# Slow Learning and Rational Expectations

Lawrence Christiano

based on joint work with Martin Eichenbaum and Ben Johannsen

$54^{th}$ Annual Conference of the Money, Macro and Finance Society

September 7, 2023

# Motivation

- In the past two decades, several events with little precedence occurred:
  - US Financial Crisis, European Sovereign Debt crisis, Covid, Ukraine, Climate Crisis.
- Our standard models assume rational expectations (RE)
  - Assumes people know *a lot* about the economy:
    - what can happen, the associated probabilities, etc.
  - Maybe Looked OK during the Great Moderation.
  - Harder to justify in unprecedented situations.

# What we Do

- Consider situation in which people don't have Rational Expectations and instead learn from observations as time passes.
  - For REE to be useful for policy analysis, require fast convergence to REE.
- Ask: What features of the economy determine speed of convergence to REE?
  - Use a reduced form example which suggests a simple *learning principle*:
    - When expectations of a variable are partially self-fulfilling, then learning converges *slowly* to REE, if at all.
- Turn to a particular 'event without precedence':
  - The drop in $R$ to its zero lower bound (ZLB) in 2009-2015.
- Ask: Is convergence fast enough for REE to be a useful laboratory in the ZLB?
  - Answer: No.
  - For the classic NK model, convergence is *extremely* slow in the ZLB.
  - Relate this result to the learning principle.

# Outline

- Simple example:
  - Learning principle.
- New Keynesian analysis of shocks and policies in the ZLB using Eggertsson-Woodford (2003) model.
  - Government spending
  - Forward Guidance
  - Interpret results using learning principle.

# Simple Example: REE

- Model analyzed in Bray and Savin (ECMA1986):

$$x_t = a + b\mathbb{E}_{t-1}x_t + \varepsilon_t,\ \varepsilon_t \sim iiN\left(0, \sigma^2\right), \sigma^2 < \infty$$

  'Workhorse model' for learning (see, e.g., Evans and Honkapohja (2001)). structures

- We consider the following parameter values: $-\infty < b < 1$
  - When $b < 0$: Muth's (1961) version of Cobweb model,
  - when $b > 0$, Lucas (1973) 'aggregate supply model'
- Rational expectations equilibrium:

$$\mathbb{E}_{t-1}x_t = E_{t-1}x_t,\ x_t = \overbrace{\frac{a}{1-b}}^{\mu} + \varepsilon_t.$$

- In REE, $x_t \sim iiN\left(\mu, \sigma^2\right)$.

# Simple Example: Learning

- Bayesian Learning about $\mu$ (assume people know the form of the REE process and value of $\sigma^2$)
  - In period 0, prior on $\mu$ is $N\left(\mu_0, \frac{\sigma^2}{\lambda_0}\right)$, where $\lambda_0 \geq 0$ is a measure of precision of prior.
  - In period $t$ observe $x_1, ..., x_t$, so Bayes' rule implies posterior $N\left(\mu_t, \frac{\sigma^2}{\lambda_0+t}\right)$ and

$$\mu_t = \mu_{t-1} + \frac{1}{\lambda_0 + t}(x_t - \mu_{t-1})$$
$$x_t = a + b\mu_{t-1} + \varepsilon_t$$

  - How people learn is a fundamental part of the law of motion of the system.
- Repeated substitution:

$$\mu_t = \frac{a}{1-b} + \sum_{j=1}^{t}\left\{\frac{z_t}{z_j}\frac{\varepsilon_j}{\lambda_0 + j}\right\} + z_t\left(\mu_0 - \frac{a}{1-b}\right)$$

where

$$z_t = \prod_{j=1}^{t}(1 - b_j),\ b_j = \frac{1-b}{\lambda_0 + j}.$$

# Simple Example: Convergence Questions

- does $\mu_t \to \mu = a/(1-b)$?
  - Yes for $b < 1$.
  - This result is known at least since Bray and Savin (1986).
- how fast does convergence occur?
  - potentially, very slowly.

