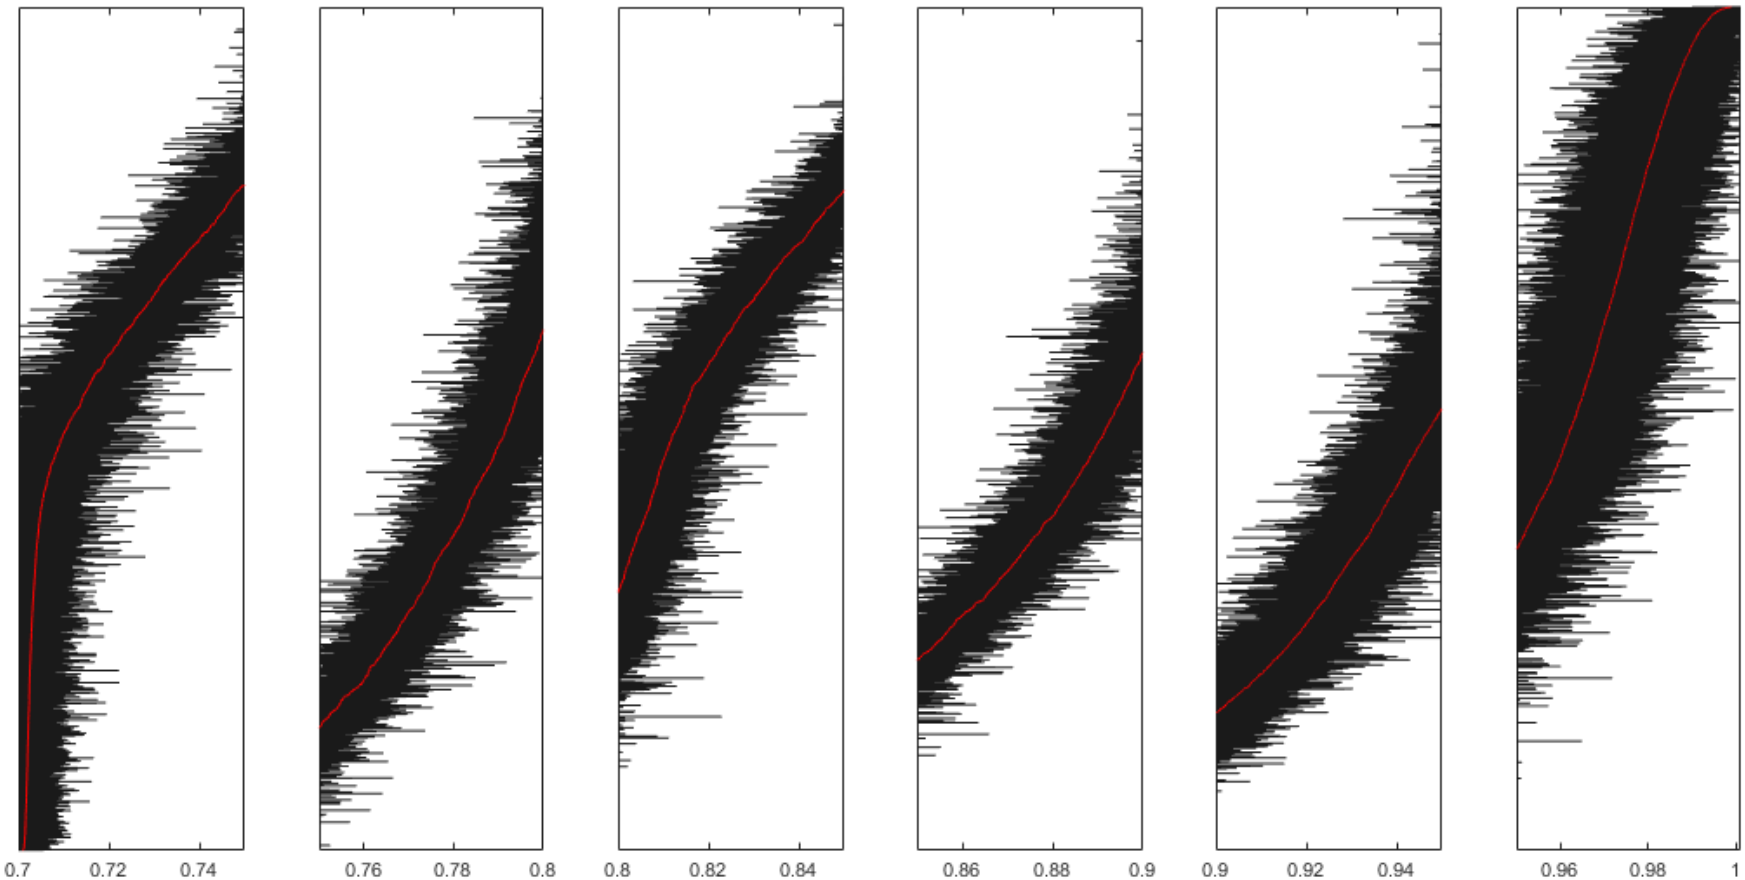

Figure: Household confidence sets, sorted by median

Risk Aversion

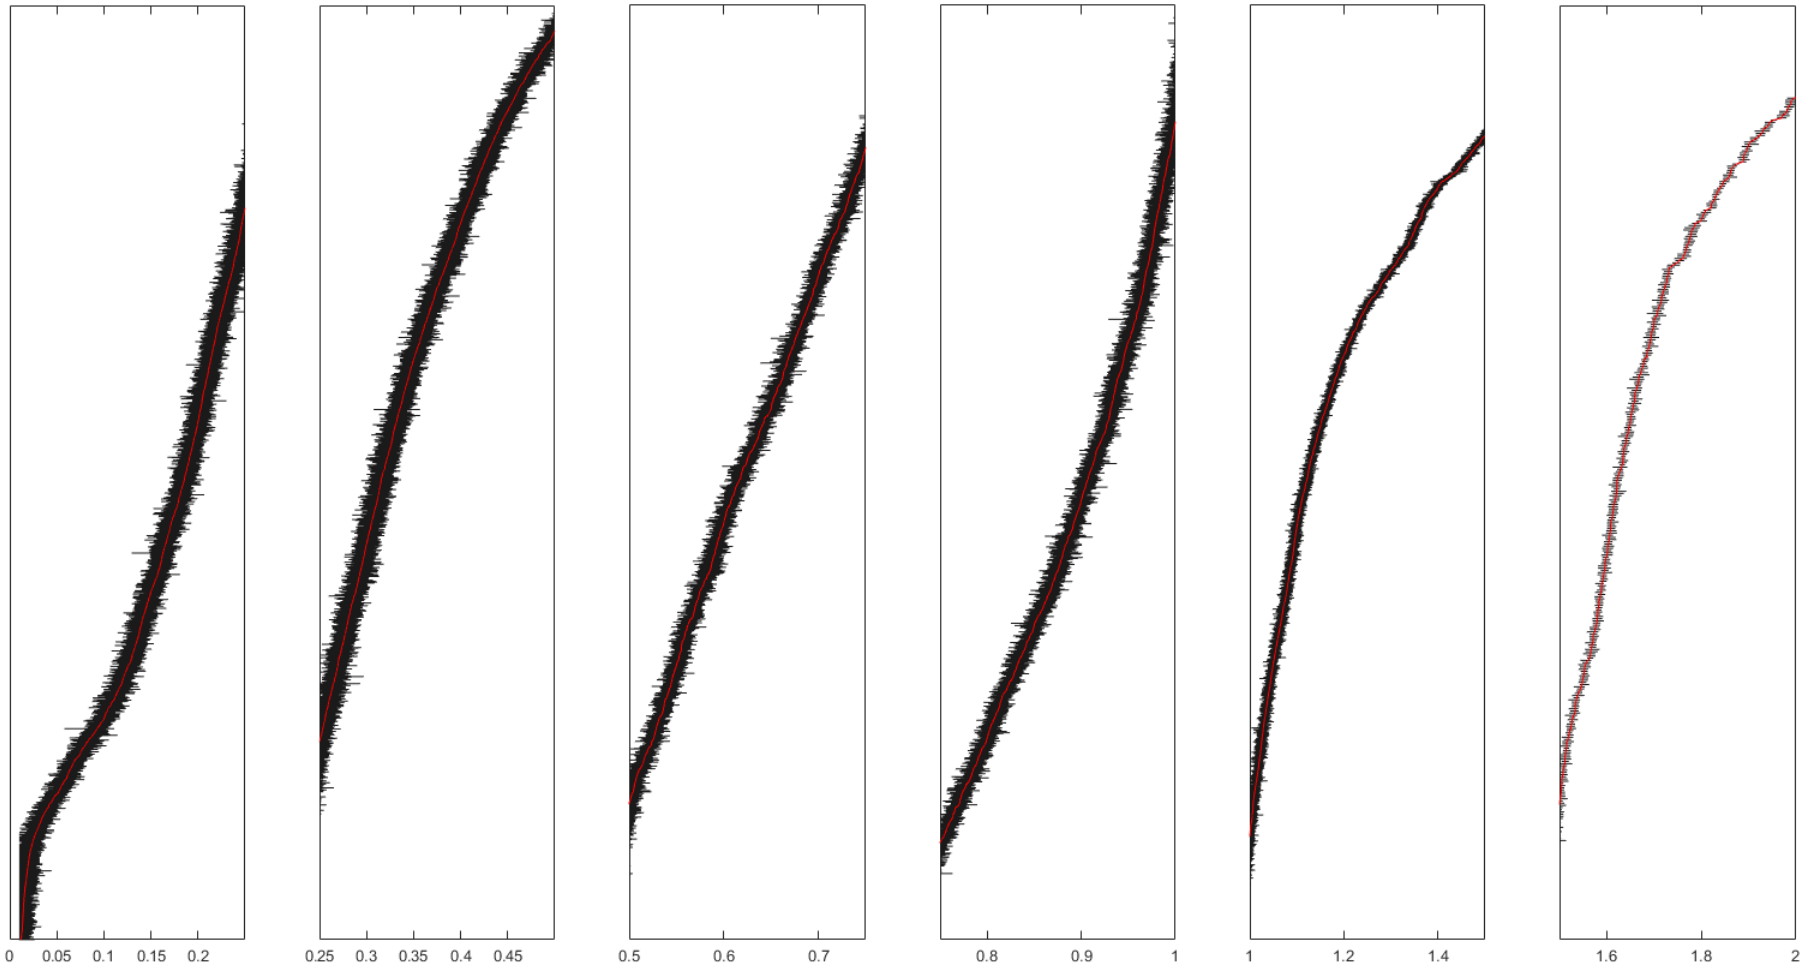

Figure: Household confidence sets, sorted by median

EIS

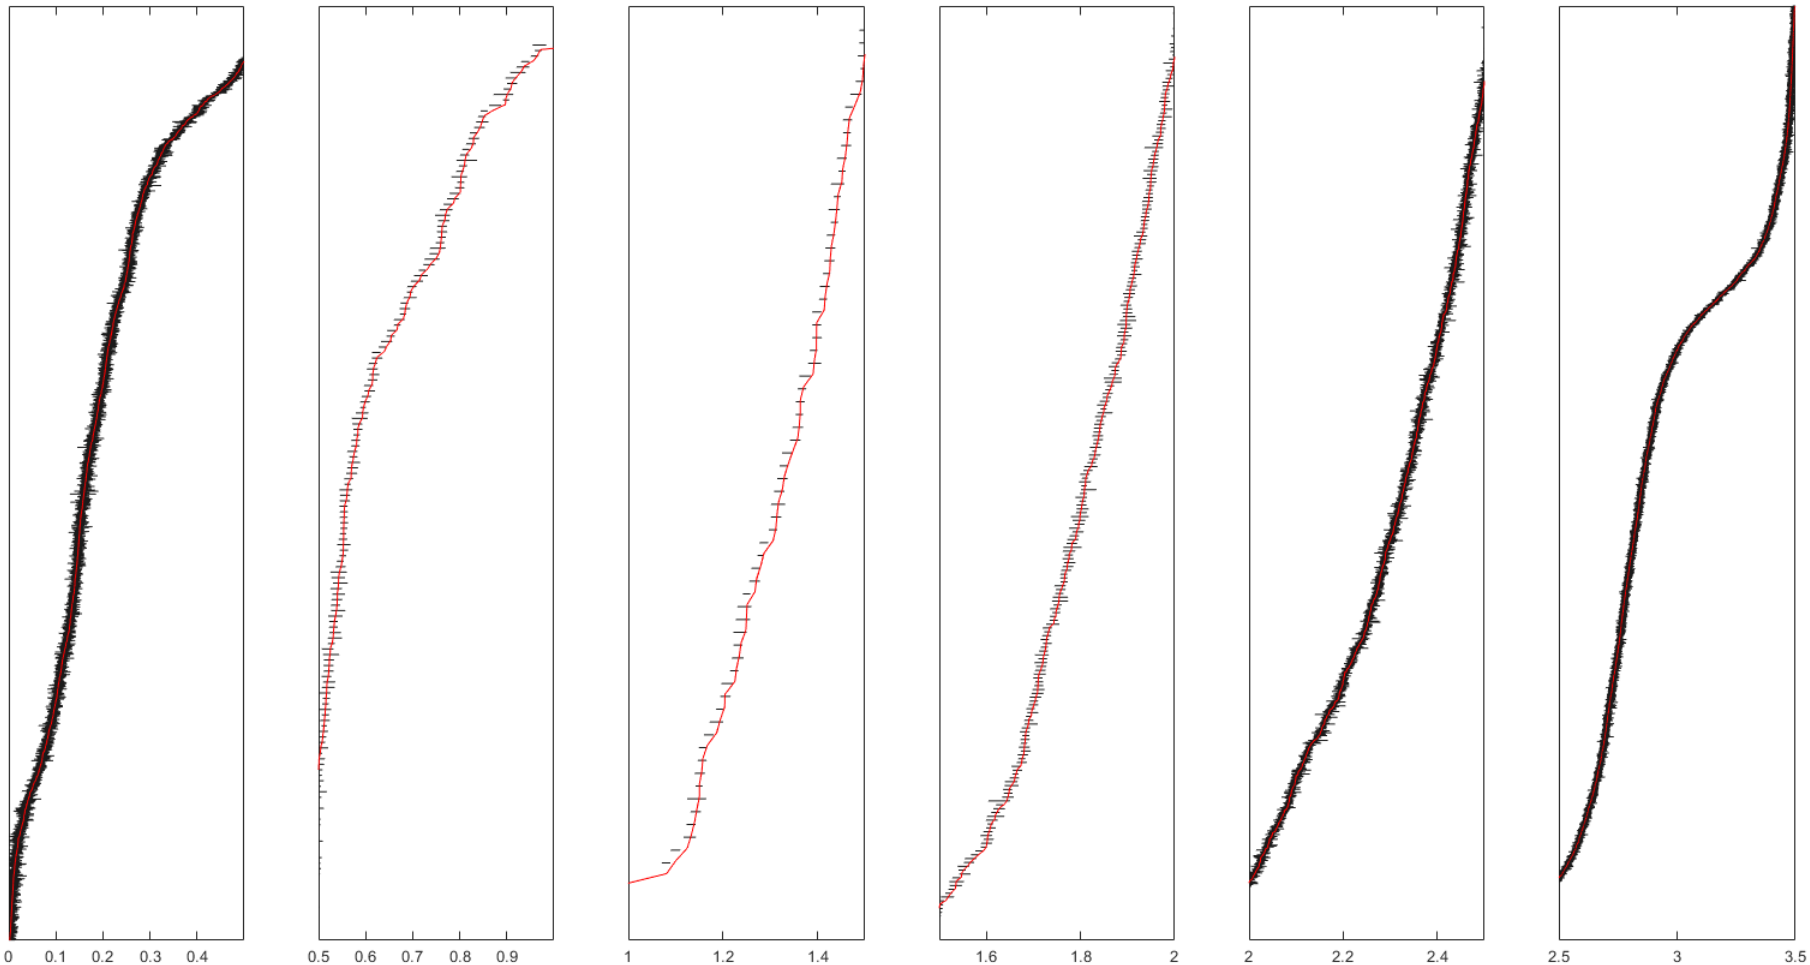

Figure: Household confidence sets, sorted by median

Preference Scatter Plots

