Economic Agents as Imperfect Problem Solvers

Cosmin Ilut   Rosen Valchev

Duke & NBER   *Boston College*

NBER SI 2020 Impulse and Propagation Mechanisms

Motivation

- Reasoning is costly – requires thought and introspection
  - solutions to math problems not immediately obvious, require thought
- Evidence: quality of decision-making varies
  - increases in effort/time spent deliberating
- Long-standing question of bounded rationality
  - behavioral economics: mistakes in behavior given beliefs about state
  - procedurally rational (Simon, '76): basic cost-benefit tradeoff of reasoning

This paper

- A tractable and intuitive model of procedurally rational reasoning
- Signal-extraction within a quadratic tracking framework

$$\min_{\hat{c}_t} \mathbb{E}_t(\hat{c}_t - c^*(y_t))^2$$

  - Unknown policy function $c^*(.)$, known y_t
  - Stochastic choice & local learning: experimental & neuro-science lit
- State and history dependent uncertainty and reasoning
  - Business-as-usual vs salient thinking

Workhorse Laboratory: Consumption-Savings problem

Endogenous state (assets) interacts with state-dependent reasoning

- Feedback between reasoning, actions/mistakes and asset evolution

$\Rightarrow$ distribution of mistakes 'matters'

1. across agents – idiosyncratic reasoning errors do not wash-out
   - ⋆ selection effect amplifies aggregate shocks
2. within history of agent – settle in "learning traps"

Properties of ergodic distribution: challenging for fully rational model

- High local MPCs, including for rich agents
- Large and persistent inequality, with trapped hand-to-mouth

Outline

- **General Framework**
  - tractable for us as analysts to describe basic reasoning features
  - intuitive implications: state and history dependency
- Consumption-Savings application
  - illustrates feedback between reasoning and endogenous state evolution
  - distribution of mistakes matter both across and within agents
  - examine joint stationary distribution of assets and beliefs

General Framework

- Optimal policy $c^*(.)$ is unknown, Bayesian non-parametric learning
- Gaussian Process distribution prior: for any pair of state values y, y':

$$c^*(\begin{bmatrix} c^*(y) \\ c^*(y') \end{bmatrix}) \sim N\left(\begin{bmatrix} c_0(y) \\ c_0(y') \end{bmatrix}, \begin{bmatrix} \sigma_0(y,y) & \sigma_0(y,y') \\ \sigma_0(y,y') & \sigma(y',y') \end{bmatrix}\right)$$

  - prior is centered around the truth

$$c_0(y) = c^{Rational}(y)$$

  - covariance function: encodes beliefs about likely function shapes

$$\sigma_0(y,y') \equiv Cov(c^*(y), c^*(y'))$$

- Equivalently, Bayesian non-parametric kernel regression:

$$c^*(y) = \sum_k \theta_k \phi_k(y) \ ; \qquad \theta_k \sim N(\mu_k, \sigma_c^2)$$

Details

Reasoning process

- As analysts, we remain agnostic about specific reasoning process
  - Model it via a signal-extraction framework that captures key trade-offs

- Agents deliberate about the best course of action today
  1. Reason about optimal action today

$$\eta(y_t) = c^*(y_t) + \varepsilon_t,\ \varepsilon_t \sim N(0, \sigma^2_{\eta,t})$$

     ⋆ no objective info as observed by econometrician (Arragones et al. 2005)

  2. Use a covariance function declining in distance $||y - y'||$:

$$\sigma_0(y, y') = \sigma_c^2 exp(-\psi(y - y')^2)$$

  ⇒ uncertainty is *state-dependent*

- 'As if' agents use "online" ("solve-as-you-go") solution method
  - in-line with experiments & neuro-science (also machine learning)

Agent's problem

$$\min_{\sigma^2_{\eta,t},\hat{c}_t} \underbrace{E_t(\hat{c}_t - c^*(y_t))^2}_{=\hat{\sigma}^2_t(y_t)} + \kappa \underbrace{\ln\left[\frac{\hat{\sigma}^2_{t-1}(y_t)}{\hat{\sigma}^2_t(y_t)}\right]}_{\text{Entropy Cost}}$$

$$s.t.\ \hat{\sigma}^2_t(y_t) \leq \hat{\sigma}^2_{t-1}(y_t)$$

1. Action: $\hat{c}_t = E_t(c^*(y_t))$
2. Reasoning: choose $\sigma^2_{\eta,t}$ so posterior variance

$$\hat{\sigma}^2_t(y_t) = \min\left\{\kappa, \hat{\sigma}^2_{t-1}(y_t)\right\}$$

$\rightarrow$ resulting signal-to-noise ratio is state and history dependent

$$\alpha_t(y_t; y^{t-1}) = \max\left\{1 - \frac{\kappa}{\hat{\sigma}^2_{t-1}(y_t)}, 0\right\}$$

$\rightarrow$ effective action

$$\hat{c}_t(y_t) = \hat{c}_{t-1}(y_t) + \alpha_t(y_t; y^{t-1})(\eta_t - \hat{c}_{t-1}(y_t))$$

Outline

- General Framework
  - tractable for us as analysts to describe basic reasoning features
  - local reduction in uncertainty & cost-benefit tradeoff
  - intuitive implications: state and history dependency
- **Consumption-Savings application**
  - illustrates feedback between reasoning and endogenous state evolution
  - distribution of mistakes matter both across and within agents
  - examine joint stationary distribution of assets and beliefs

