The Trouble with Rational Expectations in Heterogeneous Agent Models: A Challenge for Macroeconomics

---

Benjamin Moll
London School of Economics

Economic Journal Lecture, Royal Economic Society

Heterogeneous agent models with aggregate risk

- Classic papers by Krusell-Smith and Den Haan from late 90s...
- ... huge literature since then
- My argument: what we're doing "makes no sense" and the problem is rational expectations about equilibrium prices!
- Challenge = what should replace rational expectations?
  - spell out some criteria
  - discuss some promising directions

The key problem in HA models with aggregate risk

Key problem: rational expectations + general equilibrium
$\Rightarrow$ cross-sectional distribution enters household/firm decision problem

- true even though households/firms do not really care about distribution and only care about prices
- intuition: next slide
- Mean Field Games "Master equation" a.k.a. "Monster equation"

Recent work: impressive advances solving such models (DNNs etc)
(e.g. Schaab, Bilal, Bhandari-Bourany-Evans-Golosov, Han-Yang-E, Gu-Lauriere-Merkel-Payne, Gopalakrishna-Gu-Payne, Huang, Proehl)

... but this still really holds back HA literature, e.g. ~~non-linearities, crises~~

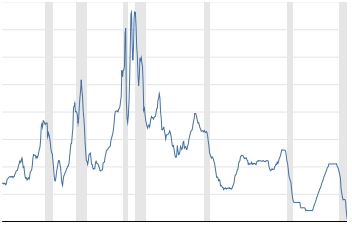
Intuition: with RE households/firms forecast prices by forecasting distributions

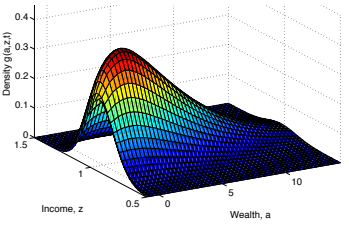
- Suppose I live in one of our models, only care about r



Intuition: with RE households/firms forecast prices by forecasting distributions

- Suppose I live in one of our models, only care about r
  - I'd realize that in equilibrium r depends on distribution G

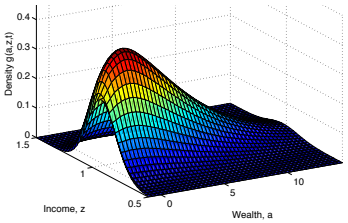




Intuition: with RE households/firms forecast prices by forecasting distributions

- Suppose I live in one of our models, only care about r
  - I'd realize that in equilibrium r depends on distribution G
  - RE $\Rightarrow$ in order to forecast r, I'd forecast entire distribution G!

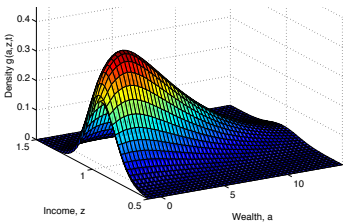




Intuition: with RE households/firms forecast prices by forecasting distributions

- Suppose I live in one of our models, only care about r
  - I'd realize that in equilibrium r depends on distribution G
  - RE $\Rightarrow$ in order to forecast r, I'd forecast entire distribution G!





- Why make our lives so hard? If best economists with best algorithms, computers cannot solve problem, stretch that real-world people do
- We're spending a lot of intellectual and computational horse power spent solving unrealistically complex problem
- $\Rightarrow$ go back to drawing board and replace RE about equilibrium prices

Plan

1. Back to the roots of RE: it was all about equilibrium prices
2. The trouble with rational expectations in heterogeneous agent models
3. What should replace RE?

Back to the roots of RE

- Back to John Muth = father of rational expectations (1961 paper)
- ... and to Lucas, Prescott, Sargent & co
- Better modeling expectations of equilibrium prices was the central goal in the development of RE

3. PRICE FLUCTUATIONS IN AN ISOLATED MARKET

We can best explain what the hypothesis is all about by starting the analysis in a rather simple setting: short-period price variations in an isolated market with a fixed production lag of a commodity which cannot be stored.[5] The market equations take the form

$$
\begin{aligned}
C_t &= -\beta p_t && \text{(Demand)}, \\
P_t &= \gamma p_t^e + u_t, && \text{(Supply)}, \\
P_t &= C_t && \text{(Market equilibrium)},
\end{aligned}
\tag{3.1}
$$

where: P_t represents the number of units produced in a period lasting as long as the production lag,

C_t is the amount consumed,

p_t is the market price in the tth period,

p_t^e is the market price expected to prevail during the tth period on the basis of information available through the $(t-1)$'st period,

u_t is an error term—representing, say, variations in yields due to weather.

