# Economic Agents as Imperfect Problem Solvers

Cosmin Ilut Rosen Valchev

*Duke & NBER* *Boston College*

$10^{th}$ International Workshop on Theoretical and Experimental Macroeconomics

June 26, 2019

# Motivation

- Reasoning is costly – requires thought and introspection
  - eg. solutions to math problems not immediately obvious, even though objective parameters are well defined and known
- Evidence: quality of decision-making
  - increases in effort/time spent deliberating (Hohnich et. al. (2017))
  - decreases in task complexity (Deck and Jahedi (2015))
  - increases in natural ability (D'Acunto et al. (2019a, 2019b)
- This paper: a tractable model of costly reasoning
  - Mistakes even when objective state is perfectly observed
  - Agnostic about particular deliberation process
    - Capture basic trade-off of reasoning effort vs accuracy
    - Information accumulation and spillover to similar situations

# Our paper

- Quadratic tracking problem $\Rightarrow$ effective action:

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

  - perfectly observed objective state variable $y_t \in \mathbb{R}$
  - unknown optimal policy function $c^*(y)$
- Have access to costly signals about optimal policy function $c^*(.)$
  - signal on $c^*(y_t)$ at $y_t$ is partially informative about $c^*(y)$ at other $y$
  - so uncertainty over best action is state and history dependent
- History and state dependent reasoning choice and effective action $\hat{c}_t$
  - reason less at usual state realizations (inertial behavior)
  - but more at unusual state realizations (salience & non-linearity of $\hat{c}_t$)
- Heterogeneity, even when controlling for observables
  - different persistent biases in inertial behavior due to histories
  - stochastic choice conditional on past

# Workhorse Laboratory: Consumption-Savings problem

- Endogenous state variable: emphasizes important interaction between history dependent reasoning choice and state evolution

- Feedback between beliefs, reasoning and asset (i.e. state) evolution
  - Agents tend to settle in "learning traps" $\Rightarrow$ inertia, endogenous habits
  - High local MPCs across the wealth distribution
  - Persistent inequality, fat-tailed wealth distribution

- Heterogeneity matters
  - Aggregate effects – errors do not wash-out (selection & systematic over-reaction)
    - state-dependent effect, no unconditional bias
  - Amplification – volatility and persistence due to aggregation effects

- Constrained optimal behavior

  $\Rightarrow$ responds to policy

# Literature

1. Imperfect actions as outcomes of procedural rationality
   - general principle: Simon (1976)
   - decision theory: Arragones et al. (2005), Ergin & Sarver (2010)
2. Limited attention resources
   - 'Rational inattention': Sims (1998, 2003), Matejka (2015)
   - Information choice: Woodford (2003), Reis (2006), Gabaix (2014)
3. Bayesian statistics
   - Gaussian processes in machine learning: Rasmussen & Williams (2006)
4. Evidence for reasoning as a friction in a consumption-savings problem
   - field data: cognitive ability (D'Acunto et al., 2019)
   - experimental: inertia and persistent mistakes (Khaw & Zorilla, 2018)
5. Challenging empirical properties of consumption responses:
   - Persistence: experience effects (Malmendier & Shen, 2018)
   - High MPC out of temporary shocks: Parker (2017), Olaffson & Pagel (2018), Ganong & Noel (2018)
   - Fat-tailed wealth distribution: De Nardi & Fella (2017)

# Outline

1. General learning framework
2. Consumption-Savings Application
3. 3 period example – analytics
4. Ergodic behavior – numerical solutions

# Basic Framework

- Quadratic tracking problem $\Rightarrow$ optimal action $c_t$:

$$c_t = \mathbb{E}(c^*(y_t))$$

- $c^*(y)$ function is unknown – Bayesian non-parametric learning
- Prior is a Gaussian Process distribution: for any $\mathbf{y} = [y_1, ..., y_N]'$

$$c^*(\mathbf{y}) \sim N\left( \begin{bmatrix} c_0(y_1) \\ \vdots \\ c_0(y_N) \end{bmatrix}, \begin{bmatrix} \sigma_0(y_1, y_1) & \dots & \sigma_0(y_1, y_N) \\ \vdots & \ddots & \vdots \\ \sigma_0(y_N, y_1) & \dots & \sigma_0(y_N, y_N) \end{bmatrix} \right)$$