Standard savings problem (in an Aiyagari economy)

- Income states: endogenous state $\left(a_{i,t-1}\right)$ and exogenous $\left(s_{i,t}\right)$
  - fixed labor supply at one unit, earn wage w
  - with iid $s_{i,t}$, sufficient state is cash on hand: $y_{i,t} \equiv (1+r)a_{i,t-1} + ws_{i,t}$
  - choose $c_{i,t}$ and $a_{i,t}$ st: budget constraint & borrowing limit

$$a_{i,t} + c_{i,t} = y_{i,t} \quad \& \quad a_{i,t} \geq 0$$

- Heterogeneity:
  - ex-ante identical: same preferences and reasoning params
  - ex-post heterogeneous: idiosyncratic income $s_{i,t}$ & reasoning errors
    - ⋆ Feedback: reasoning/beliefs $|a_{i,t-1} \rightarrow \{c_{i,t}, a_{i,t}\} \rightarrow$ reasoning

- Aggregate production function $K^{\alpha}N^{1-\alpha}$; capital depreciation rate δ

$$r = \alpha K^{\alpha-1} - \delta; w = (1-\alpha)K^{\alpha}$$

Next

Feedback between state and reasoning: distribution of mistakes 'matters'

1. across agents – idiosyncratic reasoning errors do not wash-out
2. within history of agent – settle in "learning traps" with high MPC
3. Stationary equilibrium at joint distribution of assets and beliefs

$$K = \int_i a_i di$$

Evolution of beliefs: $t = 1$ (illustration)