# A Feedback Loop and Speed of Convergence

- To understand convergence rate, recall data-generating process under learning:

$$x_t = a + b\mu_{t-1} + \varepsilon_t$$
$$\mu_t = \mu_{t-1} + \frac{1}{\lambda_0 + t}(x_t - \mu_{t-1})$$

  - There is a *feedback loop* $\mu_{t-1} \rightarrow x_t \rightarrow \mu_t \rightarrow x_{t+1} \ldots$.
  - If $1 > b > 0$ : feedback loop is positive and expectations are (partially) self-fulfilling.
    - People slow to leave their initial prior, $\mu_0$.
  - If $b < 0$ expectations self-defeating.
    - People may be quick to shift away from $\mu_0$.

- Suggests speed of convergence may be a *decreasing* function of $b$.

# Simple Example: Learning Might be Very Slow (or, Fast)

- Consider expected gap relative to REE, as fraction of initial gap:

$$z_t = \frac{E\left(\mu_t - \frac{a}{1-b}\right)}{\mu_0 - \frac{a}{1-b}} = f\left(t, \lambda_0, b\right).$$

  How long does it take to close $2/3$ of initial gap, $z_T = 1/3$?

- Answer ($\lambda_0 = 1$):

| $b$ | 0 | 0.5 | 0.75 | 0.85 | 0.95 |
|---|---|---|---|---|---|
| $T$ | 2 | 10 | 120 | 2500 | 4 billion |

- We establish asymptotic properties, as $t \to \infty$, of various features of $\mu_t$.
  - For example, $z_t \simeq \kappa t^{b-1}$, $\kappa \neq 0$ as $t \to \infty$, for $b < 1$.
- Learning principle:
  - positive feedback loop ($b > 0$): slow learning.
  - negative feedback loop ($b < 0$): relatively fast learning. ▸ constant gain

# Turning to New Keynesian Model

- Recursive Formulation of NK Model
- Results:
  - Convergence to REE under learning surprisingly (to us) slow in ZLB.
- Key findings:
  - When the ZLB model is binding, NK model corresponds to a *high-b economy*,
  - Absence of Taylor principle in ZLB implies a strong positive feedback loop in inflation expectations.
  - Convergence to a REE is very slow.
- When the ZLB doesn't bind, the NK model corresponds to a low *b* economy.
  - Taylor principle outside ZLB diminishes positive feedback loop in inflation expectations.
  - Convergence to REE is faster.

# NK Model with Learning

- Simple closed economy, NK model without capital, flexible wages, Rotemberg-sticky prices.
  - Up to period 0, economy is in unique steady state REE with
    - $\beta = 1/(1 + r_{ss})$, $ss$ ˜ 'steady state'
    - gross nominal interest rate, $R > 1$.

- In period 0, everyone discovers unexpectedly that $r$ drops to $r_\ell < r_{ss}$ (Eggertsson-Woodford, 2003).
  - People know the law of motion of $r$, $r \in (r_\ell, r_{ss})$, $r_{ss}$ is an absorbing state and $P[r_{t+1} = r_\ell | r_t = r_\ell] = p$.
  - When economy reverts to absorbing state, $r = r_{ss}$, everyone understands we're back to unique steady state REE with $R > 1$.

# Model

- What people in the model don't know:
  - how the economy will evolve over time during the ZLB.
  - the dynamic impact of government policies.
- People learn about these things as data come in.
  - Circular process: learning influenced by the data and data influenced by learning.
- Two ways that people learn:
  - Bayesian learning (also, least squares learning).
  - Constant gain learning.

# Households

- Beginning of Period State Variables for $h^{th}$ household, $h \in (0,1)$:
  - $b_h$ ˜ stock of bonds acquired in previous period.
  - $r$ ˜ discount rate observed at the beginning of the period.
  - $\Theta$ ˜ parameters governing beliefs about density of $x$.
    - ⋆ $x = [C, \pi]$ ˜ aggregate variables that allow people to deduce $R$ (nominal interest rate), $w$ (real wage), $T$ (profits net of lump sum taxes)
    - ⋆ Density of $x$ degenerate when $r = r_{ss}$, non-trivial with $r = r_{\ell}$.
- The $h^{th}$ household forms plans for $C_h, N_h, b_h^{'}$ contingent on the not-yet-realized current value of $x$.