All the variables used are deviations from equilibrium values.

In Bob Lucas' words

"[1960s-style macroeconometric models] implied behavior of actual equilibrium prices and incomes that bore no relation to, and were in general grossly inconsistent with, the price expectations that the theory imputed to individual agents." (Lucas 1995, Nobel Lecture)

"One needs a principle to reconcile the price distributions implied by the market equilibrium with the distributions used by agents to form their own views of the future. John Muth noted that [...] these distributions could not differ in a systematic way. His term for this latter hypothesis was rational expectations." (Lucas 1980, "Methods and Problems in Business Cycle Theory")

Lucas and Prescott (1971) "Investment under Uncertainty"

- Paper that first spells out RE the way we now understand it
- Muth: only price means consistent. Lucas-Prescott: whole distributions.

Briefly, we shall be concerned with a competitive industry in which product demand shifts randomly each period, and where factor costs remain stable. In this context, we attempt to determine the competitive equilibrium time paths of capital stock, investment rates, output, and output price for the industry as a whole and for the component firms. From the viewpoint of firms in this industry, forecasting future demand means simply forecasting future output prices. The usual way to formulate this problem is to postulate some forecasting rule for firms, which in turn generates some pattern of investment behavior, which in turn, in conjunction with industry demand, generates an actual price series.

To avoid this difficulty, we shall, in this paper, go to the opposite extreme, assuming that the actual and anticipated prices have the *same* probability distribution, or that price expectations are *rational*.[4] Thus we surrender, in advance, any

$$(9) \qquad v(k,u) = \sup_{x \geq 0} \left\{ s(k,u) - x + \beta \int v\left[kh\left(\frac{x}{k}\right), z \right] p(dz,u) \right\}.$$

[4] This term is taken from Muth [15], who applied it to the case where the expected and actual price (both random variables) have a common *mean value*. Since Muth's discussion of this concept applies equally well to our assumption of a common *distribution* for these random variables, it seems natural to adopt the term here.

Important goal of developing RE: operational macro theories

Lucas and Prescott (1971) “Investment under Uncertainty”:

- “[By imposing RE], we obtain an operational investment theory linking...”

Lucas (1980) “Methods and Problems in Business Cycle Theory”:

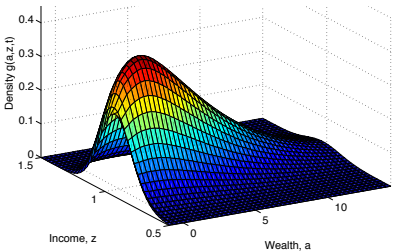
- “Our task as I see it [...] is to write a FORTRAN program that will accept specific economic policy rules as ‘input’ and will generate as ‘output’ statistics describing the operating characteristics of time series we care about, which are predicted to result from these policies.”

The trouble with rational expectations in heterogeneous agent models

Intuition: RE ⇒ forecast prices by forecasting distributions

- Suppose I live in one of our models, only care about interest rate r
  - I'd realize that in equilibrium r depends on distribution G
  - RE ⇒ in order to forecast r, I'd forecast entire distribution G!





- Why make our lives so hard?
- Real-world people do not forecast prices by forecasting distributions
- **Next: explain this in a bit more detail using specific example**

Example: forecasting equilibrium w and r

- Start with rep agent economy (RBC model) then add heterogeneity
- Preferences:

$$\mathbb{E}_0 \sum_{t=0}^{T} \beta^t U(c_t, n_t)$$

- Technology:

$$y_t = z_t F(k_t, \ell_t), \qquad k_{t+1} = i_t + (1-\delta)k_t$$

- Resource constraints:

$$c_t + i_t = y_t, \qquad \ell_t = n_t, \quad \text{all } t$$

- Notes:
  - time horizon T can be finite or ∞. Useful case: two periods $t = 0, 1$
  - aggregate productivity z_t is stochastic (Markov process)

Representative agent case: competitive equilibrium

Quantities and prices $\{w_t, r_t\}$ such that

1. Households maximize

$$\max_{\{c_t, n_t, a_{t+1}\}} \mathbb{E}_0 \sum_{t=0}^{T} \beta^t U(c_t, n_t) \quad \text{s.t.}$$

$$c_t + a_{t+1} = w_t n_t + (1 + r_t) a_t$$

2. Firms maximize

$$\max_{\{i_t, \ell_t, k_{t+1}\}} \mathbb{E}_0 \sum_{t=0}^{T} R_{0\to t}^{-1} \left(z_t F(k_t, \ell_t) - w_t \ell_t - i_t\right) \quad \text{s.t.}$$

$$k_{t+1} = i_t + (1 - \delta) k_t \quad \text{with } R_{0\to t} = \prod_{s=1}^{t} (1 + r_s)$$