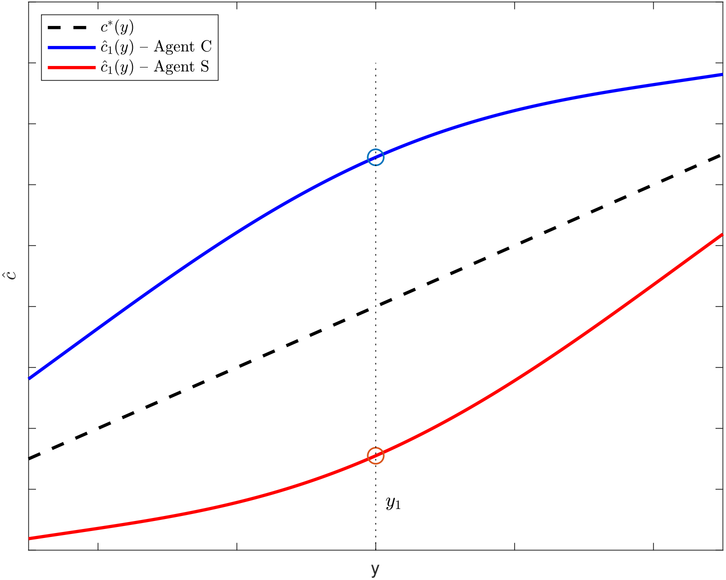

$t = 2$ State dependent uncertainty

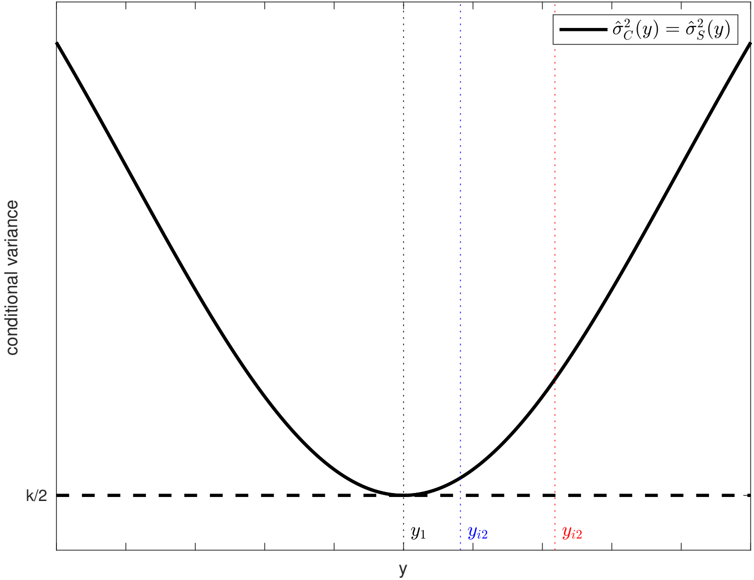

$t = 2$ Beliefs: aggregate effect of errors from selection

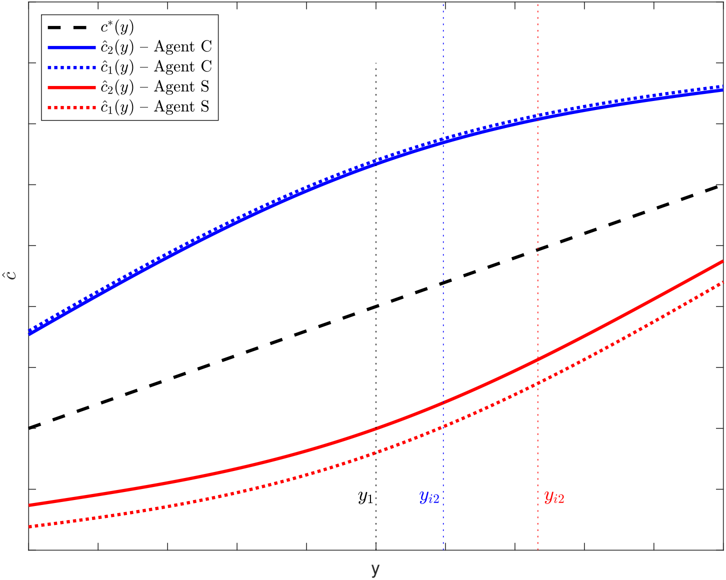

Low initial consumption reasoning signal

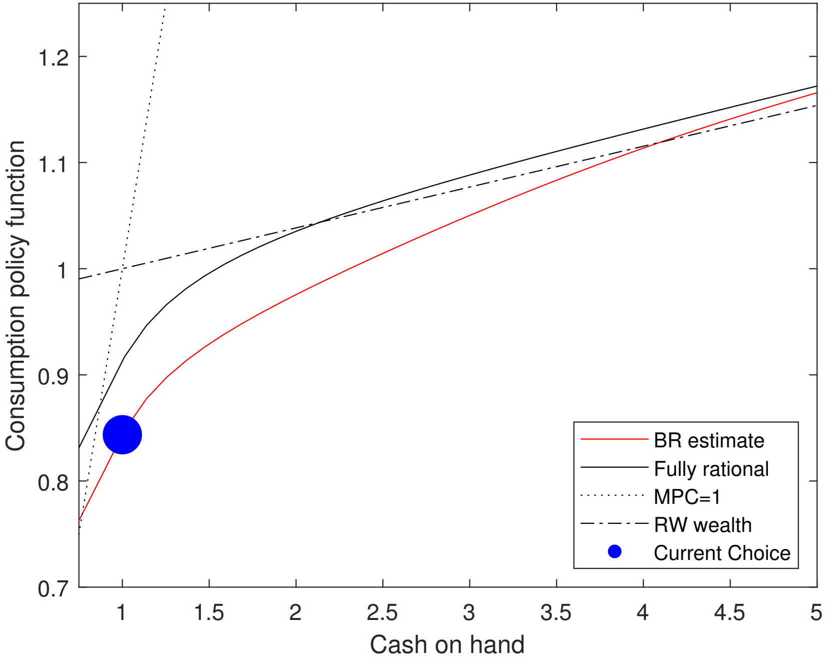

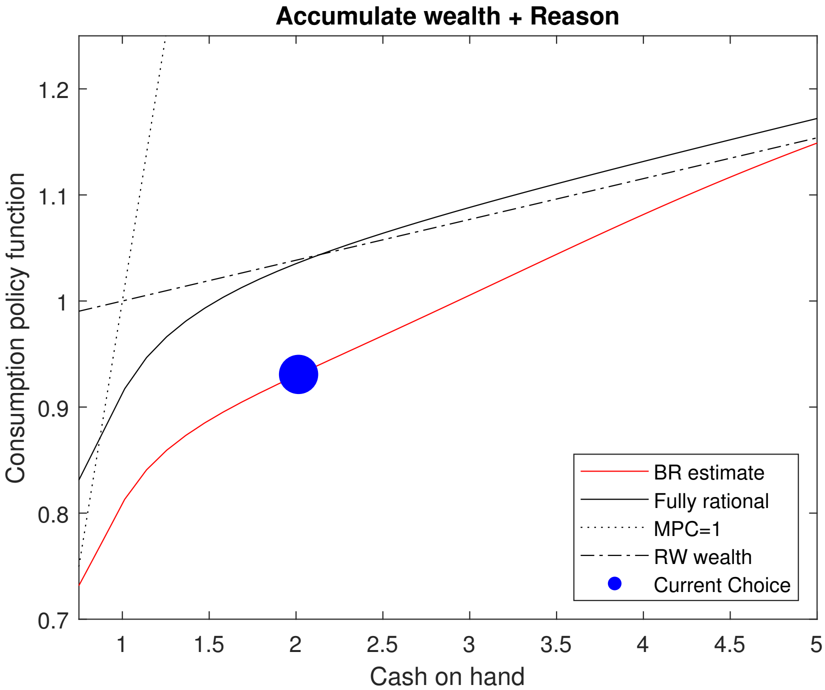
Accumulate wealth + Reason
Consumption policy function
Cash on hand
1.2
1.1
1
0.9
0.8
0.7
1
1.5
2
2.5
3
3.5
4
4.5
5
BR estimate
Fully rational
MPC=1
RW wealth
Current Choice