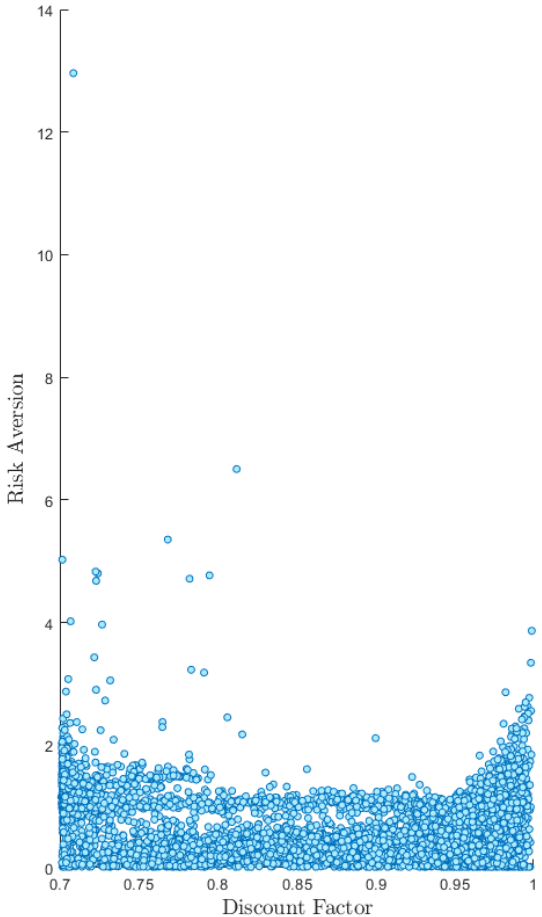

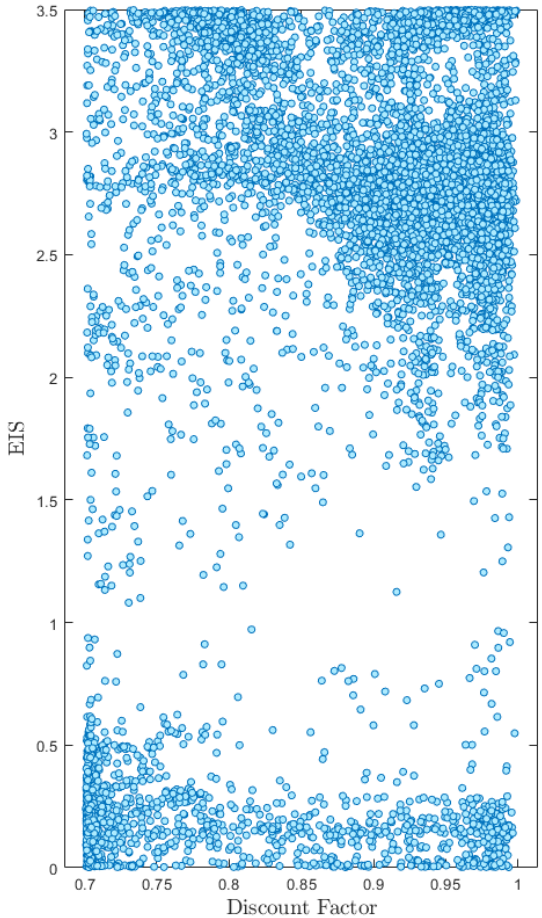

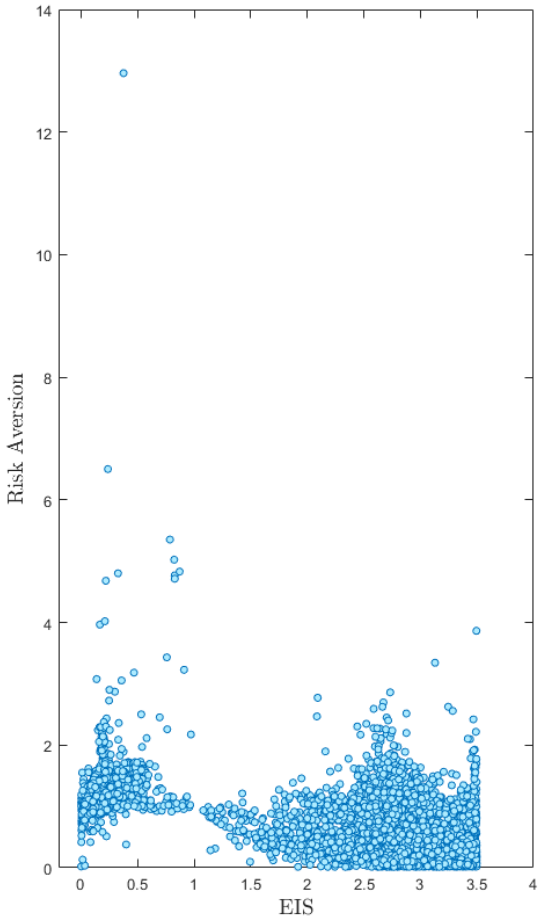

back

Set Valued Preferences - Inference

- The joint (quasi) posterior of (α, ϑ) can be written as:

$$
\begin{aligned}
& p_{N,T}(\alpha, \{\vartheta, U\}_{i=1..N} \mid \mathcal{Y}_{N,T}) && (3)\\
= \; & p\left(\alpha \mid \{\vartheta\}_{i=1..N}, \mathcal{Y}_N^{T=2008}\right) p\left(\{\vartheta, U\}_{i=1..N} \mid \mathcal{Y}_{N,T}\right) && (4)\\
= \; & p\left(\alpha \mid \{\vartheta, U\}_{i=1..N}, \mathcal{Y}_N^{T=2008}\right) \prod_{i=1..N} p\left(\vartheta_i, U_i \mid \mathcal{Y}_T^i\right) && (5)
\end{aligned}
$$

  where $p(\vartheta_i, U_i) \propto \pi(\vartheta_i, U_i) e^{-TQ_T(\vartheta_i, U_i)}$

- The algorithm of Chen, Christensen and Tamer (2018) uses the CU-GMM criterion

$$Q_T(\vartheta_i, U_i) = \frac{1}{2}(\bar{m}(y_{i,t,t+1}, Y_{t,t+1}, \vartheta) - U_i)^T V^{-1}(\bar{m}(y_{i,t,t+1}, Y_{t,t+1}, \vartheta) - U_i)$$

  to obtain draws of partially identified parameters ϑ_i.