# Household $x$-Contingent Plan

- For a range of values of $x = [C, \pi]$ the $h^{th}$ household chooses $C_h, N_h, b_h^{'}$ to solve:

$$\max_{C_h, N_h, b_h'} \{\log (C_h) - \frac{\chi}{2} (N_h)^2 + \frac{1}{1+r_\ell} \left[ (1-p) V_h^{ss} \left( b_h^{'} \right) + p \mathbb{E} V_h (b_h', \Theta', x') \right] \},$$

  subject to the budget constraint:

$$C_h + \frac{b_h'}{R(x)} \leq \frac{b_h}{\pi(x)} + w(x) N_h + T(x),$$

  where $V_h$ and $V_h^{ss}$ denote the value functions in case $r = r^\ell$ or $r = r^{ss}$ in the next period, respectively. ▸ EquilibriumFunction

- Here,
  - $\mathbb{E}$ denotes the expectation operator over marginal data density of $x'$, conditional on $r' = r_\ell$, $\Theta$, $x$.
  - $\Theta'$, next period's belief parameters constructed by combining $\Theta, x$.

# Beliefs, $\Theta$

- Because they see the same aggregate data, firms and households have same beliefs about the distribution of $x = [C, \pi]$.
- People think that both elements of $\log x$ are independently drawn from a different Normal distribution.
  - They are uncertain about the mean and variance of each Normal.
  - Their joint prior over the means and variances of $C$ and $\pi$ are (truncated) Normal inverse Wishart.
- The vector $\Theta$ denotes the parameters that characterize these prior distributions.

# Evolution of Beliefs over Time

- In making their $x-$contingent decisions, people internalize that $\Theta'$ is a function of $\Theta$ and the observed value of $x$:

$$\Theta' = f(\Theta, x).$$

  Here, $f$ has an analytic representation for each of the three learning rules used.

- The people in our model are 'internally rational' in the sense of Adam and Marcet 2011.
  - Actually, the slides present results for a short-cut that Cogley and Sargent call 'anticipated utility'.

- In period 0, $\Theta_0$ are free parameters.

# Household Value Function

- Value function satisfies the following fixed point property:

$$V_h\left(b_h, \Theta, x\right) = \max_{C_h, N_h, b_h'} \left\{ \log\left(C_h\right) - \frac{\chi}{2}\left(N_h\right)^2 + \frac{1}{1+r_\ell}\left[\left(1-p\right) V_h^{ss}\left(b_h^{'}\right) + p\mathbb{E} V_h\left(b_h', \Theta', x'\right)\right]\right\},$$

  subject to the budget constraint.

- That households can map from $x$ into the aggregate variables required for their budget constraints corresponds to our assumption that they are good at *static* general equilibrium reasoning.
  - However, they are not good at *dynamic* general equilibrium reasoning.
  - Their beliefs about the future are distorted.

# Production and Firms

- Dixit-Stiglitz formalization standard in NK model.
  - Final good created by aggregating intermediate goods produced by monopolists.
- Intermediate good firms have sticky prices in the sense of Rotemberg.
- Intermediate firms' problem expressed in recursive form.
- Have same beliefs as households.

# Government

- Fiscal policy:
  - Baseline: $G = G_{ss} >$ fixed for all $r$.
  - Alternative: $G = G_\ell > G_{ss}$, $r = r_\ell$, $G = G_{ss}$, $r = r_{ss}$.
  - Government uses lump sum taxes to balance budget in each period.
- Monetary policy:

$$R = \max\left\{1, \frac{1}{\beta} + \alpha\left(\pi - 1\right)\right\}, \alpha > 1$$

- We also consider perturbations on this policy, including forward guidance.