  - e.g. prior is centered around the truth, state-by-state:

$$c_0(y) = c^*(y)$$

  - covariance function: decreasing correlation with distance

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

# Reasoning

- Deliberating about best course of action: obtain signals

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

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

- Recursive conditional expectation for best action at state $y$

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

# Optimal deliberation

- Myopic trade off

$$\min_{\sigma^2_{\eta,t}} V_{cc}\sigma^2_t(y_t) + \kappa \ln\left[\frac{\hat{\sigma}^2_{t-1}(y_t)}{\hat{\sigma}^2_t(y_t)}\right]$$
$$s.t.\ \sigma^2_t(y_t) \leq \sigma^2_{t-1}(y_t)$$

- Optimal solution: choose $\sigma^2_{\eta,t}$ so posterior variance

$$\sigma^{*2}_t(y_t) = \min\left\{\frac{\kappa}{V_{cc}}, \sigma^2_{t-1}(y_t)\right\}, \forall y_t$$

- Resulting signal-to-noise ratio is state and history dependent

$$\alpha^*_t(y_t; \eta_t, \eta^{t-1}) = \max\left\{1 - \frac{\kappa/V_{cc}}{\sigma^2_{t-1}(y_t)}, 0\right\}$$

- Leading to the effective action

$$\hat{c}_t(y_t) = \hat{c}_{t-1}(y_t) + \alpha^*_t(y_t; \eta_t, \eta^{t-1})(\eta_t - \hat{c}_{t-1}(y_t))$$

# Consumption Savings Application

- State dependent reasoning interacts with *endogenous* state evolution
- Workhorse laboratory: standard consumption savings problem
- Assuming quadratic utility $u(.)$, $\beta = \frac{1}{1+r}$ and iid income $w_t$

$$\Rightarrow c^*(w_t, a_t) = \frac{r}{1+r}(\underbrace{a_t + w_t}_{y_t})$$

- Ex-ante identical continuum of agents with $a_0^i = 0$, $w_{it} = w_t$
  - Only heterogeneity in idiosyncratic reasoning errors
  - Ex-ante prior equals truth: $c_0(y) = c^*(y)$

# Evolution of Beliefs

- Time 1 optimal reasoning choice:

$$\hat{\sigma}_{i1}^2(y_1) = \frac{\kappa}{V_{cc}} \Rightarrow \sigma_{\eta^i,1}^2 = \bar{\sigma}_{\eta,1}^2 = \frac{\kappa\sigma_c^2}{W_{cc}\sigma_c^2 - \kappa}$$

- Heterogeneous reasoning signals due to idiosyncratic error:

$$\eta_{i,1} = \underbrace{\frac{r}{1+r}(a_0 + w_1)}_{=c^*(y_1)} + \varepsilon_{i1} \ , \ \varepsilon_{i1} \sim N(0, \sigma_{\eta^i,1}^2)$$

- Time 1 consumption:

