Discussion of *HANKSSON*

Monetary Policy Conference 2026
Banca D'Italia

Cristiano Cantore

Sapienza University of Rome

18/06/2026

Summary

- Does heterogeneity matter for aggregate demand?

Summary

- Does heterogeneity matter for aggregate demand?
  - Amplification vs. dampening
  - Should central banks care about heterogeneity?

Summary

- Does heterogeneity matter for aggregate demand?
  - Amplification vs. dampening
  - Should central banks care about heterogeneity?

- The distribution and evolution of consumption is crucial. This paper is about **model-consistent measurement** vs consumption imputation:
  - **Individual-level matched** Norwegian administrative data: income, wealth, consumption
  - Income shock: unemployment-induced (as opposed to lotteries)

Summary

- Does heterogeneity matter for aggregate demand?
  - Amplification vs. dampening
  - Should central banks care about heterogeneity?
- The distribution and evolution of consumption is crucial. This paper is about **model-consistent measurement** vs consumption imputation:
  - **Individual-level matched** Norwegian administrative data: income, wealth, consumption
  - Income shock: unemployment-induced (as opposed to lotteries)
- Two sufficient statistics:
  1. **A**MPC vs. $\overline{\text{MPC}}$

Summary

- Does heterogeneity matter for aggregate demand?
  - Amplification vs. dampening
  - Should central banks care about heterogeneity?
- The distribution and evolution of consumption is crucial. This paper is about **model-consistent measurement** vs consumption imputation:
  - **Individual-level matched** Norwegian administrative data: income, wealth, consumption
  - Income shock: unemployment-induced (as opposed to lotteries)
- Two sufficient statistics:
  1. **A**MPC vs. $\overline{\text{MPC}}$
  2. Cyclicality of consumption inequality: $\underbrace{\dfrac{dc_t^H}{dc_t}}_{\beta_c^H}$ vs. $\underbrace{\dfrac{dc_t^S}{dc_t}}_{\beta_c^S}$

Summary

- Does heterogeneity matter for aggregate demand?
  - Amplification vs. dampening
  - Should central banks care about heterogeneity?

- The distribution and evolution of consumption is crucial. This paper is about **model-consistent measurement** vs consumption imputation:
  - **Individual-level matched** Norwegian administrative data: income, wealth, consumption
  - Income shock: unemployment-induced (as opposed to lotteries)

- Two sufficient statistics:
  1. **A**MPC vs. $\overline{\text{MPC}}$
  2. Cyclicality of consumption inequality: $\underbrace{\frac{dc_t^H}{dc_t}}_{\beta_c^H}$ vs. $\underbrace{\frac{dc_t^S}{dc_t}}_{\beta_c^S}$

independent of MPCs

accounts for cyclical savings

My take

- Asking a first-order question.
- Important and long-lasting contribution to Normative and Positive macroeconomics!

My take

- Asking a first-order question.
- Important and long-lasting contribution to Normative and Positive macroeconomics!
- **Main results**:
  1. **A**MPC $\approx \overline{\text{MPC}}$: To analyze aggregate behavior, you do not need to keep track of heterogeneity! (Campbell and Mankiw 1989; Bilbiie 2008, 2020; Werning 2015)

My take

- Asking a first-order question.
- Important and long-lasting contribution to Normative and Positive macroeconomics!
- **Main results**:
  1. **A**MPC $\approx \overline{\text{MPC}}$: To analyze aggregate behavior, you do not need to keep track of heterogeneity! (Campbell and Mankiw 1989; Bilbiie 2008, 2020; Werning 2015)
  2. Disposable income is the key: *Automatic stabilizers do stabilize!*

My take

- Asking a first-order question.
- Important and long-lasting contribution to Normative and Positive macroeconomics!
- **Main results**:
  1. **A**MPC $\approx \overline{\text{MPC}}$: To analyze aggregate behavior, you do not need to keep track of heterogeneity! (Campbell and Mankiw 1989; Bilbiie 2008, 2020; Werning 2015)
  2. Disposable income is the key: *Automatic stabilizers do stabilize!*
  3. Consistent with recent literature on high micro MPCs and low aggregate multipliers (Orchard et al. 2025)