# Market Clearing in a *Period Learning Equilibrium*

- Given $r = r_\ell$ and $\Theta$,
- The vector, $x = [C, \pi]$, is adjusted to ensure goods, bonds and labor markets clear in a way that is consistent with private sector optimization and government policy.
  - The approach is inspired by Eusepi, Gibbs and Preston, 2022.
  - Concept similar to 'Period Equilibrium' in García-Schmidt and Woodford 2019.

# Learning Equilibrium

- As long as $r = r_\ell$, economy is a sequence of period learning equilibria.
- When $r = r_{ss}$ economy jumps to $R > 1$ REE steady state.

# Is Rational Expectations a Useful Guide for Policy Analysis Under Learning?

- First,
  - Does learning select one of the (multiple) REE in the ZLB?
- Second,
  - How quickly does convergence occur?
- Third,
  - are predictions of REE about macro stabilization policies robust to learning?
- Related issue: there are also multiple *steady state* REE's in the NK model (BSGU).
  - Based on our experiments and the literature, we will focus on the zero inflation steady state.

# Multiple REE in ZLB

- Scenario:
  - Economy was in zero inflation steady state up to period 0
  - Unexpectedly, discount rate shock happens and everyone correctly believes that economy goes back to zero inflation steady state with constant probability.
  - Well known: there are two stationary rational expectations ZLB equilibria.
    - ★ In our model, can characterize a ZLB equilibrium as a zero of a function of inflation alone, $f(\pi_\ell) = 0$.
    - ★ This function has an 'inverse U', Laffer curve shape.

- Parameter values

$$p = 0.80,\ r_\ell = -0.0015\,(-0.6APR),\ G_{ss} = 0.20,\ r_{ss} = 0.005\,(2.0APR),$$

$$Y_{ss} = N_{ss} = 1,\ \varepsilon = 7,\ \phi = 110,\ \chi = 1.25,\ \alpha = 1.5$$

# REE Equilibria in ZLB

- Two ZLB equilibria
  - Bad-ZLB (A) equilibrium: substantial deflation, very high real rate, very low consumption.
  - Good-ZLB (B) equilibrium: more modest deflation, reduced consumption and high in real rate.

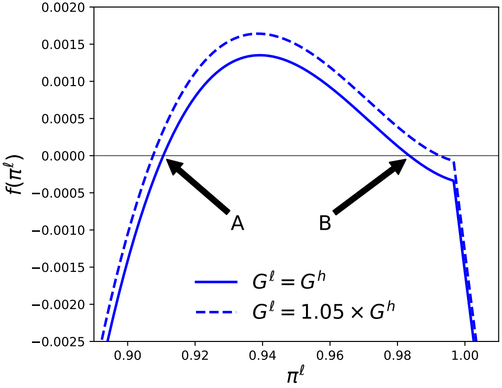

# Does Learning Select One of the Two Equilibria?

- Bad-ZLB equilibrium is locally unstable under learning.
  - When priors means are centered (priors on variance positive) on Bad-ZLB, you go to Good-ZLB.
- Good-ZLB equilibrium is 'globally' stable under learning.
  - When prior mean of $x$ is centered on steady state, on Good-ZLB or on Bad-ZLB: converge to Good-ZLB.

# Experiment #1: Slow Learning in the ZLB

- $r$ drops and $G$ remains unchanged.

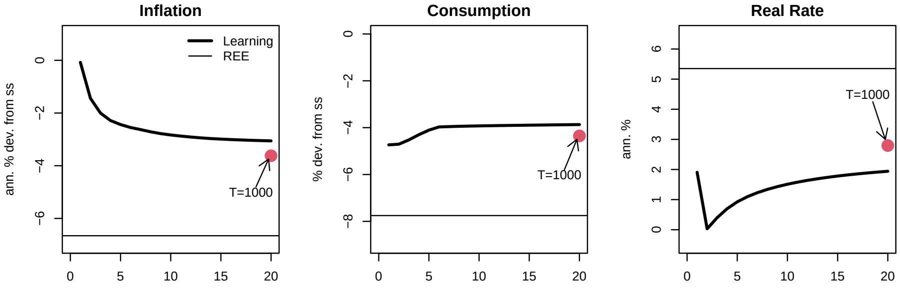


- Key results:
  - Economic impact of the shock under learning is small compared with REE.
    - ★ Learning is **extremely** slow.
  - Learning moves the model in the 'right' empirical direction:
    - ★ addresses 'missing deflation puzzle'.