3. Markets clear

$$k_t = a_t, \qquad \ell_t = n_t, \quad \text{all } t$$

Key difficulty: households, firms need to forecast w and r

Quantities and prices $\{w_t, r_t\}$ such that

1. Households maximize

$$\max_{\{c_t, n_t, a_{t+1}\}} \mathbb{E}_0 \sum_{t=0}^{T} \beta^t U(c_t, n_t) \quad \text{s.t.}$$

$$c_t + a_{t+1} = w_t n_t + (1 + r_t) a_t$$

2. Firms maximize

$$\max_{\{i_t, \ell_t, k_{t+1}\}} \mathbb{E}_0 \sum_{t=0}^{T} R_{0\to t}^{-1} \left(z_t F(k_t, \ell_t) - w_t \ell_t - i_t\right) \quad \text{s.t.}$$

$$k_{t+1} = i_t + (1 - \delta) k_t \quad \text{with } R_{0\to t} = \prod_{s=1}^{t} (1 + r_s)$$

3. Markets clear

$$k_t = a_t, \qquad \ell_t = n_t, \quad \text{all } t$$

Key difficulty: households, firms need to forecast w and r

Focus on wages $\{w_t\}$ for now

1. Households maximize

$$\max_{\{c_t, n_t, a_{t+1}\}} \mathbb{E}_0 \sum_{t=0}^{T} \beta^t U(c_t, n_t) \quad \text{s.t.}$$

$$c_t + a_{t+1} = w_t n_t + (1 + r_t) a_t$$

2. Firms maximize

$$\max_{\{i_t, \ell_t, k_{t+1}\}} \mathbb{E}_0 \sum_{t=0}^{T} R_{0\to t}^{-1} \left(z_t F(k_t, \ell_t) - w_t \ell_t - i_t\right) \quad \text{s.t.}$$

$$k_{t+1} = i_t + (1 - \delta) k_t \quad \text{with } R_{0\to t} = \prod_{s=1}^{t} (1 + r_s)$$

3. Markets clear

$$k_t = a_t, \qquad \ell_t = n_t, \quad \text{all } t$$

Key difficulty: households, firms need to forecast w and r

Focus on wages $\{w_t\}$ for now

1. Households maximize

$$\Rightarrow \quad \text{Labor supply} = n(w_t, a_t)$$

2. Firms maximize

$$\max_{\{i_t, \ell_t, k_{t+1}\}} \mathbb{E}_0 \sum_{t=0}^{T} R_{0\to t}^{-1} \left(z_t F(k_t, \ell_t) - w_t \ell_t - i_t\right) \quad \text{s.t.}$$

$$k_{t+1} = i_t + (1-\delta)k_t \quad \text{with } R_{0\to t} = \prod_{s=1}^{t}(1+r_s)$$

3. Markets clear

$$k_t = a_t, \qquad \ell_t = n_t, \quad \text{all } t$$

Key difficulty: households, firms need to forecast w and r

Focus on wages $\{w_t\}$ for now

1. Households maximize

$$\Rightarrow \quad \text{Labor supply} = n(w_t, a_t)$$

2. Firms maximize

$$\Rightarrow \quad \text{Labor demand} = \ell(w_t, k_t, z_t)$$

3. Markets clear

$$k_t = a_t, \qquad \ell_t = n_t, \quad \text{all } t$$

Key difficulty: households, firms need to forecast w and r

Focus on wages $\{w_t\}$ for now

1. Households maximize

$$\Rightarrow \quad \text{Labor supply} = n(w_t, a_t)$$

2. Firms maximize

$$\Rightarrow \quad \text{Labor demand} = \ell(w_t, k_t, z_t)$$

3. Markets clear

$$k_t = a_t, \qquad \ell_t = n_t, \quad \text{all } t$$

$$\Rightarrow \quad \text{Equilibrium wage} = w^*(k_t, z_t)$$

Note: RE equilibrium is pretty complicated even in this rep agent economy

Solution methods for representative agent case

Two (global) solution methods:

1. Tackle competitive equilibrium directly
    - actually pretty hard even in rep agent case, e.g. (k, K) trick
2. Solve via planning problem
    - no prices so completely sidesteps key difficulty
    - frequent approach in literature (e.g. RBC model)