$$\hat{c}_{i1} = \frac{r}{1+r}y_1 + \alpha_1\varepsilon_{i1}$$

# $t = 1$ Conditional Beliefs

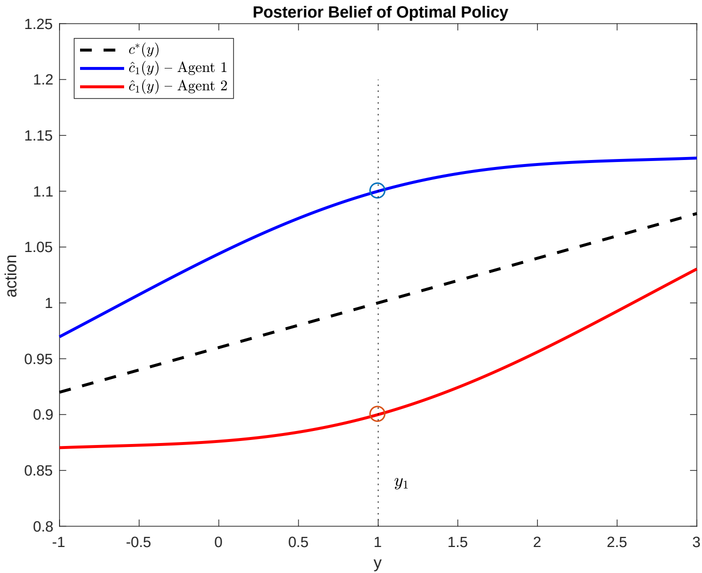

# Asset Evolution

- Since reasoning errors are iid, no aggregate effect at time 1:

$$\bar{c}_1 = \int \hat{c}_{i1}(y_1) di = c^*(y_1)$$

- But heterogeneity in time 1 action, leads to time 2 wealth dispersion:

$$y_{i2} = y_1 + w_2 - (1+r)\alpha_1 \varepsilon_{i1}$$

- $\Rightarrow$ agents face different uncertainty:

$$\hat{\sigma}^2_{i1}(y_{i2}) = \sigma^2_c(1 - \alpha_1 \exp(-2\psi(y_{i2} - y_1)^2)$$

- $\Rightarrow$ different choice of reasoning effort in time 2:

$$\sigma^2_{\eta^i,2} = \frac{\kappa \hat{\sigma}^2_{i1}(y_{i2})}{W_{cc}\hat{\sigma}^2_{i1}(y_{i2}) - \kappa} \Rightarrow \alpha_{i2} = 1 - \frac{\kappa / W_{cc}}{\hat{\sigma}^2_{i1}(y_{i2})}$$