My take

- Asking a first-order question.
- Important and long-lasting contribution to Normative and Positive macroeconomics!
- **Main results**:
  1. $\mathbf{A}\text{MPC} \approx \overline{\text{MPC}}$: To analyze aggregate behavior, you do not need to keep track of heterogeneity! (Campbell and Mankiw 1989; Bilbiie 2008, 2020; Werning 2015)
  2. Disposable income is the key: *Automatic stabilizers do stabilize!*
  3. Consistent with recent literature on high micro MPCs and low aggregate multipliers (Orchard et al. 2025)
  4. You need to be careful when conducting policy analysis using HANK models calibrated on labor earnings!

Comments Overview

1. MPC estimation and unemployment shock
2. Household-level robustness
3. The extensive margin of heterogeneity
4. Minor comments/clarifications

I. MPC estimation: unemployment shock vs. lotteries

- The **A**MPC result relies on MPC heterogeneity estimated from unemployment-induced income changes.

I. MPC estimation: unemployment shock vs. lotteries

- The **A**MPC result relies on MPC heterogeneity estimated from unemployment-induced income changes.
- Unemployment is not randomly assigned:

I. MPC estimation: unemployment shock vs. lotteries

- The **A**MPC result relies on MPC heterogeneity estimated from unemployment-induced income changes.
- Unemployment is not randomly assigned:
  - It may be correlated with preferences, liquidity, occupation, and local labor demand.

I. MPC estimation: unemployment shock vs. lotteries

- The **A**MPC result relies on MPC heterogeneity estimated from unemployment-induced income changes.
- Unemployment is not randomly assigned:
  - It may be correlated with preferences, liquidity, occupation, and local labor demand.
- **Selection** could affect both:

I. MPC estimation: unemployment shock vs. lotteries

- The **A**MPC result relies on MPC heterogeneity estimated from unemployment-induced income changes.
- Unemployment is not randomly assigned:
  - It may be correlated with preferences, liquidity, occupation, and local labor demand.
- **Selection** could affect both:
  - The estimated MPC distribution.

I. MPC estimation: unemployment shock vs. lotteries

- The AMPC result relies on MPC heterogeneity estimated from unemployment-induced income changes.
- Unemployment is not randomly assigned:
  - It may be correlated with preferences, liquidity, occupation, and local labor demand.
- **Selection** could affect both:
  - The estimated MPC distribution.
  - Its covariance with individual income betas.

I. MPC estimation: unemployment shock vs. lotteries

- The **A**MPC result relies on MPC heterogeneity estimated from unemployment-induced income changes.
- Unemployment is not randomly assigned:
  - It may be correlated with preferences, liquidity, occupation, and local labor demand.
- **Selection** could affect both:
  - The estimated MPC distribution.
  - Its covariance with individual income betas.
- Concern: selection into unemployment, group-specific income persistence, and job-loss severity could mechanically drive part of the result.

I. MPC estimation: unemployment shock vs. lotteries

- The **A**MPC result relies on MPC heterogeneity estimated from unemployment-induced income changes.
- Unemployment is not randomly assigned:
  - It may be correlated with preferences, liquidity, occupation, and local labor demand.
- **Selection** could affect both:
  - The estimated MPC distribution.
  - Its covariance with individual income betas.
- Concern: selection into unemployment, group-specific income persistence, and job-loss severity could mechanically drive part of the result.
- Suggestion: If available, report Y and C betas by industry of occupation and/or local labor-market exposure.

II. Household-level robustness

- The main analysis is individual-level, but consumption and wealth are measured at the household level and, for couples, assigned as one half of the household total.