Heterogeneous agents case: competitive equilibrium

Quantities and prices $\{w_t, r_t\}$ such that

1. Households: heterogeneous in (a_{it}, y_{it}), y_{it} = id. risk, distribution $G_t(a, y)$

$$\max_{\{c_{it}, n_{it}, a_{it+1}\}} \mathbb{E}_0 \sum_{t=0}^{T} \beta^t U(c_{it}, n_{it}) \quad \text{s.t.}$$

$$c_{it} + a_{it+1} = w_t y_{it} n_{it} + (1 + r_t) a_{it}$$

2. Firms (as before): rep firm optimally chooses $\{\ell_t, k_t\}$ given $\{w_t, r_t\}$

3. Markets clear

$$k_t = \int a dG_t(a, y), \qquad \ell_t = \int n_t(a, y) dG_t(a, y), \quad \text{all } t$$

Note: households/firms do not care about dist'n G_t, only care about prices

Key difficulty: households, firms need to forecast w and r

Focus on wages $\{w_t\}$ for now

1. Households: heterogeneous in (a_{it}, y_{it}), y_{it} = id. risk, distribution $G_t(a, y)$

$$\Rightarrow \quad \text{Household } i\text{'s labor supply} = n(w_t, a_{it}, y_{it})$$

2. Firms (as before)

$$\Rightarrow \quad \text{Labor demand} = \ell(w_t, k_t, z_t)$$

3. Markets clear

$$k_t = \int a dG_t(a, y), \qquad \ell_t = \int n_t(a, y) dG_t(a, y), \quad \text{all } t$$

Key difficulty: households, firms need to forecast w and r

Focus on wages $\{w_t\}$ for now

1. Households: heterogeneous in (a_{it}, y_{it}), y_{it} = id. risk, distribution $G_t(a, y)$

   $$\Rightarrow \quad \text{Household } i\text{'s labor supply} = n(w_t, a_{it}, y_{it})$$

2. Firms (as before)

   $$\Rightarrow \quad \text{Labor demand} = \ell(w_t, k_t, z_t)$$

3. Markets clear

   $$k_t = \int a dG_t(a, y), \qquad \ell_t = \int n_t(a, y) dG_t(a, y), \quad \text{all } t$$

   $$\Rightarrow \quad \text{Equilibrium wage} = w^*(G_t(a, y), z_t)$$

Note: equilibrium prices depend on entire cross-sectional distribution G_t!

Generic feature of heterogeneous agent models: $p_t = \mathcal{P}^*(G_t(x), z)$

Rational expectations: forecast prices by forecasting distributions

See this clearly in special case with two time periods $t = 0, 1$

1. Households solve

$$V_0(a, y, G, z) = \max_{c,n,a'} U(c, n) + \beta \mathbb{E}[V_1(a', y', G', z')|y, G, z] \quad \text{s.t.}$$

$$c + a' = w_0^*(G, z)yn + (1 + r_0^*(G, z))a$$

$$V_1(a', y', G', z') = \max_{c',n'} U(c', n') \quad \text{s.t.} \quad c' = w_1^*(G', z')y'n' + (1 + r_1^*(G', z'))a'$$

   where G' = cross-sectional distribution at $t = 1$, satisfying $G' = \mathcal{T}_{s_0} G$

2. Firm investment decision: similar problem featuring
   - prices $w_1^*(G', z')$ and $r_1^*(G', z')$
   - value function $J_1(k, G', z')$

MFG "Monster equation", makes solution extremely hard

Why make our lives so hard? Clearly people do not do this...

Solution methods for heterogeneous agent case

1. Linearization or MIT shocks: typical approach in particular in HANK literature
   - certainty (equivalence) for prices so sidesteps key difficulty
   - but not suitable for inflation debate, financial crises, asset pricing, ...
2. Krusell-Smith/DenHaan
   - forecast prices by forecasting moments of distributions, e.g. mean:

$$\bar{a}_t = \int a dG_t(a, y) \quad \text{instead of} \quad G_t(a, y)$$

   - bounded rationality interpretation
   - but do we think people do **that**? I personally also don't
   - exception: moment = price, more momentarily
     (Gomes-Michaelides, Favilukis-Ludvigson-VanNieuwerburgh, Kaplan-Mitman-Violante, Lee-Wolpin, Llull, Storesletten-Telmer-Yaron ...)
3. Tackling full RE equilibrium: impressive advances in recent literature

(e.g. Schaab, Bilal, Bhandari-Bourany-Evans-Golosov, Han-Yang-E, Gu-Lauriere-Merkel-Payne, Gopalakrishna-Gu-Payne, Huang, Proehl)

   - unrealistically complex: too much intellectual/computational horse power

Taking stock and what next?