# Intuition: In ZLB there is a Positive Feedback Loop Between Inflation and Inflation Expectations

- Suppose firms and households *expect lower inflation* in the future during ZLB episode.
  - Other things the same, firms want to reduce prices now.
  - Households: $R = 1$ in ZLB, so low inflation expectations $\rightarrow$ real rate high $\rightarrow$ labor supply increased $\rightarrow$ marginal cost of production down $\rightarrow$ inflation down.

- In sum: Households and firms complement each other in creating a positive feedback loop that makes the NK model behave like a 'high-b' economy.

# What if we Ignore the ZLB?

- Outside ZLB, Taylor Principle operates to prevent expectations from having a big impact on inflation.
  - Taylor Principle works to detach inflation from expectations of inflation.
  - It effectively makes $b$ small.

- Suggests that if we ignore the ZLB, so the Taylor principle remains active when $r$ falls, then convergence of the learning model to the REE should go more quickly (consistent with Ferrero 2007).

# Experiment #2: Fall in $r$ With and Without ZLB

- We do see faster convergence when don't impose the ZLB in Experiment #1, consistent with 'learning principle' intuition.

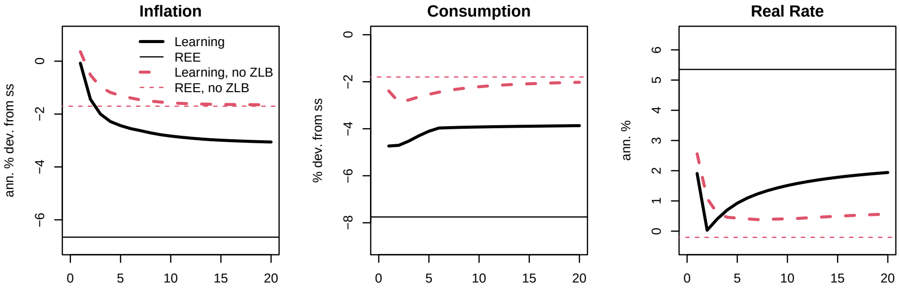

# Experiment #3: Increase in $G$ During ZLB

- Standard result in rational expectations (REE) literature:
  - multiplier on government spending can be very large in the ZLB.
    - Depends on going to the Good ZLB
    - Bad ZLB not stable under learning.
  - But, large multiplier in REE happens chiefly by raising expected inflation.
    - If learning is backward-looking, then this inflation expectation channel broken.
- Our finding:
  - We find that the multiplier under learning is very small, compared to REE.
  - Rational expectations generates *very* misleading prediction about the effects of government spending.

# Experiment #3: Impact of an Increase in Government Spending

- In the REE ZLB, government purchases have a big effect, by raising expected inflation expectations

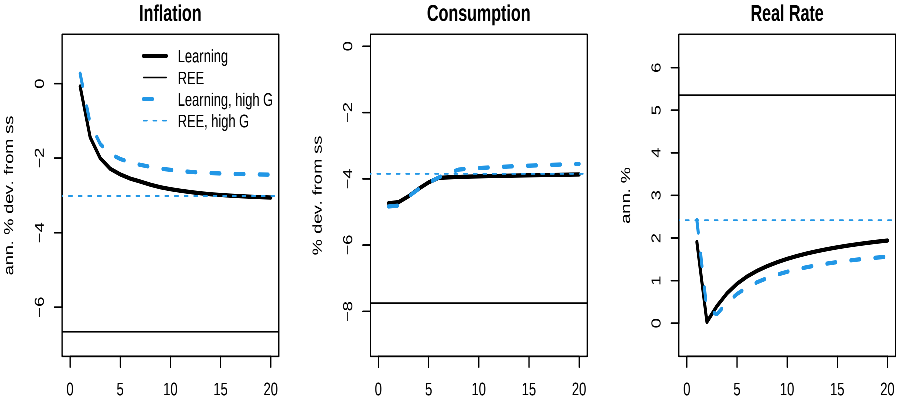


- Government purchases do very little in the learning equilibrium.