II. Household-level robustness

- The main analysis is individual-level, but consumption and wealth are measured at the household level and, for couples, assigned as one half of the household total.
- **Household composition and within-household sharing** may matter for:

II. Household-level robustness

- The main analysis is individual-level, but consumption and wealth are measured at the household level and, for couples, assigned as one half of the household total.
- **Household composition and within-household sharing** may matter for:
  - HtM status and liquid-wealth measures.

II. Household-level robustness

- The main analysis is individual-level, but consumption and wealth are measured at the household level and, for couples, assigned as one half of the household total.
- **Household composition and within-household sharing** may matter for:
  - HtM status and liquid-wealth measures.
  - Flat consumption betas across groups.

II. Household-level robustness

- The main analysis is individual-level, but consumption and wealth are measured at the household level and, for couples, assigned as one half of the household total.
- **Household composition and within-household sharing** may matter for:
  - HtM status and liquid-wealth measures.
  - Flat consumption betas across groups.
- Some household-level checks are already in the paper:

II. Household-level robustness

- The main analysis is individual-level, but consumption and wealth are measured at the household level and, for couples, assigned as one half of the household total.
- **Household composition and within-household sharing** may matter for:
  - HtM status and liquid-wealth measures.
  - Flat consumption betas across groups.
- Some household-level checks are already in the paper:
  - Income betas under the tax-transfer counterfactual.

II. Household-level robustness

- The main analysis is individual-level, but consumption and wealth are measured at the household level and, for couples, assigned as one half of the household total.
- **Household composition and within-household sharing** may matter for:
  - HtM status and liquid-wealth measures.
  - Flat consumption betas across groups.
- Some household-level checks are already in the paper:
  - Income betas under the tax-transfer counterfactual.
  - The BPP-Commault robustness uses a household sample.

II. Household-level robustness

- The main analysis is individual-level, but consumption and wealth are measured at the household level and, for couples, assigned as one half of the household total.
- **Household composition and within-household sharing** may matter for:
  - HtM status and liquid-wealth measures.
  - Flat consumption betas across groups.
- Some household-level checks are already in the paper:
  - Income betas under the tax-transfer counterfactual.
  - The BPP-Commault robustness uses a household sample.
- A useful extension would be to replicate the main sufficient statistics at the household level throughout:

II. Household-level robustness

- The main analysis is individual-level, but consumption and wealth are measured at the household level and, for couples, assigned as one half of the household total.
- **Household composition and within-household sharing** may matter for:
  - HtM status and liquid-wealth measures.
  - Flat consumption betas across groups.
- Some household-level checks are already in the paper:
  - Income betas under the tax-transfer counterfactual.
  - The BPP-Commault robustness uses a household sample.
- A useful extension would be to replicate the main sufficient statistics at the household level throughout:
  - Table 1: **A**MPC and $\mathrm{Cov}(MPC, \beta)$.

II. Household-level robustness

- The main analysis is individual-level, but consumption and wealth are measured at the household level and, for couples, assigned as one half of the household total.
- **Household composition and within-household sharing** may matter for:
  - HtM status and liquid-wealth measures.
  - Flat consumption betas across groups.
- Some household-level checks are already in the paper:
  - Income betas under the tax-transfer counterfactual.
  - The BPP-Commault robustness uses a household sample.
- A useful extension would be to replicate the main sufficient statistics at the household level throughout:
  - Table 1: **A**MPC and $\mathrm{Cov}(MPC, \beta)$.
  - Table 2: consumption betas by HtM status.

II. Household-level robustness

- The main analysis is individual-level, but consumption and wealth are measured at the household level and, for couples, assigned as one half of the household total.
- **Household composition and within-household sharing** may matter for:
  - HtM status and liquid-wealth measures.
  - Flat consumption betas across groups.
- Some household-level checks are already in the paper:
  - Income betas under the tax-transfer counterfactual.
  - The BPP-Commault robustness uses a household sample.
- A useful extension would be to replicate the main sufficient statistics at the household level throughout:
  - Table 1: **A**MPC and $\mathrm{Cov}(MPC, \beta)$.
  - Table 2: consumption betas by HtM status.
- This would also address transfers and benefits that are effectively received or pooled at the household level.