Goal of Muth, Lucas & co when developing RE: operational macro theories

RE achieves exactly this goal in representative agent models

But **RE ⇒ het. agent models with aggregate risk "not operational"**

- attributes to people extreme ability to think through equilibrium
- means that people forecast prices by forecasting distributions
- thereby making solution extremely hard

We should go back to drawing board:

- replace RE about equilibrium prices in HA models
- existing attempts (e.g. KS 98) but we need to be more systematic
- Payoff: kill two birds with one stone
  1. make models operational (solution feasible)
  2. ... and more empirically realistic / more interesting

What should replace RE?

What should replace RE?

- I only know the problem, not the solution!
- But spell out some criteria that I find reasonable
- Common element: form expectations about prices directly
  - natural solution
  - different from RE
  - but how discipline prob. distributions to compute price expectations?
- Note: keep RE about non-equilibrium variables, e.g. idiosyncratic y_{it}

Natural solution: form expectations about prices directly

In the 2-period example

$$V_0(a, y, G, z) = \max_{c,n,a'} U(c, n) + \beta \mathbb{E}[V_1(a', y', G', z')|y, G, z] \quad \text{s.t.}$$
$$c + a' = w_0^*(G, z)yn + (1 + r_0^*(G, z))a$$
$$V_1(a', y', G', z') = \max_{c',n'} U(c', n') \quad \text{s.t.} \quad c' = w_1^*(G', z')y'n' + (1 + r_1^*(G', z'))a'$$

where G' = cross-sectional distribution at $t = 1$

Natural solution: form expectations about prices directly

In the 2-period example

$$V_0(a, y, w, r) = \max_{c,n,a'} U(c, n) + \beta \widetilde{\mathbb{E}}[V_1(a', y', w', r')|\cdot] \quad \text{s.t.}$$

$$c + a' = wyn + (1 + r)a$$

$$V_1(a', y', w', r') = \max_{c',n'} U(c', n') \quad \text{s.t.} \quad c' = w'y'n' + (1 + r')a'$$

where subjective expectation $\widetilde{\mathbb{E}}$ computed using probability distribution

$$\mathbb{P}(w', r'| \cdot)$$

Natural solution: form expectations about prices directly

In the 2-period example

$$V_0(a, y, w, r) = \max_{c,n,a'} \; U(c, n) + \beta \widetilde{\mathbb{E}}[V_1(a', y', w', r')|\cdot] \quad \text{s.t.}$$

$$c + a' = wyn + (1 + r)a$$

$$V_1(a', y', w', r') = \max_{c',n'} \; U(c', n') \quad \text{s.t.} \quad c' = w'y'n' + (1 + r')a'$$

where subjective expectation $\widetilde{\mathbb{E}}$ computed using probability distribution

$$\mathbb{P}(w', r'|\cdot)$$

Note: different from Krusell-Smith (forecast prices using moments)

- exception: moment = price

  (Gomes-Michaelides, Favilukis-Ludvigson-VanNieuwerburgh, Kaplan-Mitman-Violante, Lee-Wolpin, Llull, Storesletten-Telmer-Yaron ...)

Challenge: discipline price expectations $\mathbb{P}$

Price expectations $\widetilde{\mathbb{E}}[V(x', p')| \cdot]$ computed using probability distribution

$$\mathbb{P}(p'| \cdot)$$

Challenge: navigating the "wilderness of non-rational expectations"

Sargent (2008) AEA Presidential Address:

- "There is such a bewildering variety of ways to imagine discrepancies between objective and subjective distributions"
- "There is an infinite number of ways to be wrong, but only one way to be correct"
- "Desire to retain discipline of RE" $\Rightarrow$ "cautious modifications of RE"

Three criteria for price expectations $\mathbb{P}$

Price expectations $\widetilde{\mathbb{E}}[V(x', p')|\cdot]$ computed using probability dist'n $\mathbb{P}(p'|\cdot)$

Three criteria for $\mathbb{P}$:

1. Simplify solution of het. agent models (make them operational)
   - eliminates models that nest RE: $\mathbb{P}^{\theta}$ with $\mathbb{P}^{\theta=0} = \mathbb{P}^{RE}$ (e.g. diagnostic)
2. Consistency with empirical evidence
   - large literature, e.g. survey expectations
     (e.g. Manski, Armantier-et-al, Weber-DAcunto-Gorodnichenko-Coibion, DAcunto-Weber, Handbook of Economic Expectations)
   - large heterogeneity (disagreement) $\neq$ RE "communism" $\Rightarrow \mathbb{P}_i(p')$
3. (Some) consistency between beliefs and model reality
   - $\mathbb{P}$ "not too far" from objective price dist'n $||\mathbb{P}(p') - \mathbb{P}^{obj}(p')|| < \varepsilon$
   - fixed point problem, expectations respond to policy (Lucas critique)
   - perhaps don't need consistency for entire $\mathbb{P}$, e.g. only $\mathbb{E}[p']$?