Accumulate wealth + Reason

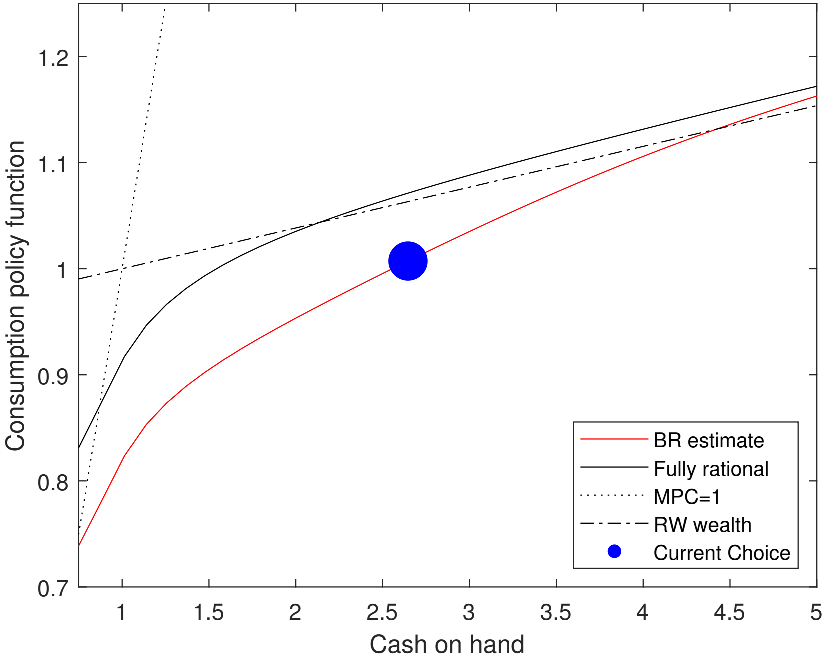

Accumulate wealth + Reason

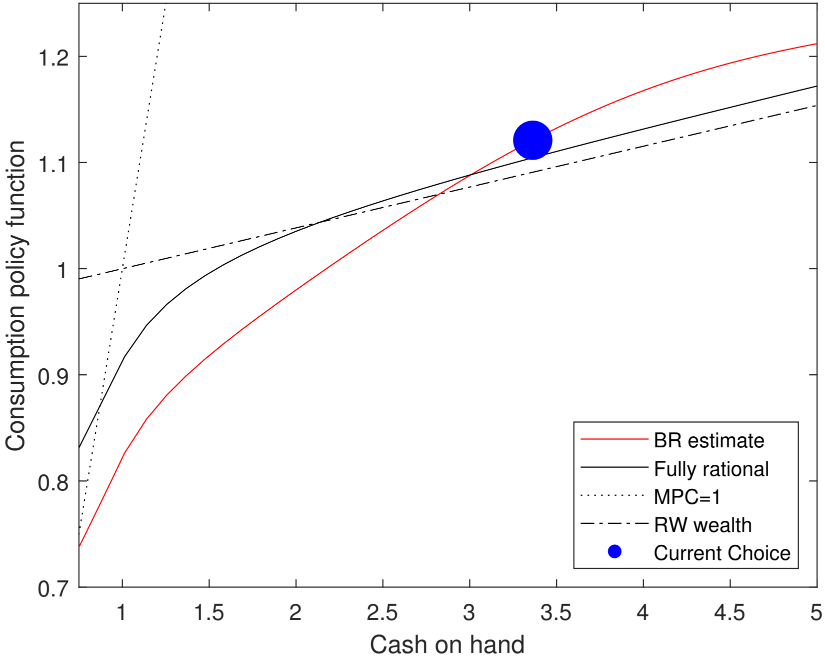

Locally stable wealth (no reasoning)

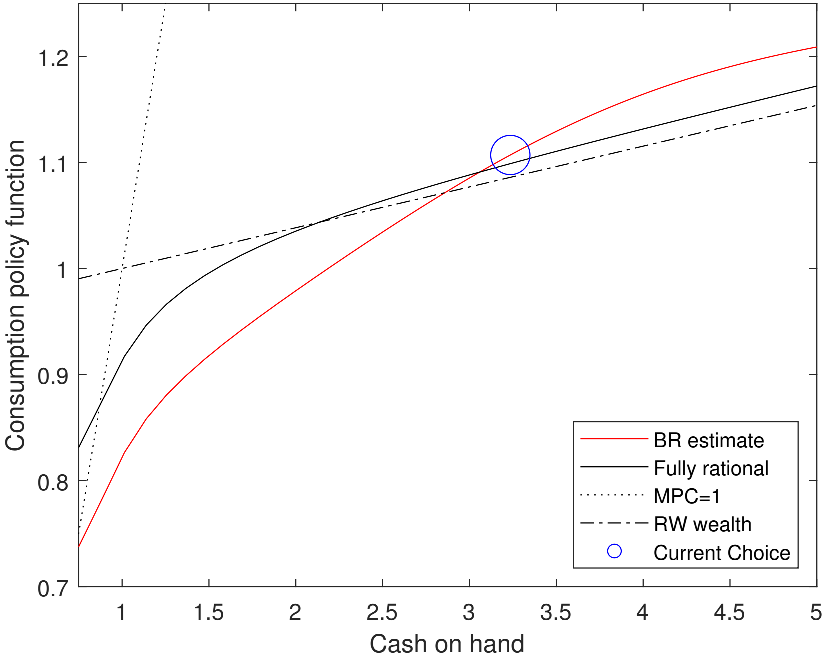

Locally stable wealth (no reasoning)