back

CRRA preferences and consumption

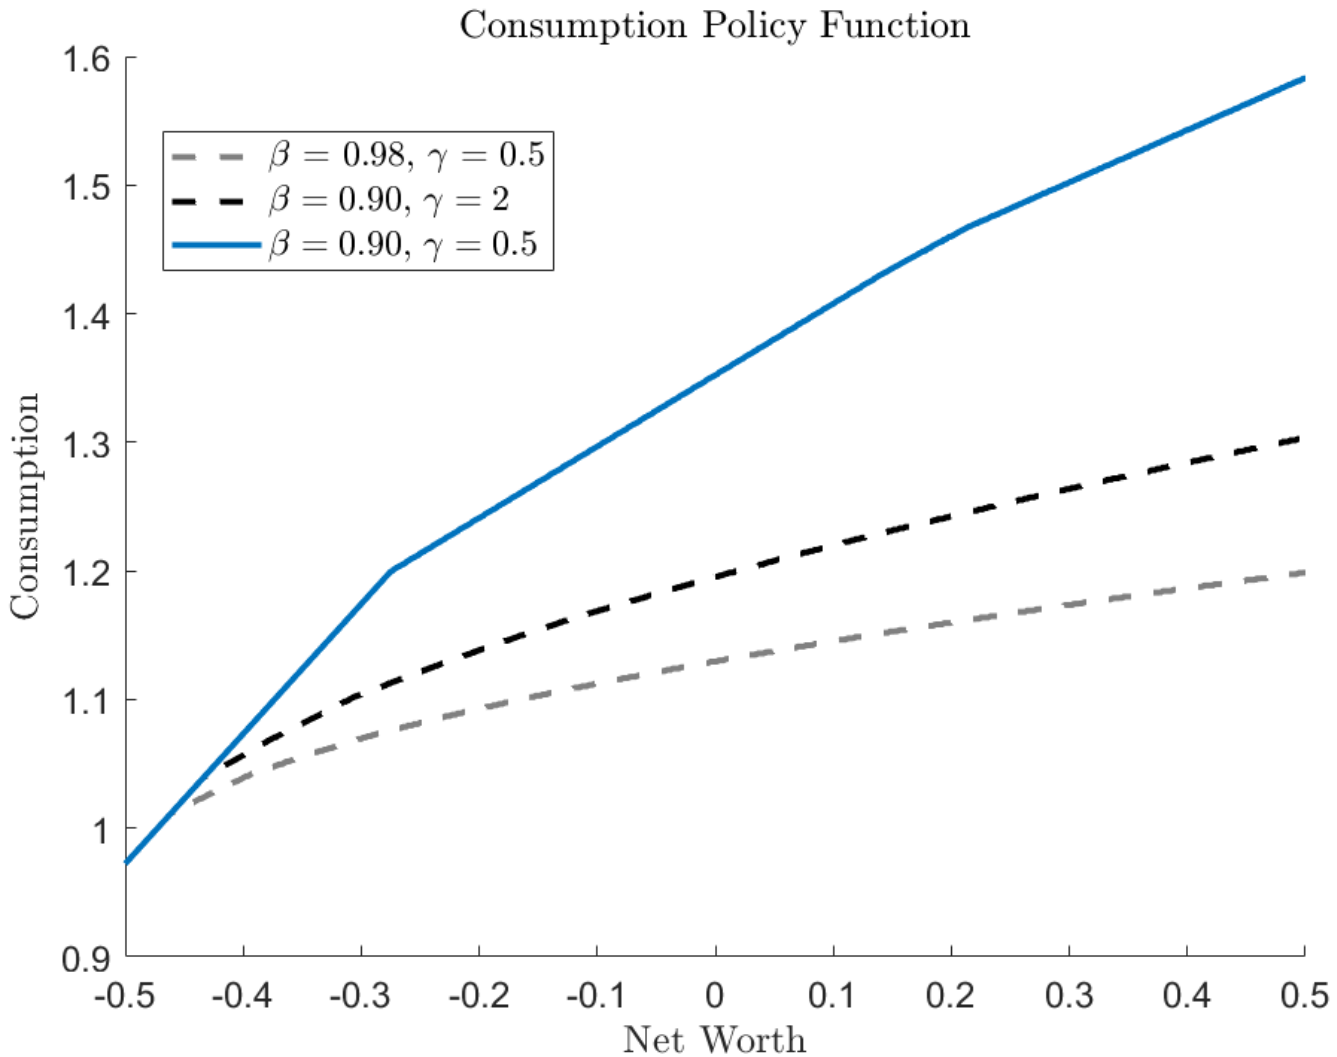

back

Effects of characteristics on MPC

Marginal Propensity to Consume

	Coef.	q_5	q_{95}
Const.	0.654	0.563	0.745
Married	0.069	-0.006	0.143
High Educ.	-0.086	-0.153	-0.019
High Prof.	0.105	0.028	0.183
Income	0.027	-0.001	0.055
House Owning	-0.113	-0.230	0.004
Unemp.	0.056	-0.018	0.130
HH Size	-0.049	-0.078	-0.019
Liq. Constr.	-0.014	-0.087	0.059

back

Constrained versus Unconstrained

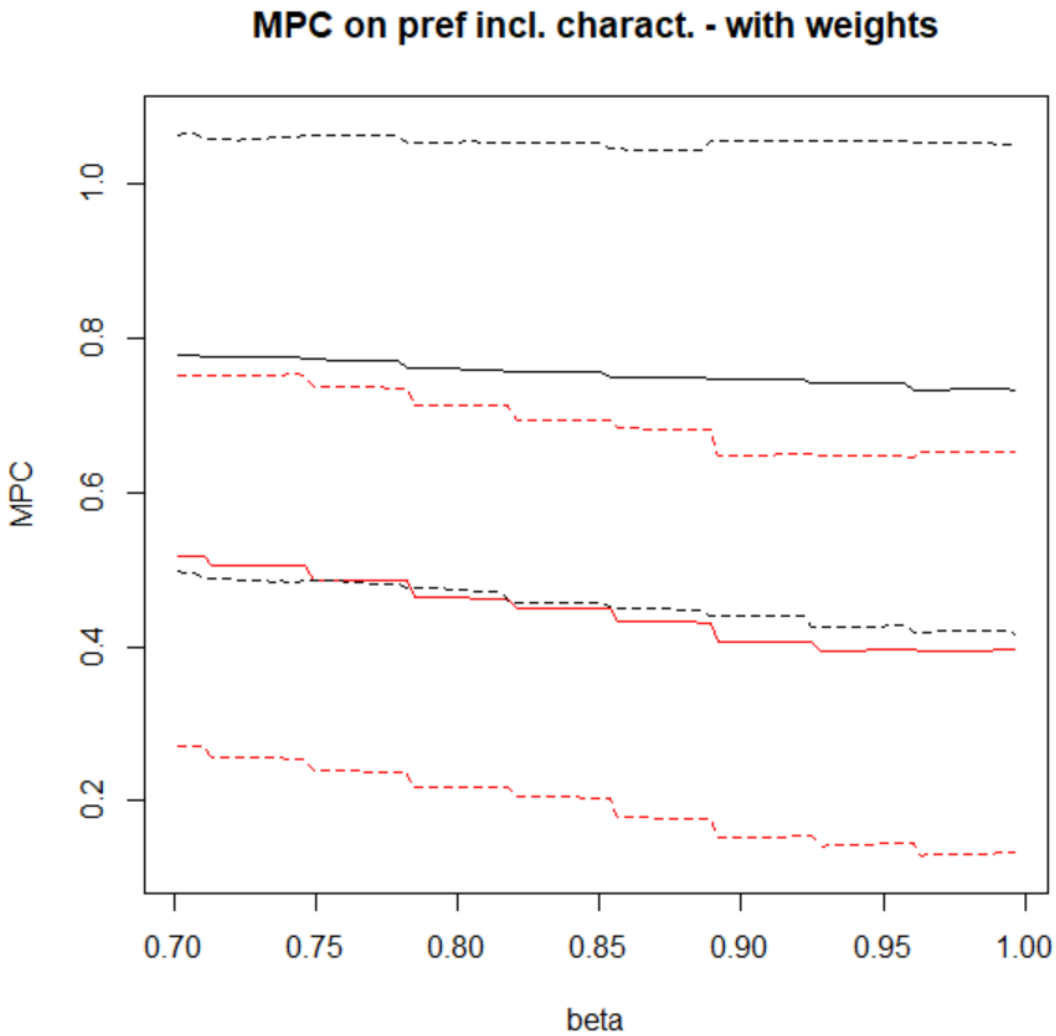

Figure: The marginal effect of discount factor on MPC while aggregating over the other preferences.