Another requirement: compatibility with non-linear models and recursive methods

- Existing non-RE literature: mostly linear or linearized models
- But goal here: develop HA models with aggregate non-linearities
- $\Rightarrow$ need compatibility with non-linear models
- Similarly, need compatibility with recursive methods (Bellman equations)
- Lie at heart of HA solution methods for two reasons:
  1. micro non-linearities like borrowing constraints
  2. state space $\Rightarrow$ solve one optimization problem rather than millions

Some promising directions

- Temporary equilibrium and internal rationality (but only intermediate step)
  Hicks, Grandmont, Woodford, Piazzesi-Schneider, Adam-Marcet
- Survey expectations and hypothetical vignettes
  Manski, Malmendier-Nagel, Coibion-Gorodnichenko, Haaland-Roth-Wohlfart, ...
- Least-squares learning and restricted perceptions equilibrium
  Bray, Marcet-Sargent, Woodford, Evans-Honkapohja,...
- Reinforcement learning ($\neq$ deep neural networks)
  "optimal control of incompletely-known Markov decision processes" (Sutton-Barto)
- Big world hypothesis
  "agent magnitudes smaller than environment, cannot perceive state of world and action values" (Javed-Sutton)
- Heuristics and simple models
  Tversky-Kahnemann, Molavi,...
- ...

All of these: interesting in RA models but potentially larger payoff in HA models

3. PRICE FLUCTUATIONS IN AN ISOLATED MARKET

We can best explain what the hypothesis is all about by starting the analysis in a rather simple setting: short-period price variations in an isolated market with a fixed production lag of a commodity which cannot be stored.[5] The market equations take the form

$$
\begin{aligned}
C_t &= -\beta p_t && \text{(Demand)}, \\
P_t &= \gamma p_t^e + u_t, && \text{(Supply)}, \\
P_t &= C_t && \text{(Market equilibrium)},
\end{aligned}
\tag{3.1}
$$

where: P_t represents the number of units produced in a period lasting as long as the production lag,

C_t is the amount consumed,

p_t is the market price in the tth period,

p_t^e is the market price expected to prevail during the tth period on the basis of information available through the $(t-1)$'st period,

u_t is an error term—representing, say, variations in yields due to weather.

All the variables used are deviations from equilibrium values.

Least-squares learning in Muth 1961

- Producers estimate $p^e_{t+1} = \widehat{\mathbb{E}}_t[p_{t+1}]$ from past data on $p_1, \ldots, p_t$
- Simplest case: no serial correlation, no underlying state variable
- General case: serial correlation $u_{t+1} = \rho u_t + \varepsilon_t$, perceived law of motion

$$p_{t+1} = \theta_0 + \theta_1 u_t + \varepsilon_{t+1}, \quad u_t = \text{state variable}, \ \varepsilon_t = \text{i.i.d}$$

Least-squares learning in Muth 1961

- Producers estimate $p^e_{t+1} = \widehat{\mathbb{E}}_t[p_{t+1}]$ from past data on $p_1, ..., p_t$
- Simplest case: no serial correlation, no underlying state variable
- Natural approach: estimate p^e_t as mean of past observations

$$\widehat{p}^e_t = \frac{1}{t}\sum_{s=1}^{t} p_s$$

- Recursive implementation

$$\widehat{p}^e_{t+1} = \widehat{p}^e_t + \frac{1}{t}\left[p_t - \widehat{p}^e_t\right]$$

- Converges to rational expectations as $t \to \infty$ (Marcet-Sargent,...)

Least-squares learning in Muth 1961

- Producers estimate $p^e_{t+1} = \widehat{\mathbb{E}}_t[p_{t+1}]$ from past data on $p_1, \ldots, p_t$
- Simplest case: no serial correlation, no underlying state variable
- General case: serial correlation $u_{t+1} = \rho u_t + \varepsilon_t$, perceived law of motion

$$p_{t+1} = \theta_0 + \theta_1 u_t + \varepsilon_{t+1}, \quad u_t = \text{state variable}, \ \varepsilon_t = \text{i.i.d}$$

Least-squares learning in Muth 1961

- Producers estimate $p^e_{t+1} = \widehat{\mathbb{E}}_t[p_{t+1}]$ from past data on $p_1, ..., p_t$
- Simplest case: no serial correlation, no underlying state variable
- General case: serial correlation $u_{t+1} = \rho u_t + \varepsilon_t$, perceived law of motion

$$p_{t+1} = x_{t+1}^\top \theta + \varepsilon_{t+1}, \quad x_{t+1} = \begin{bmatrix} 1 \\ u_t \end{bmatrix}, \quad \theta = \begin{bmatrix} \theta_0 \\ \theta_1 \end{bmatrix}$$