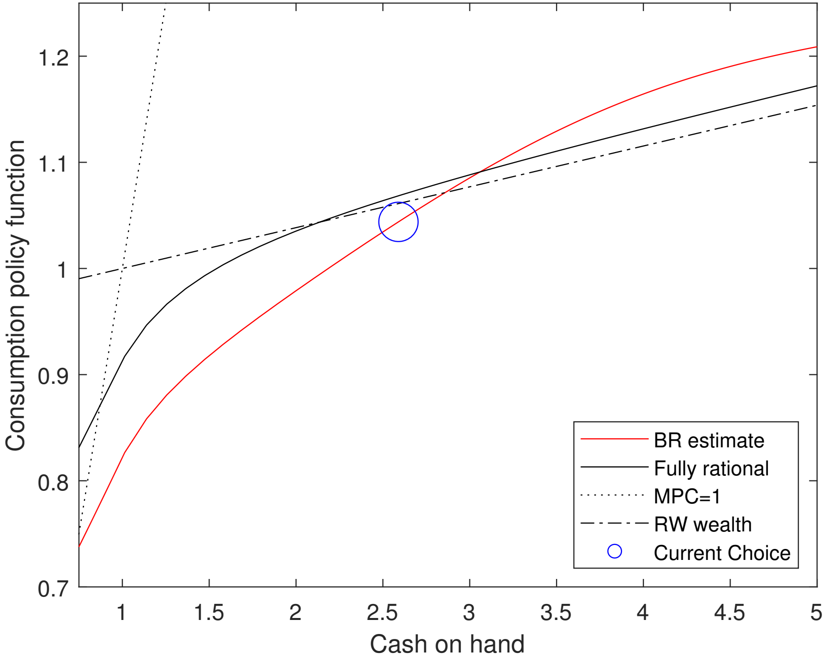

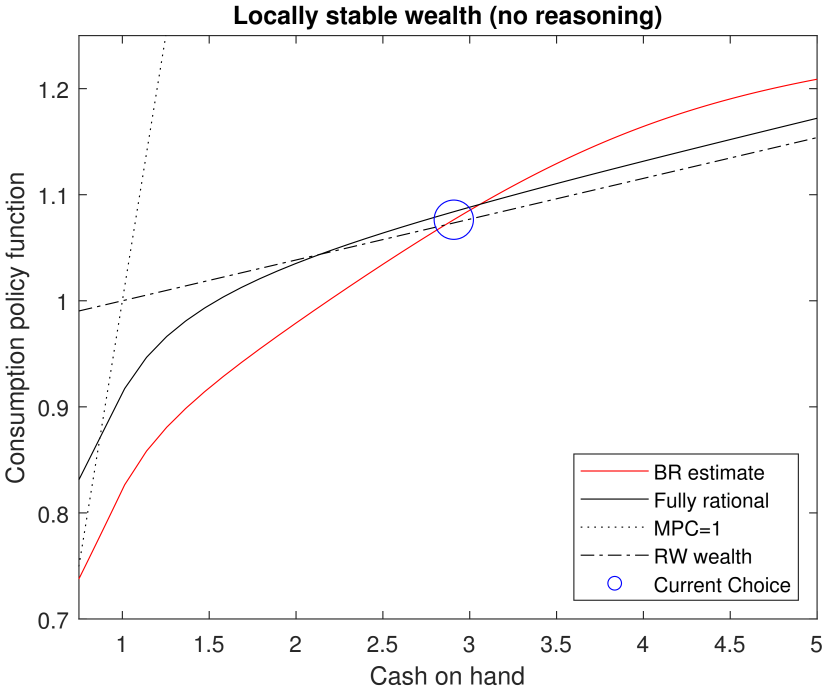
Locally stable wealth (no reasoning)
Consumption policy function
Cash on hand
1.2
1.1
1
0.9
0.8
0.7
1
1.5
2
2.5
3
3.5
4
4.5
5
BR estimate
Fully rational
MPC=1
RW wealth
Current Choice

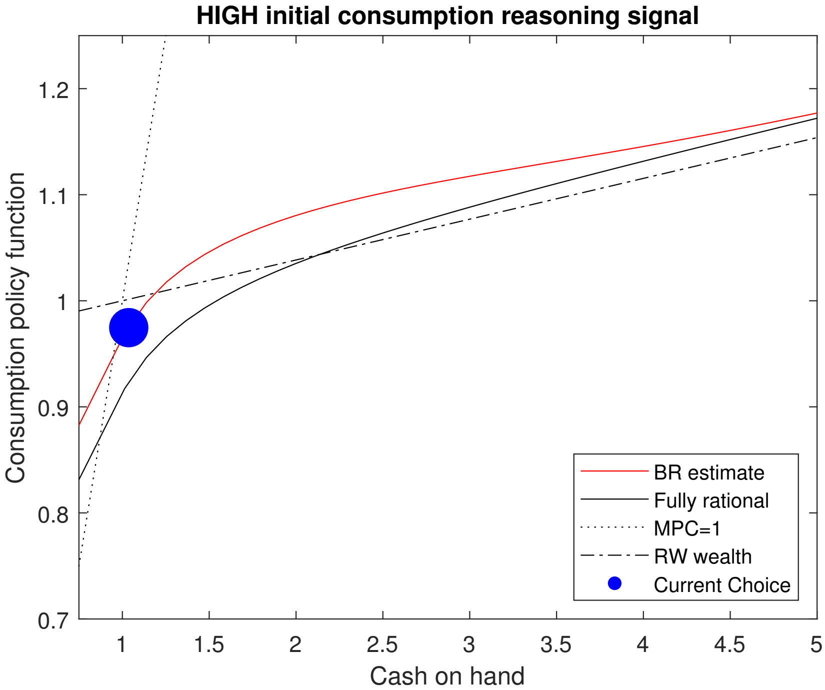
HIGH initial consumption reasoning signal
Consumption policy function
Cash on hand
1.2
1.1
1
0.9
0.8
0.7
1
1.5
2
2.5
3
3.5
4
4.5
5
BR estimate
Fully rational
MPC=1
RW wealth
Current Choice