Constrained versus Unconstrained

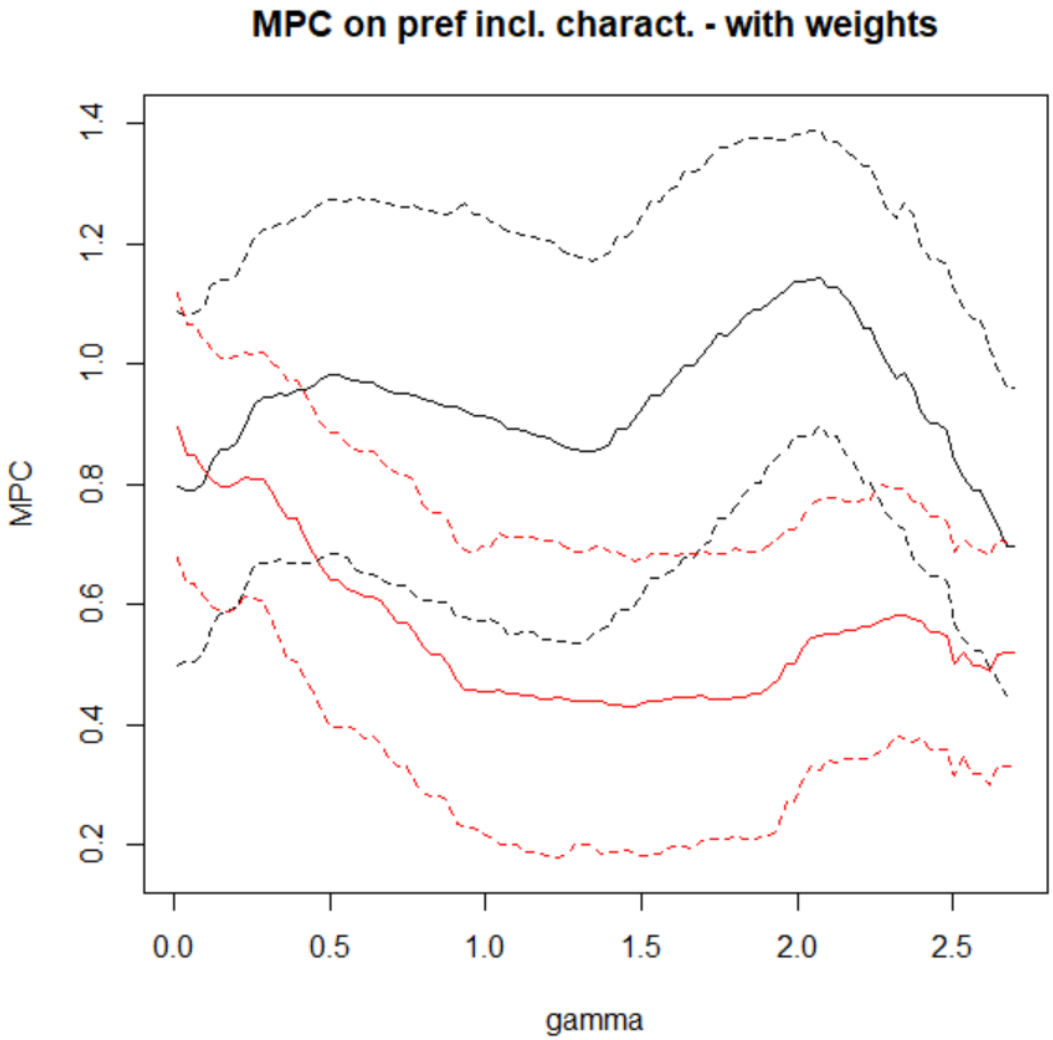

Figure: The marginal effect of risk aversion on MPC while aggregating over the other preferences.

Constrained versus Unconstrained

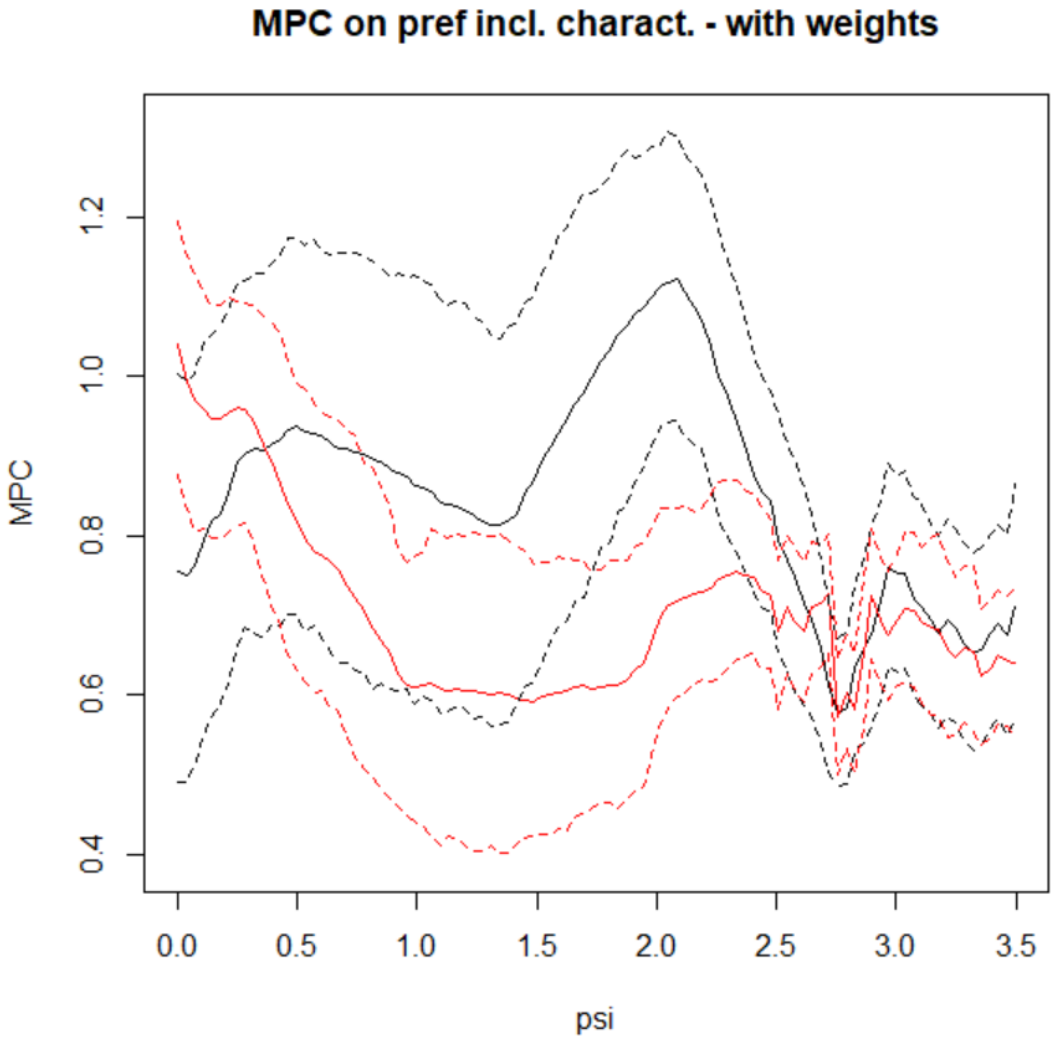

Figure: The marginal effect of intertemporal elasticity of substitution on MPC while aggregating over the other preferences.