Least-squares learning in Muth 1961

- Producers estimate $p^e_{t+1} = \widehat{\mathbb{E}}_t[p_{t+1}]$ from past data on $p_1, ..., p_t$
- Simplest case: no serial correlation, no underlying state variable
- General case: serial correlation $u_{t+1} = \rho u_t + \varepsilon_t$, perceived law of motion

$$p_{t+1} = x_{t+1}^\top \theta + \varepsilon_{t+1}, \quad x_{t+1} = \begin{bmatrix} 1 \\ u_t \end{bmatrix}, \quad \theta = \begin{bmatrix} \theta_0 \\ \theta_1 \end{bmatrix}$$

- Estimate $\widehat{\mathbb{E}}_t[p_{t+1}]$ by least squares (simplest case: intercept θ_0 only)

$$\widehat{\theta}_t = \left[\sum_{s=1}^{t} x_s x_s^\top \right]^{-1} \left[\sum_{s=1}^{t} x_s p_s \right]$$

- **Recursive least squares**: compute $\widehat{\theta}_{t+1}$ from $\widehat{\theta}_t$ and new data (p_t, x_t)

Least-squares learning in Muth 1961

- General case: serial correlation $u_{t+1} = \rho u_t + \varepsilon_t$, perceived law of motion

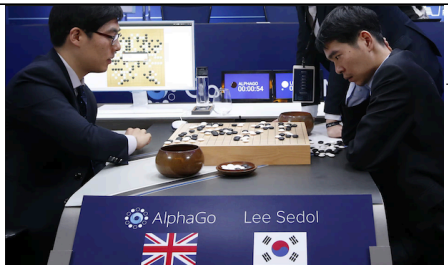
$$p_{t+1} = x_{t+1}^{\top}\theta + \varepsilon_{t+1}, \quad x_{t+1} = \begin{bmatrix} 1 \\ u_t \end{bmatrix}, \quad \theta = \begin{bmatrix} \theta_0 \\ \theta_1 \end{bmatrix}$$

- **Recursive least squares**: compute $\widehat{\theta}_{t+1}$ from $\widehat{\theta}_t$ and new data (p_t, x_t)
- Marcet-Sargent: special case of **stochastic approximation**

$$\text{NewEstimate} \leftarrow \text{OldEstimate} + \text{StepSize}\,[\text{Target} - \text{OldEstimate}]$$

- Stochastic approximation literature $\Rightarrow$ convergence results
- Connection to Krusell-Smith/DenHaan
- Jacobson (2025) "Beliefs, Aggregate Risk, and the U.S. Housing Boom"

Another stochastic approximation method: Reinforcement learning

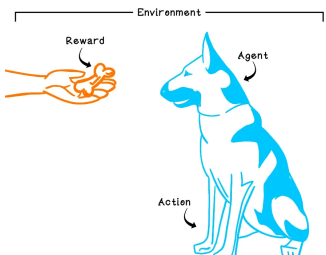




Playing Atari with Deep Reinforcement Learning

Volodymyr Mnih Koray Kavukcuoglu David Silver Alex Graves Ioannis Antonoglou
Daan Wierstra Martin Riedmiller
DeepMind Technologies
{vlad,koray,david,alex.graves,ioannis,daan,martin.riedmiller} @ deepmind.com



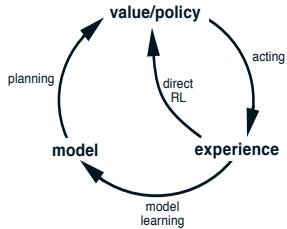


MAR 5, 2025 5:00 AM

Pioneers of Reinforcement Learning Win the Turing Award

RL: learning value functions of incompletely-known Markov decision processes

- See Sutton-Barto textbook for great introduction
- Interesting feature: learning can be completely "model-free"



- Temporal difference learning (= stochastic approximation method)

$$V(S_t) \leftarrow V(S_t) + \alpha[\underbrace{R_t + \beta V(S_{t+1})}_{\text{Target}} - V(S_t)]$$

- Temporal difference learning about equilibrium prices?
- Perhaps with behavioral twist so don't just converge to RE

Least-squares learning in HA models

Least-squares learning in heterogeneous agent models

- Spell out concrete example of approach satisfying criteria 1 and 3
- Material from “Mean Field Games without Rational Expectations” with Lenya Ryzhik
- Our version of Maggie Jacobson (2025) “Beliefs, Aggregate Risk, and the U.S. Housing Boom”