Accumulate some wealth + Reason

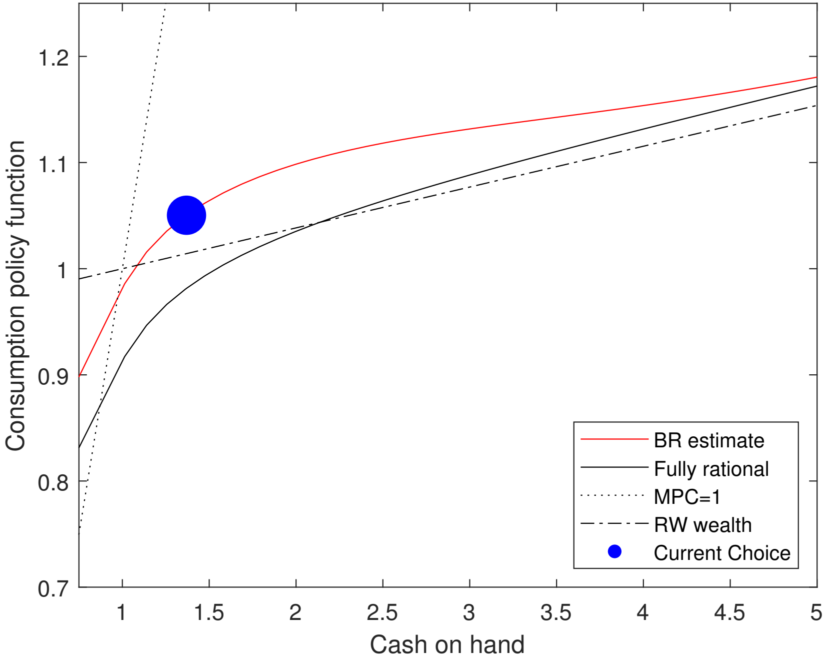

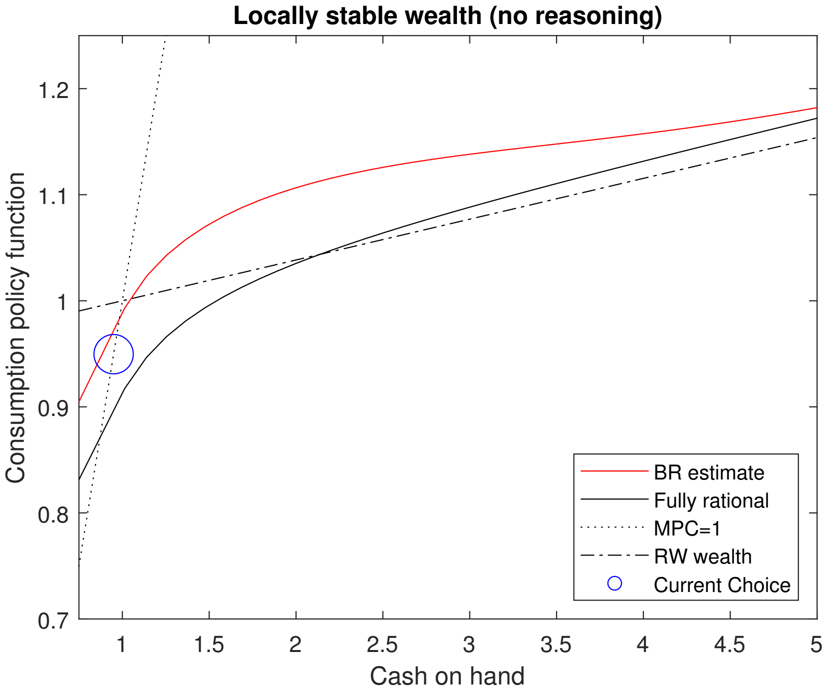
Locally stable wealth (no reasoning)
Consumption policy function
Cash on hand
1.2
1.1
1
0.9
0.8
0.7
1
1.5
2
2.5
3
3.5
4
4.5
5
BR estimate
Fully rational
MPC=1
RW wealth
Current Choice

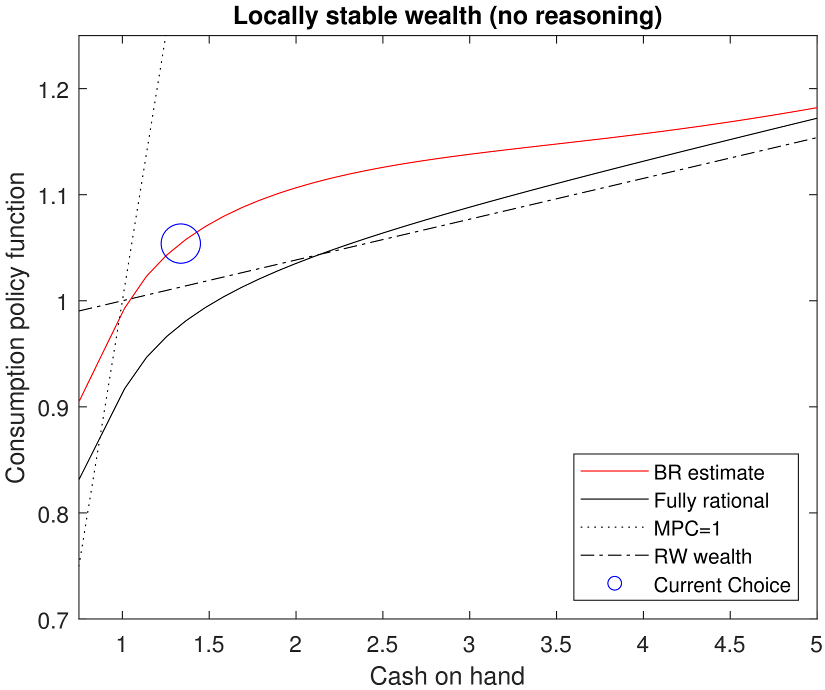

Locally stable wealth (no reasoning)
Consumption policy function
Cash on hand
0.7
0.8
0.9
1
1.1
1.2
1
1.5
2
2.5
3
3.5
4
4.5
5
BR estimate
Fully rational
MPC=1
RW wealth
Current Choice