# $t = 2$ Uncertainty

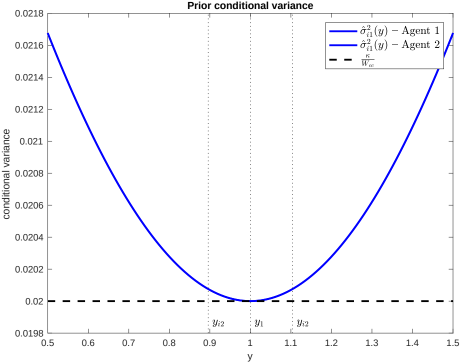


(a) Prior Conditional Variance

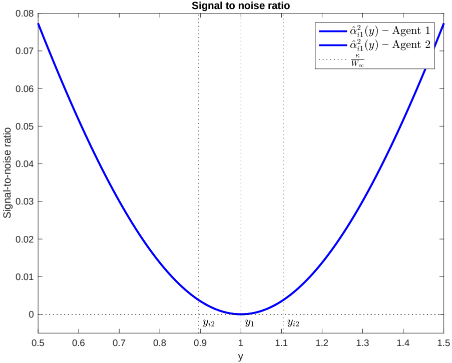


(b) Optimal Signal-to-Noise ratio

# $t = 2$ Uncertainty

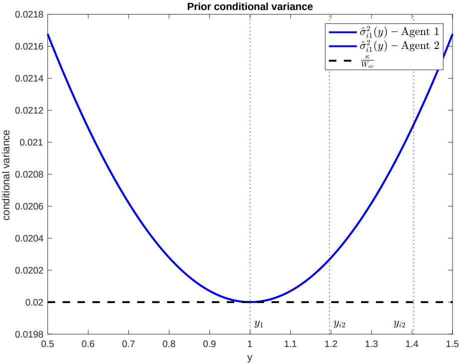


(c) Prior Conditional Variance

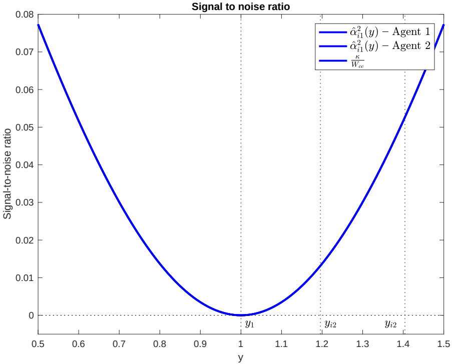


(d) Optimal Signal-to-Noise ratio

# $t = 2$ Conditional Belief

- Conditional Belief

$$\hat{c}_{i2}(y_{i2}) = (1 - \alpha_{i2})\hat{c}_{i1}(y_{i2}) + \alpha_{i2}(\frac{r}{1+r}y_{i2} + \varepsilon_{i2})$$

  - Different reasoning efforts $\Rightarrow$ different $\sigma^2_{\eta^i}$
  - But also different weights put on previous signals $\eta_{i1}$

- Aggregate effect of reasoning errors:

$$\bar{c}_2 = c^*(\int y_{i2}di) + \underbrace{\int (1 - \alpha_{i2})\alpha_1 \exp(-\psi(w_2 - (1+r)\alpha_1\varepsilon_{i1})^2)\varepsilon_{i1}di}_{>0 \iff w_2>0} + \underbrace{\int \alpha_{i2}\varepsilon_{i2}di}_{=0}$$

# $t = 2$ Conditional Beliefs

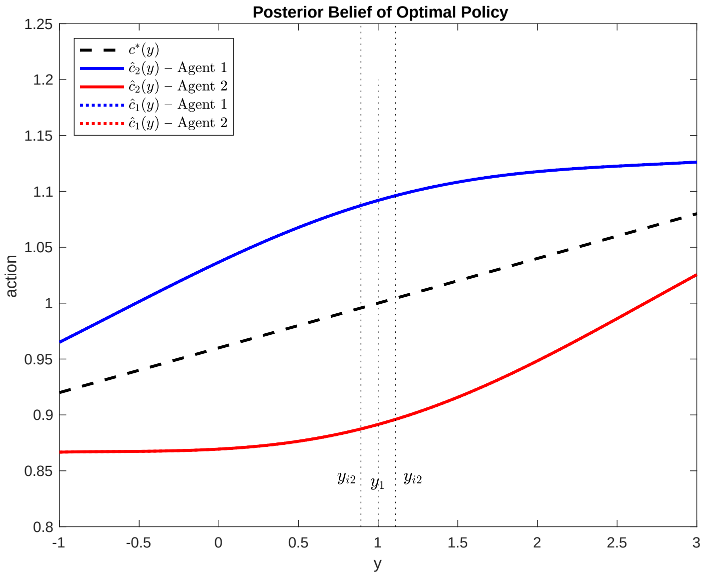

# $t = 2$ Conditional Beliefs

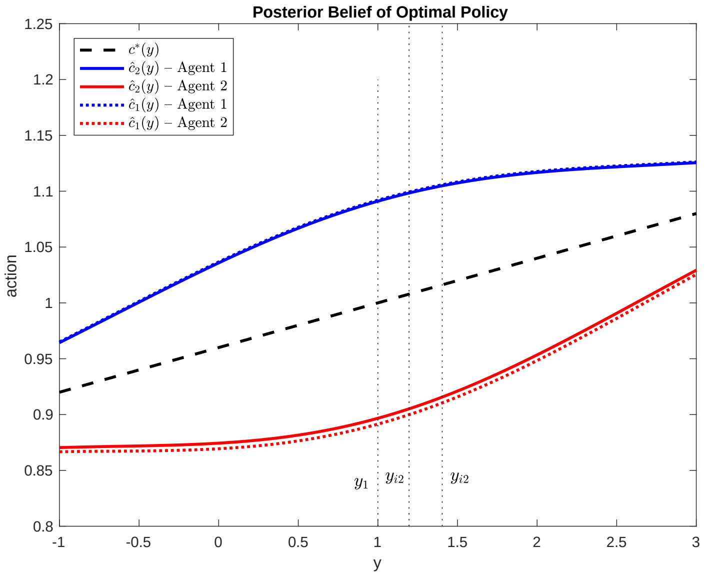

# Stochastic Choice

$$\hat{c}_{i2}(y_{i2}) = (1 - \alpha_{i2})\hat{c}_{i2}(y_{i2}) + \alpha_{i2}(\frac{r}{1+r}y_{i2} + \varepsilon_{i2})$$