# The G Multiplier In ZLB

- Here is the multiplier, $\frac{dY}{dG}$, directly.

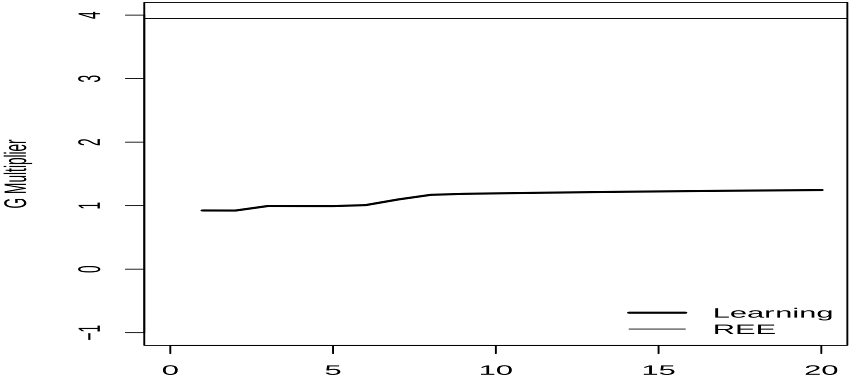


- A huge difference between REE and learning.
- Next, turn on Taylor principle by ignoring ZLB.

# The G Multiplier Outside ZLB

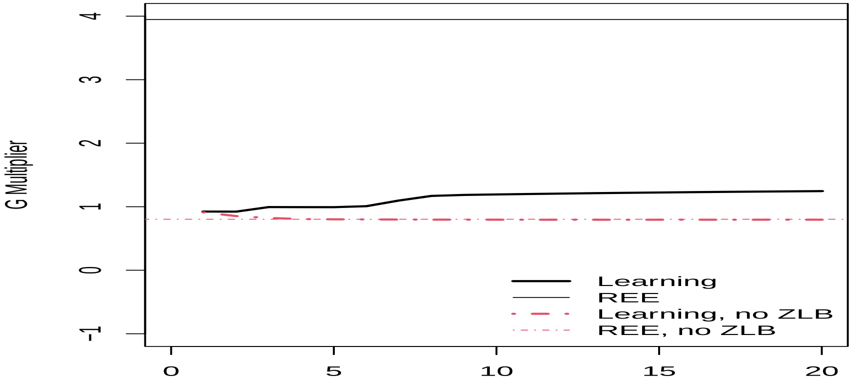

# Forward Guidance

- Simple representation of forward guidance:
  - Monetary authority announces that when $r$ jumps back up to $r_{ss}$, $R$ remains at zero for *one* period.
- REE
  - Forward guidance has huge impact on ZLB equilibrium via cascading effects on expected future inflation.
- Learning
  - The future interest rate cut does stimulate a little.
  - But, there is little amplification via expected inflation
  - No forward guidance puzzle.

# Experiment #4: Monetary Versus Fiscal Policy On the Way to the ZLB

- In our experiments with learning, when $r$ falls to $r^{\ell} < 0$, the economy takes some time (one period) to hit a binding ZLB.
- If the monetary authority reacts by pushing the economy into ZLB *immediately*, it moves allocations close to their first best levels.