With characteristics

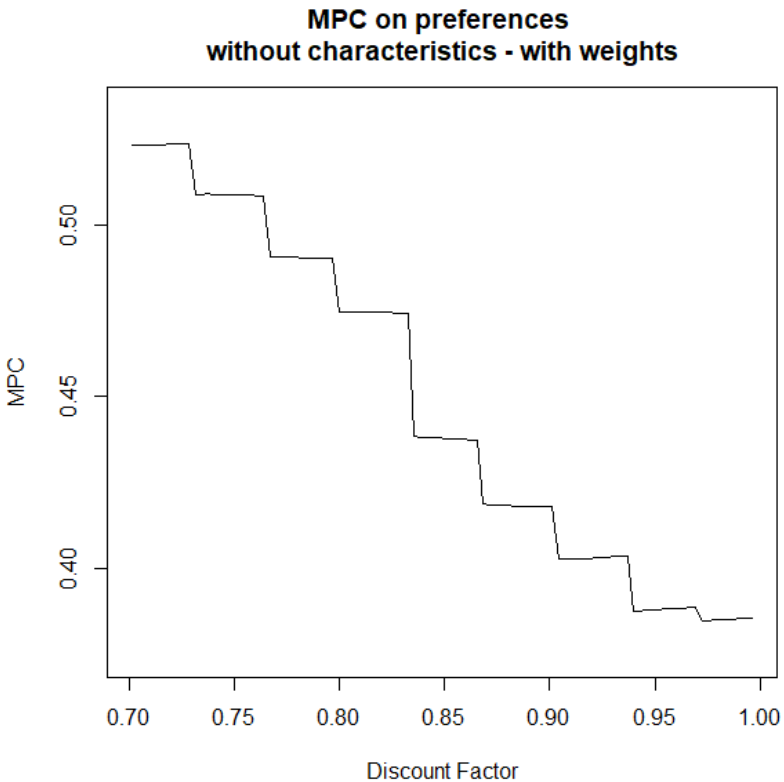

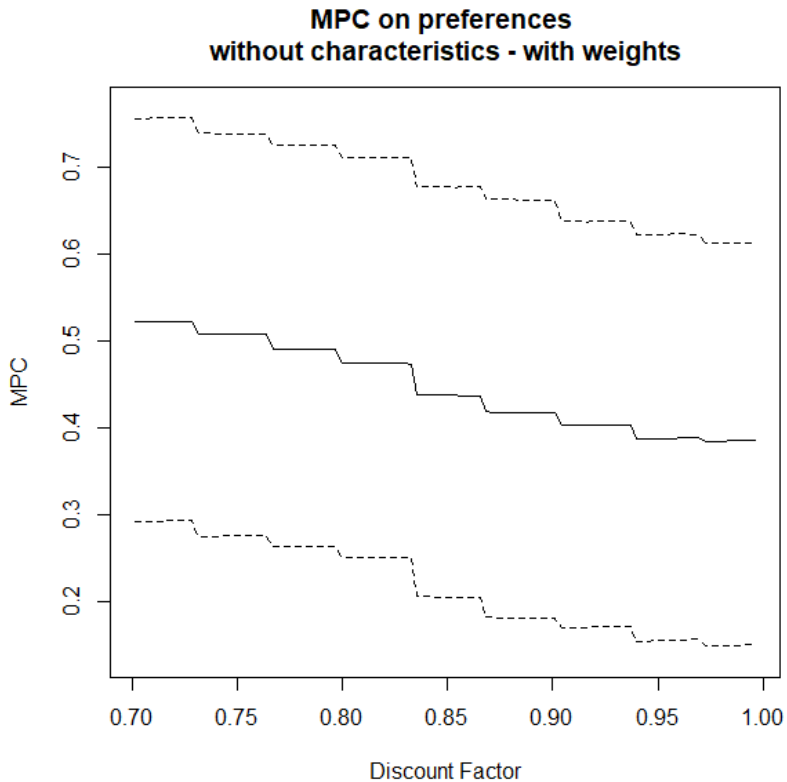

Figure: The marginal effect of discount factor on MPC while aggregating over the other preferences.