A general heterogeneous agent model with aggregate risk

- Agents i with states $X_{i,t} \in \mathbb{R}^n$, "aggregate state" $Z_t \in \mathbb{R}^k$
- State of the economy: density $g_t(x)$ and Z_t
- Each agent i chooses control $\alpha_{i,t} \in A \subset \mathbb{R}^n$ to maximize

$$V_{i,0} = \max_{\alpha_i \in A} \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t R(X_{i,t}, Z_t, \alpha_{i,t}, p_t)\right] \quad \text{subject to}$$

$$X_{i,t+1} \sim \mathcal{T}_x(\cdot|X_{i,t}, Z_t, \alpha_{i,t}, p_t), \quad Z_{t+1} \sim \mathcal{T}_z(\cdot|Z_t), \quad p_t = P^*(g_t, Z_t)$$

  - V: value function, R: reward function, $0 < \beta \leq 1$: discount factor
  - $\mathcal{T}_x$: transition prob. of $X_{i,t}$, independent across i
  - $\mathcal{T}_z$: transition prob. of Z_t
- Key: reward R depends on density g_t through low-dim. price functional

$$p_t = P^*(g_t, Z_t), \qquad P^* : \mathcal{P}(\mathbb{R}^n) \times \mathbb{R}^k \to \mathbb{R}^\ell$$

- Given stochastic policy $\alpha_{i,t} \sim \pi_t(\cdot|X_{i,t})$

$$g_{t+1}(x) = \sum_{\tilde{x},\tilde{\alpha}} g_t(\tilde{x})\pi_t(\tilde{\alpha}|\tilde{x})\mathcal{T}_x(x|\tilde{x}, Z_t, \tilde{\alpha}, p_t) \quad \text{or} \quad g_{t+1} = \mathbf{A}^\top_{\pi_t, Z_t} g_t$$

Least-squares learning about prices

- Perceived law of motion (PLM) = Markov process

$$\widehat{p}_{s+1} \sim \widehat{T}_p(\cdot|\widehat{p}_s, Z_s, \theta), \quad s \geq t, \quad \widehat{p}_t = p_t,$$

  where $\theta \in \mathbb{R}^d$ = parameter vector

- Example: VAR for prices (and possibly GDP,...)
- Learn θ over time from past observations of p_t
- Agents form estimate $\widehat{\theta}_t$ of θ, update it using learning rule

$$\widehat{\theta}_{t+1} = L(p_t, \widehat{\theta}_t)$$

- Example: recursive least squares estimator

$$\widehat{\theta}_t = \left[\sum_{s=1}^{t} x_s x_s^\top\right]^{-1} \left[\sum_{s=1}^{t} x_s p_s\right]$$

Least-squares learning in heterogeneous agent models

- Given current estimate $\widehat{\theta}_t = \theta$: solve finite-dimensional Bellman equation

$$\widehat{V}(x, z, p; \theta) = \max_{\alpha} \; R(x, z, \alpha, p) + \beta \mathbb{E}_{x', z', p'}[\widehat{V}(x', z', p'; \theta) | x, z, p] \quad \text{s.t.}$$

$$x' \sim \mathcal{T}_x(\cdot | x, z, \alpha, p), \quad z' \sim \mathcal{T}_z(\cdot | z), \quad p' \sim \widehat{\mathcal{T}}_p(\cdot | p, z, \theta).$$

- $\Rightarrow$ perceived policy $\widehat{\pi}_t(x, z, p, \theta)$. Actual policy at time t:

$$\pi_t(x, Z_t) := \widehat{\pi}_t(x, Z_t, p_t; \widehat{\theta}_t)$$

- Then solve forward-in-time system:

$$g_{t+1} = \mathbf{A}^{\top}_{\pi_t, Z_t} g_t$$
$$Z_{t+1} \sim \mathcal{T}_z(\cdot | Z_t)$$
$$\widehat{\theta}_{t+1} = L(p_t, \widehat{\theta}_t)$$

- Note: Least-squares learning $\Rightarrow$ completely sidestep "Monster equation"
- Difference to KS with prices: "in one sweep" rather than inner+outer loop
- Satisfies criteria 1 & 3 (computation & Lucas critique) but not 2 (empirics)

Summary: the trouble with RE in het. agent models

- Too much intellectual and computational horsepower solving unrealistically complex problem $\Rightarrow$ we should drop RE about equilibrium prices
- Open question: what should replace RE?
- ... how discipline $\mathbb{P}(p' | \cdot)$ to compute price expectations $\widetilde{\mathbb{E}}[V(x', p') | \cdot]$?
- Spelled out three criteria for $\mathbb{P}$
- Discussed some promising directions

Thanks!