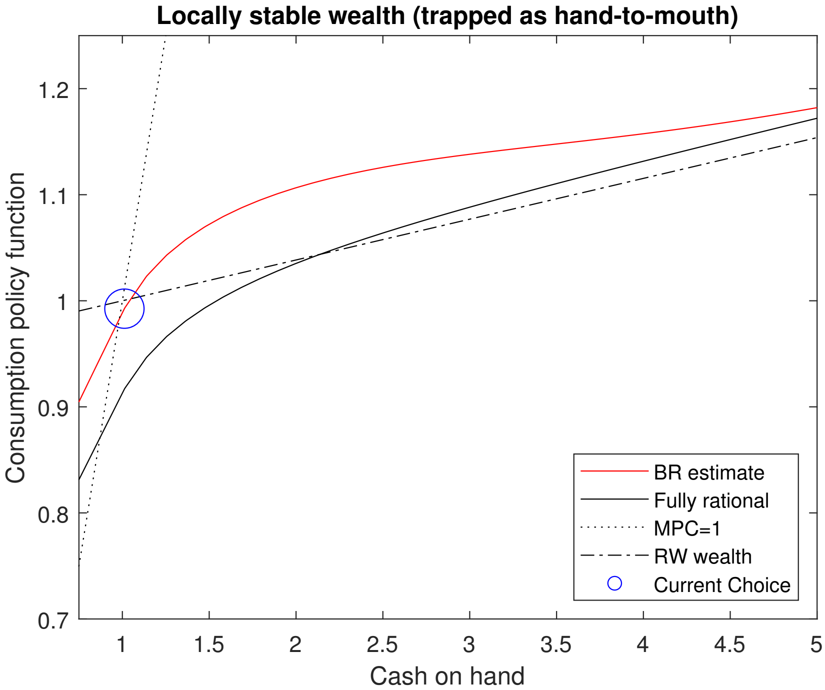

Locally stable wealth (trapped as hand-to-mouth)
Consumption policy function
Cash on hand
0.7
0.8
0.9
1
1.1
1.2
1
1.5
2
2.5
3
3.5
4
4.5
5
BR estimate
Fully rational
MPC=1
RW wealth
Current Choice

Stationary distribution

- Standard Aiyagari (1994)

$$u(c) = \ln(c), \quad \beta = 0.96, \quad \alpha = 0.36, \quad \delta = 0.08$$

$$\ln(s_{i,t}) \sim N(-\frac{\sigma_s^2}{2}, \sigma_s^2), \quad \sigma_s = 0.2$$

- Illustration with reasoning parameters

$$\kappa = 0.97; \ \sigma_c^2 = 0.74; \psi = 0.05; \theta = 0.02$$

1. Set $\{\sigma_c^2, \psi\}$ equal to an econometrician's estimates from simulated η_{it}
   - ⋆ solving this fixed-point restricts to model-consistent priors
2. With probability θ: iid shock that resets information to time 0 prior ($\theta > 0$: plausible & computationally needed for ergodicity)
3. Set κ so that bottom 20% asset share = 0% (respecting fp in $\{\sigma_c^2, \psi\}$)

Wealth distribution

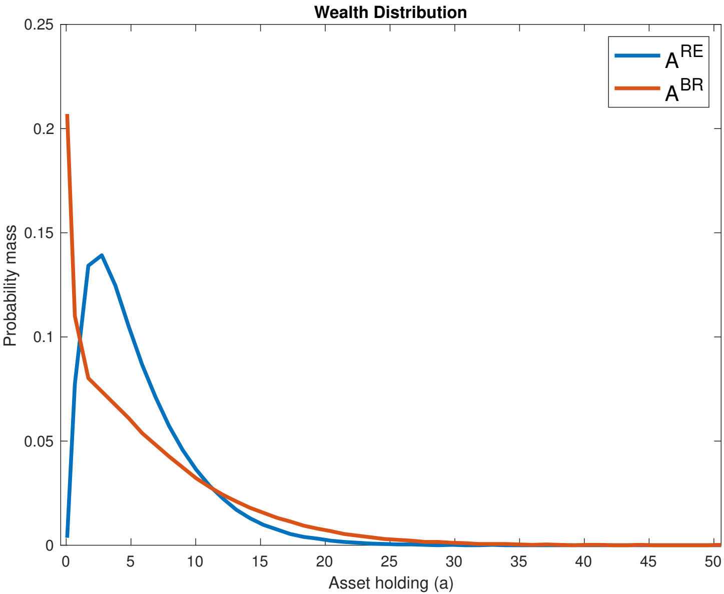

Key implications	Data	**Benchmark**	Rational	Low κ	Low ψ
INEQUALITY	PSID				
Gini coefficient	0.77	0.57	0.39	0.44	0.68
Hand-to-Mouth ($a_i \leq \frac{w}{6}$)	0.23	0.22	0.01	0.03	0.34
$Prob(HtM_{t+2} \mid HtM_t)$	0.65	0.90	0.36	0.60	0.91
$E(\Delta c_{t+2} \mid HtM_t)$	0.003	0.001	0.015	0.01	0.001
MPC	Lit				
Mean	0.2-0.6	0.28	0.05	0.14	0.32
Mean $\mid$ *not-HtM*	0.2-0.6	0.17	0.04	0.13	0.11
Mean $\mid$ top 20% a_i	0.2-0.6	0.15	0.04	0.12	0.04