III. SS with cyclical HtM

- In the paper's fixed-share TANK benchmark:

$$\hat{c}_t = E_t\hat{c}_{t+1} - \tilde{\lambda}_c(\hat{\gamma}_t - E_t\hat{\gamma}_{t+1}) - \sigma r_t, \qquad \hat{\gamma}_t \equiv \hat{c}_t^S - \hat{c}_t^H \ \& \ \Gamma \equiv \frac{C_S}{C_H}$$

- A one-time real-rate cut implies:

$$\frac{d\hat{c}_t}{d(-r_t)} = \sigma\frac{1}{1+\tilde{\lambda}_c\gamma_c}, \qquad \gamma_c \equiv \frac{d\hat{\gamma}_t}{d\hat{c}_t}$$

III. SS with cyclical HtM

- In the paper's fixed-share TANK benchmark:

$$\hat{c}_t = E_t\hat{c}_{t+1} - \tilde{\lambda}_c(\hat{\gamma}_t - E_t\hat{\gamma}_{t+1}) - \sigma r_t, \qquad \hat{\gamma}_t \equiv \hat{c}^S_t - \hat{c}^H_t \ \&\ \Gamma \equiv \frac{C_S}{C_H}$$

- A one-time real-rate cut implies:

$$\frac{d\hat{c}_t}{d(-r_t)} = \sigma\frac{1}{1+\tilde{\lambda}_c\gamma_c}, \qquad \gamma_c \equiv \frac{d\hat{\gamma}_t}{d\hat{c}_t}$$

- With countercyclical idiosyncratic risk and HtM share:

$$\hat{\lambda}_t = \rho\hat{\lambda}_{t-1} - \eta\hat{c}_t + e_t,$$

$$\hat{c}_t = \Omega E_t\hat{c}_{t+1} - \tilde{\lambda}_c(\hat{\gamma}_t - (1+\frac{\tilde{s}_\gamma - 1}{\tilde{\lambda}_c})E_t\hat{\gamma}_{t+1}) - \tilde{\lambda}_c(\Gamma-1)(1-\rho)\hat{\lambda}_t - \sigma r_t.$$

III. SS with cyclical HtM

- In the paper's fixed-share TANK benchmark:

$$\hat{c}_t = E_t\hat{c}_{t+1} - \tilde{\lambda}_c(\hat{\gamma}_t - E_t\hat{\gamma}_{t+1}) - \sigma r_t, \qquad \hat{\gamma}_t \equiv \hat{c}_t^S - \hat{c}_t^H \ \& \ \Gamma \equiv \frac{C_S}{C_H}$$

- A one-time real-rate cut implies:

$$\frac{d\hat{c}_t}{d(-r_t)} = \sigma\frac{1}{1+\tilde{\lambda}_c\gamma_c}, \qquad \gamma_c \equiv \frac{d\hat{\gamma}_t}{d\hat{c}_t}$$

- With countercyclical idiosyncratic risk and HtM share:

$$\hat{\lambda}_t = \rho\hat{\lambda}_{t-1} - \eta\hat{c}_t + e_t,$$

$$\hat{c}_t = \boxed{\Omega} E_t\hat{c}_{t+1} - \tilde{\lambda}_c(\hat{\gamma}_t - (1 + \boxed{\frac{\tilde{s}_\gamma - 1}{\tilde{\lambda}_c}})E_t\hat{\gamma}_{t+1}) - \boxed{\tilde{\lambda}_c(\Gamma - 1)(1-\rho)}\hat{\lambda}_t - \sigma r_t.$$

Appendix derivation

III. The extensive margin of inequality

Cantore, Guerriero & Nisticò (2026): The Dark Side of HANK

- For the same one-time cut:

$$\frac{d\hat{c}_t}{d(-r_t)} = \sigma \frac{1}{1+\tilde{\lambda}_c\gamma_c+\tilde{\lambda}_c(\Gamma-1)(1-\rho)\ell_\lambda}, \qquad \ell_\lambda \equiv \frac{d\hat{\lambda}_t}{d\hat{c}_t}$$