- Two key sources of choice heterogeneity
  - Current period reasoning errors ⇒ stochastic choice
    - iid hence wash out in aggregate
  - History of reasoning choices and errors ⇒ systematic behavioral differences
    - experience effects, habits
    - Aggregate effects due to selection in learning/weight on previous errors
- Additionally, state variables would also differ endogenously leading to different actions even under optimality

# Asset Evolution: $t = 3$ Uncertainty

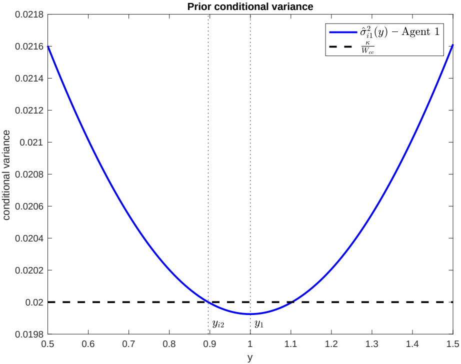


(e) Prior Conditional Variance

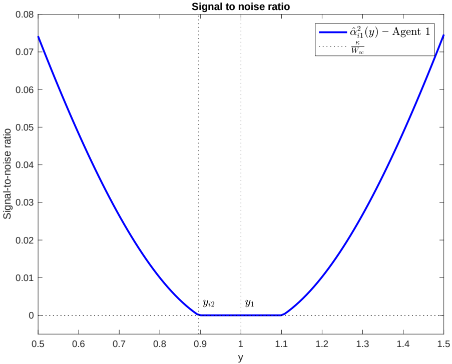


(f) Optimal Signal-to-Noise ratio

# Evolution of Assets and Beliefs

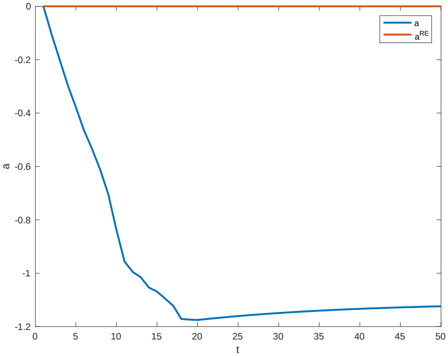


(g) Prior Conditional Variance

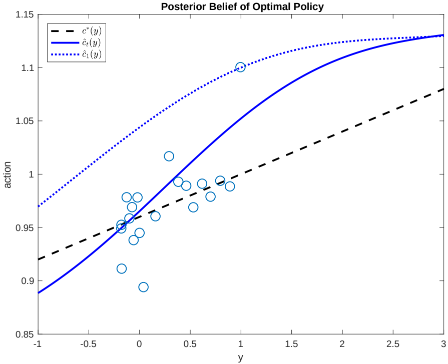


(h) Optimal Signal-to-Noise ratio

# Learning Traps

- Agents eventually settle in learning trap situations where no more reasoning occurs
- Learning stops before converging to truth
  - finite number of signals + non-random incidence
  - Initial reasoning errors do not average out
- Wealth steady state is defined by upward crossing of perceived optimal and true optimal
  - This creates a stable root in dynamics of wealth
  - Since the state (i.e. wealth) does not move around much, agent perceives no more need to keep actively learning
- Implications
  1. Persistent inequality with fat-tailed wealth distribution
  2. High local MPCs
     $\Rightarrow$ over-reaction is the norm
  3. Endogenous habit – past action dependence