Reference: PSID, Aguiar, Bils & Boar

Uncertainty shock (e.g. Covid-19, Great Recession ...)

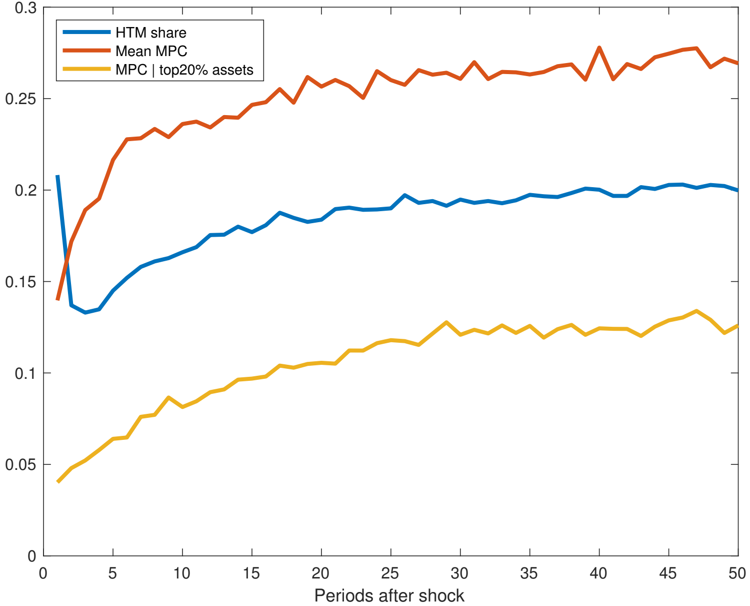

Conclusion

1. Tractable & intuitive model of costly reasoning about policy function

2. State and reasoning interaction: distribution of mistakes matters

Appendix: General Framework

- Generic recursive dynamic problem

$$V(y_t) = \max_{c_t} u(c_t) + \beta E_t(V(y_{t+1}))$$

subject to $$y_{t+1} = F(y_t, c_t, \varepsilon^y_{t+1})$$

- Typically infeasible to solve exactly, approx. via basis functions $\phi_k(y)$

$$c^*(y) = \sum_k \theta_k \phi_k(y)$$

- Economic agent as imperfect problem solver:
  - invest cognitive effort to reduce uncertainty about unknown θ_k
- Bounded but procedurally rational (Simon, 1976): trade-off
  - costly: more information contained in new reasoning signal
  - beneficial: lower posterior uncertainty

Bayesian learning: intuitive and tractable

Bayesian: prior over θ_k + (costly) reasoning signals $\rightarrow$ posterior belief

- Operationalize reduction in uncertainty over θ_k as signal-extraction
- Tractable: prior $\theta_k \sim N(\mu_k, \sigma_c^2)$ and $\lim_{k\rightarrow\infty}$

  $\Rightarrow$ Gaussian Process distrib. over space of $c^*(y)$: for any pair (y, y')

$$c^*\left(\begin{bmatrix} y \\ y' \end{bmatrix}\right) \sim N\left(\begin{bmatrix} c_0(y) \\ c_0(y') \end{bmatrix}, \begin{bmatrix} \sigma_0(y,y) & \sigma_0(y,y') \\ \sigma_0(y,y') & \sigma_0(y',y') \end{bmatrix}\right)$$

  - sequence of prior means μ_k for θ_k are chosen such that e.g.

$$c_0(y) \equiv E(c^*(y)) = c^{Rational}(y)$$

  - induced variance-covariance function

$$\sigma_0(y,y') \equiv Cov(c^*(y), c^*(y')) = \sigma_c^2 \int \phi_k(y)\phi_k(y')dk$$

Reasoning process

- Agent uses "solve-as-you-go" approach in solving problem
  - popular in machine-learning, in line with experiments & neuro-science

- Focus on characterizing $c^*(.)$ in neighborhood of y_t, means
  1. Use "local" basis functions ϕ_k, implying

$$\sigma_0(y, y') = \sigma_c^2 exp(-\psi(y - y')^2)$$

  2. Reason about optimal action today: solve for today's relevant θ_k's

$$\eta(y_t) = c^*(y_t) + \varepsilon_t, \ \varepsilon_t \sim N(0, \sigma^2_{\eta,t})$$

    - no objective info as observed by econometrician (Arragones et al. 2005)

- Recursive conditional expectations and variances, e.g.

$$\hat{c}_t(y) = \hat{c}_{t-1}(y) + \frac{\hat{\sigma}_{t-1}(y, y_t)}{\hat{\sigma}_{t-1}(y, y_t) + \sigma^2_{\eta,t}} \left[\eta_t - \hat{c}_{t-1}(y)\right]$$

Back