- Therefore amplification requires:

$$-1 < \underbrace{\tilde{\lambda}_c\gamma_c}_{\textit{intensive margin}} + \underbrace{\boxed{\tilde{\lambda}_c(\Gamma-1)(1-\rho)\ell_\lambda}}_{\textit{extensive margin}} < 0$$

- A countercyclical HtM share, $\ell_\lambda < 0$, **amplifies** the effect of the rate cut.
  - Γ (> 1) amplifies.
  - ρ $(\in (0, 1))$ dampens.

III. The extensive margin of inequality

Cantore, Guerriero & Nisticò (2026): The Dark Side of HANK

- For the same one-time cut:

$$\frac{d\hat{c}_t}{d(-r_t)} = \sigma \frac{1}{1 + \tilde{\lambda}_c \gamma_c + \tilde{\lambda}_c (\Gamma - 1)(1 - \rho)\ell_\lambda}, \qquad \ell_\lambda \equiv \frac{d\hat{\lambda}_t}{d\hat{c}_t}$$

- Therefore amplification requires:

$$-1 < \underbrace{\tilde{\lambda}_c \gamma_c}_{\textit{intensive margin}} + \underbrace{\boxed{\tilde{\lambda}_c (\Gamma - 1)(1 - \rho)\ell_\lambda}}_{\textit{extensive margin}} < 0$$

- A countercyclical HtM share, $\ell_\lambda < 0$, **amplifies** the effect of the rate cut.
  - Γ (> 1) amplifies.
  - ρ ($\in (0, 1)$) dampens.
- Can your data measure this? Consumption is observed at high frequency, but wealth-based HtM status appears annual. An annual time series of the HtM share could be informative.

Minor comments/clarifications

1. What if MPCs are state dependent/time-varying?
   * Can you estimate the MPC distribution separately in high- and low-unemployment years, or interact the unemployment-induced income shock with aggregate/local unemployment?
   * If MPCs are countercyclical, does the **A**MPC and multiplier conclusion remain close to neutral in recessions?

Minor comments/clarifications

1. What if MPCs are state dependent/time-varying?
   * Can you estimate the MPC distribution separately in high- and low-unemployment years, or interact the unemployment-induced income shock with aggregate/local unemployment?
   * If MPCs are countercyclical, does the **A**MPC and multiplier conclusion remain close to neutral in recessions?

2. You argue that transfers, more than taxes, flatten income betas. Can you decompose them by unemployment insurance, pensions/disability, child benefits, and other transfers? This would make the automatic-stabilizer mechanism more persuasive.

Thank you!

Appendix: transitions and groups

- Two groups: savers S and hand-to-mouth households H.
- **Cyclical and independent** transition probabilities:

$$\hat{s}_t = \eta_S \hat{c}_t + e_{S,t} \qquad \hat{h}_t = -\eta_H \hat{c}_t + e_{H,t}$$

  where s_t is the probability of remaining a saver and h_t is the probability of remaining HtM.
- Law of motion for the HtM population share λ_t:

$$\lambda_t = h_t \lambda_{t-1} + (1 - s_t)(1 - \lambda_{t-1}) = (s_t + h_t - 1)\lambda_{t-1} + 1 - s_t.$$

- Log-linearized share dynamics:

$$\hat{\lambda}_t = \rho \hat{\lambda}_{t-1} - \eta \hat{c}_t + e_t,$$

  where

$$\rho = s + h - 1 = 1 - \frac{1-s}{\lambda} \in (0,1), \qquad \eta = \frac{(1-\lambda)s}{\lambda}\eta_S + h\eta_H > 0, \qquad e_t = he_{H,t} - \frac{(1-\lambda)s}{\lambda}e_{S,t}.$$