With characteristics

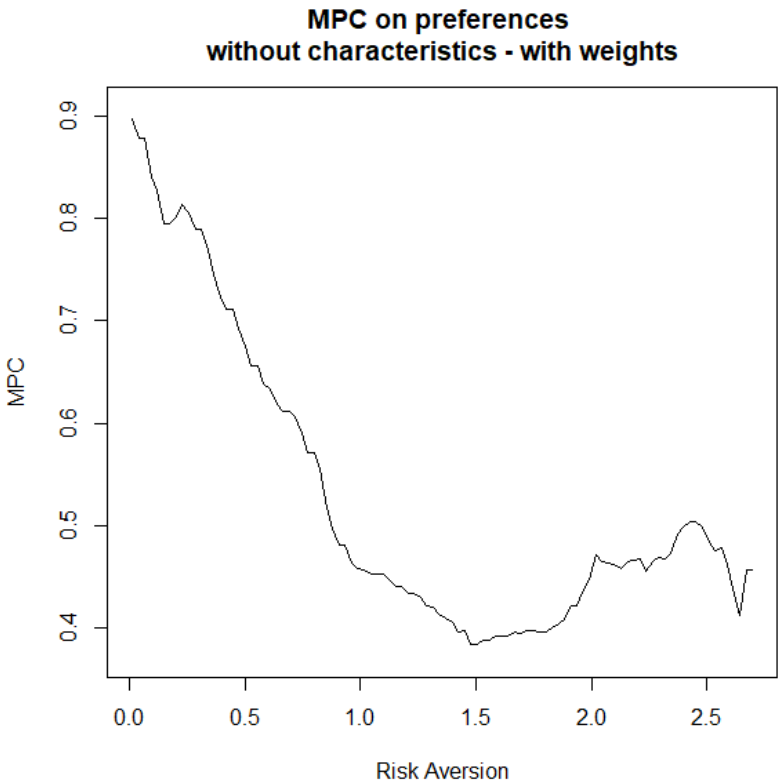

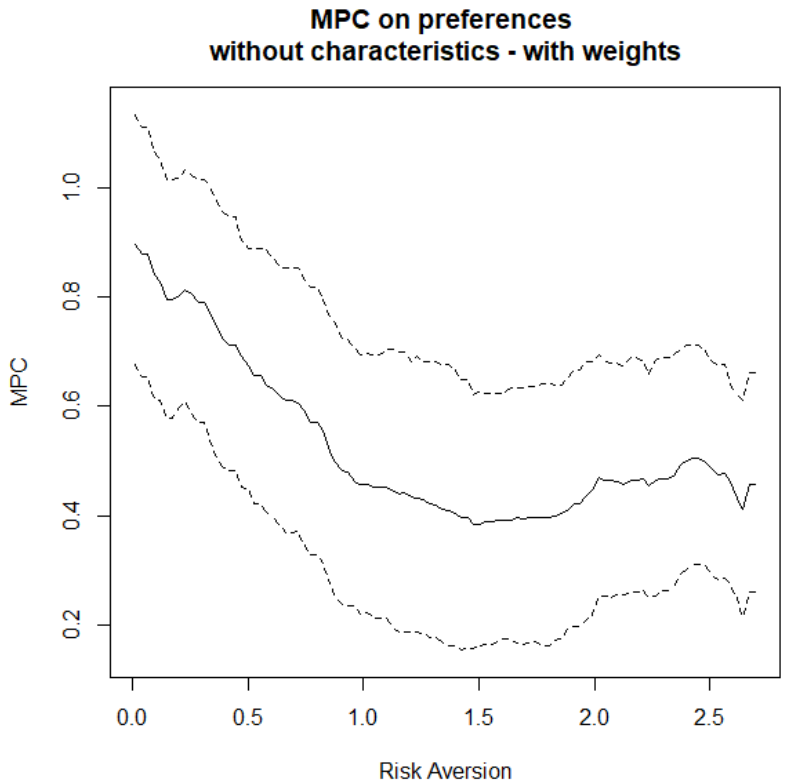

Figure: The marginal effect of risk aversion on MPC while aggregating over the other preferences.

With characteristics

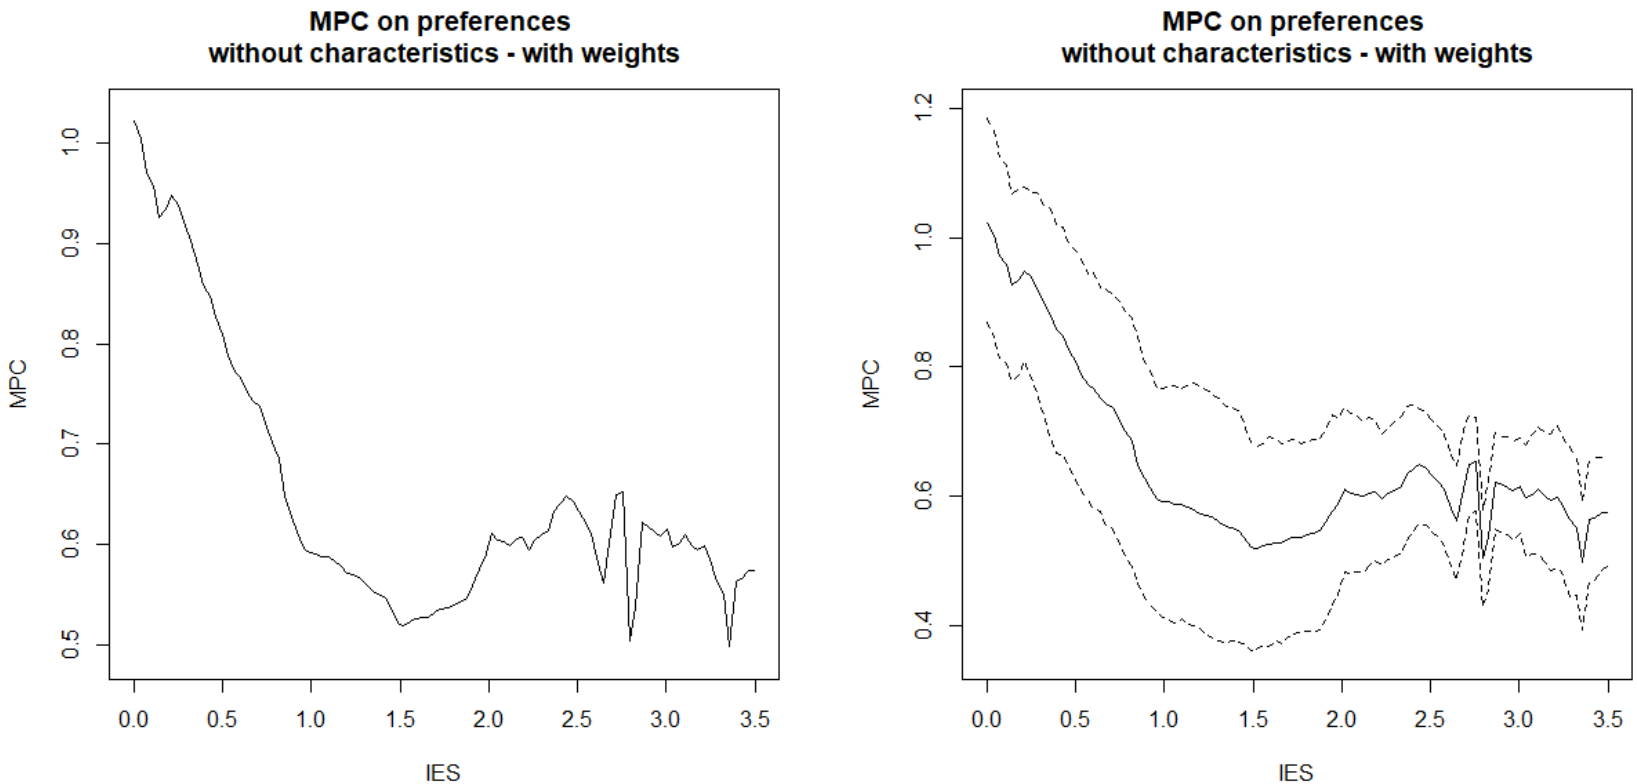

Figure: The marginal effect of intertemporal elasticity of substitution on MPC while aggregating over the other preferences.

One asset MPC

In a one asset model, the MPC expression simplifies to

$$\mu_i \quad = \quad \tau + \psi_i \left(\frac{1}{\beta_i} - 1 - \tau \right) - \gamma_i (1 - \psi_i) \frac{1}{2} \sigma_w^2 \tag{6}$$

- Given a declining consumption profile ($\frac{1}{\beta_i} - 1 - \tau > 0$), higher ψ increases the MPC.
- Higher γ decreases the MPC if $\psi < 1$.
- Indirect effects through correlation across preferences.

back