# Ergodic Mean Policy Function

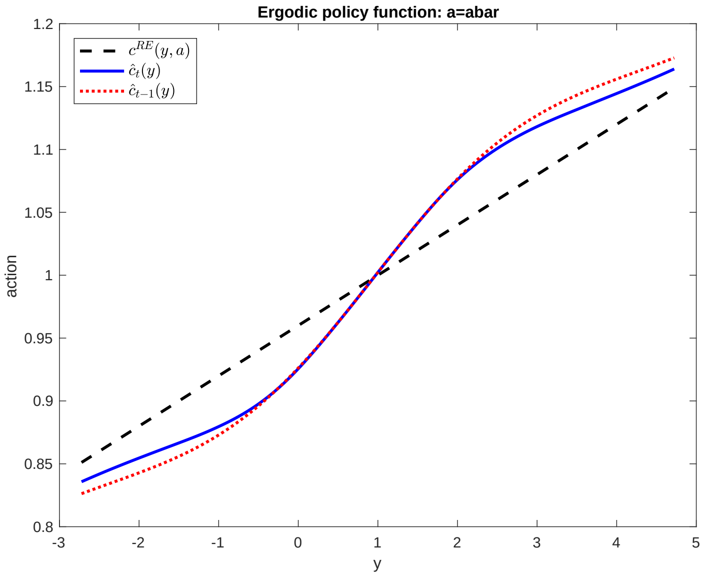

# Ergodic Moments

| Moments | Mean | Std | Kurt | $10^{th}$ | $50^{th}$ | $90^{th}$ |
|---|---|---|---|---|---|---|
| Consumption ($C_t$) | 1.01 | 0.09 | 5.10 | 0.86 | -0.02 | 1.14 |
| Wealth ($Y_t$) | 1 | 2.27 | 5.75 | -1.3 | -0.29 | 3 |
| MPC ($\frac{\partial C_t}{\partial Y_t}$) | 0.082 | 0.065 | 7.07 | 0.016 | 0.082 | 0.15 |
| Reasoning Effort ($\mathbb{1}(\alpha > 0)$) | 0.23 | 0.39 | 4.75 | 0 | 0.05 | 0.85 |

- In full information economy $MPC = 0.04$

# Impulse Response

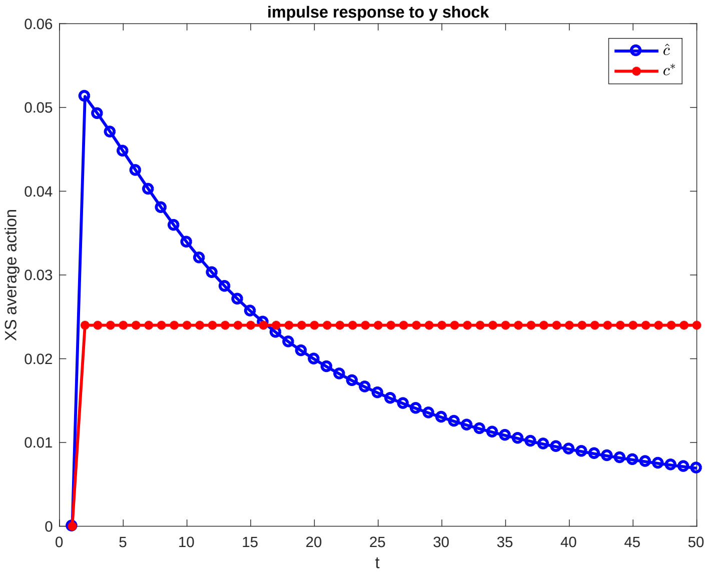

# Conclusion

- State and history dependent reasoning choices and actions
  - individual level: inertia, over-reaction, stochastic choice, biases
  - aggregate time series: shock amplification
- Study feedback between reasoning choice and endogenous states
  - reasoning signal determines state evolution (eg. consumption-savings)
  - observed state affects reasoning choice
  - future work: general equilibrium effects
- Policy implications
  - procedural rationality: patterns of errors respond to environment
  - state-dependent policy effects
- Experimental work
  - testable implications for state and history dependent reasoning errors