Back

Appendix: saver Euler

- Saver Euler with idiosyncratic transition risk:

$$C_{S,t}^{-1/\sigma} = \beta(1+i_t)E_t\left\{\frac{s_{t+1}C_{S,t+1}^{-1/\sigma} + (1-s_{t+1})C_{H,t+1}^{-1/\sigma}}{1+\pi_{t+1}}\right\}.$$

- Define $\hat{\gamma}_t \equiv \hat{c}_{S,t} - \hat{c}_{H,t}$, $r_t \equiv \hat{i}_t - E_t\hat{\pi}_{t+1}$, and

$$s_\gamma \equiv s + (1-s)\Gamma^{1/\sigma}, \qquad \tilde{s}_\gamma \equiv \frac{s}{s_\gamma}.$$

- Log-linear saver Euler:

$$\hat{c}_{S,t} = E_t\hat{c}_{S,t+1} - (1-\tilde{s}_\gamma)E_t\hat{\gamma}_{t+1} - \sigma r_t + \tilde{s}_\gamma\sigma\left(\Gamma^{1/\sigma} - 1\right)E_t\hat{s}_{t+1}.$$

Back

Appendix: aggregation

- $C_{H,t} = (1 - \tau)WN_{H,t}$ and $C_t = Y_t$
- Aggregate consumption:

$$C_t = (1 - \lambda_t)C_{S,t} + \lambda_t C_{H,t}.$$

- Define the steady-state HtM consumption share and relative consumption:

$$\tilde{\lambda}_c \equiv \lambda \frac{C_H}{C}, \qquad \Gamma \equiv \frac{C_S}{C_H}.$$

- Log-linear aggregation:

$$\hat{c}_t = (1 - \tilde{\lambda}_c)\hat{c}_{S,t} + \tilde{\lambda}_c \hat{c}_{H,t} - \tilde{\lambda}_c(\Gamma - 1)\hat{\lambda}_t.$$

- Using $\hat{\gamma}_t = \hat{c}_{S,t} - \hat{c}_{H,t}$:

$$\hat{c}_{S,t} = \hat{c}_t + \tilde{\lambda}_c \hat{\gamma}_t + \tilde{\lambda}_c(\Gamma - 1)\hat{\lambda}_t.$$

Back

Appendix: aggregate Euler

- Substitute the aggregation identity into the saver Euler:

$$\begin{aligned}\hat{c}_t = E_t\hat{c}_{t+1} - \tilde{\lambda}_c\hat{\gamma}_t + \left[\tilde{s}_\gamma - (1-\tilde{\lambda}_c)\right]E_t\hat{\gamma}_{t+1} \\ -\tilde{\lambda}_c(\Gamma-1)(\hat{\lambda}_t - E_t\hat{\lambda}_{t+1}) - \sigma r_t + \tilde{s}_\gamma\sigma\left(\Gamma^{1/\sigma}-1\right)E_t\hat{s}_{t+1}.\end{aligned}$$

- The cyclical transition probabilities imply

$$E_t\hat{s}_{t+1} = \eta_S E_t\hat{c}_{t+1},$$

and

$$E_t\hat{\lambda}_{t+1} = \rho\hat{\lambda}_t - \eta E_t\hat{c}_{t+1}.$$

Back

Appendix: aggregate Euler

- Hence:

$$\hat{c}_t = \Omega E_t \hat{c}_{t+1} - \tilde{\lambda}_c \left(\hat{\gamma}_t - \left(1 + \frac{\tilde{s}_\gamma - 1}{\tilde{\lambda}_c} \right) E_t \hat{\gamma}_{t+1} \right) - \tilde{\lambda}_c (\Gamma - 1)(1 - \rho) \hat{\lambda}_t - \sigma r_t.$$

- where

$$\Omega = 1 - \tilde{\lambda}_c (\Gamma - 1)\eta + \tilde{s}_\gamma \sigma \left(\Gamma^{1/\sigma} - 1 \right) \eta_S.$$

Back