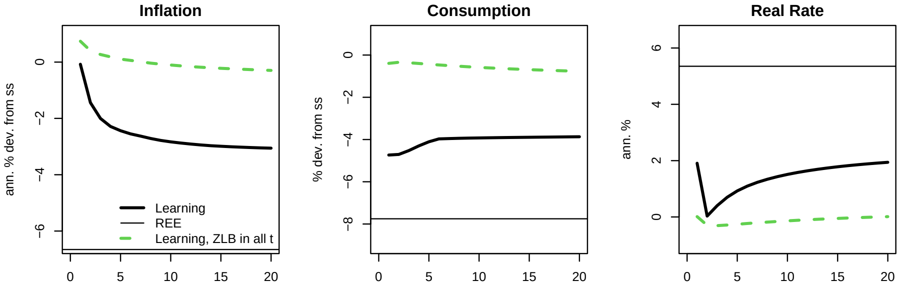


- Message: potentially, monetary policy can do even *more* than fiscal policy in dealing with a ZLB episode.
  - Requires moving monetary policy very quickly.

# Conclusion

- The preceding analysis is a cautionary tale about how REE analysis may convey misleading policy advice:
  - Could encourage fiscal authorities to rely excessively on fiscal policy.
  - Could cause monetary authority to pass up an opportunity to move quickly while $R$ remains $> 1$.
- The learning principle may suggest other circumstances in which REE delivers conclusions not robust to learning.
- Analysis confirms the wisdom of exploring the implication of replacing REE by alternative micro-founded learning mechanisms (see Gabaix, Angeletos, Fahri, Werning, Ilut, Schneider, Liu, Sastry, Shleifer, Woodford, ...).

# Appendix Materials

# Period Price and Profit Functions

- Households (and firms) observe $x = \left[\begin{array}{cc} C, & \pi \end{array}\right]$
  - from $x$ (as well as $r, G(r)$) they are able to deduce the variables needed to define their current-period budget constraint.
- GDP $(Y)$, aggregate employment $(N)$, real wage $(w)$, marginal firm cost $(s)$, profits, taxes net of profits $(T)$:

$$N = Y = (C + G(r))\left(1 + \frac{\phi}{2}(\pi - 1)^2\right)$$

$$w = \chi N C,\ s = (1 - \nu) w,\ R = \max\left\{1, 1 + r^h + \alpha(\pi - 1)\right\}.$$

We assume the government issues no debt and finances its expenditures with lump sum taxes:

$$G(r) + \nu w N,$$

where $\nu w N$ represents the subsidy paid to intermediate good firms.

# Period Price and Profit Functions, cnt'd

- Finally, profits net of taxes implied by $x$ and $r$ are:

$$T = \overbrace{(1-s)\,Y - \frac{\phi}{2}(\pi - 1)^2\,(C + G(r))}^{\text{profits for intermediate good producers}} - \overbrace{(G(r) + \nu w Y)}^{\text{lump sum taxes}}.$$

- Note: none of these mappings use bond market clearing or the household's intertemporal Euler equation. Go Back

# Cobweb Model

- Model of competitive market and a time lag in production.
  - John Muth, 'Rational Expectations and the Theory of Price Movements', ECMA, July 1961.
  - Coase and Fowler, 'Bacon Production and the Pig-Cycle in Great Britain', Economica, May, 1935.
- Demand:

$$d_t = m_I - m_p p_t + \upsilon_{1t}$$

- Supply decided in period $t$ before $\upsilon_{1t}$ is observed:

$$s_t = r_I + r_p \mathbb{E}_{t-1} p_t + \upsilon_{2t}$$

- Equilibrium, $d_t = s_t$:

$$\overbrace{p_t}^{x_t} = \overbrace{\frac{m_I - r_I}{m_p}}^{a} \overbrace{- \frac{r_p}{m_p}}^{+b} \mathbb{E}_{t-1} p_t + \overbrace{\frac{\upsilon_{1t} - \upsilon_{2t}}{m_p}}^{\varepsilon_t}$$

# Lucas Model

- Aggregate output:

$$q_t = \overline{q} + \pi \left( p_t - \mathbb{E}_{t-1} p_t \right) + \zeta_t$$

- Velocity equation:

$$m_t + \upsilon_t = p_t + q_t$$

- Monetary policy:

$$m_t = \bar{m} + u_t.$$

- Substitute second two equations into first, to obtain equilibrium condition:

$$\overbrace{p_t}^{x_t} = \overbrace{\frac{\bar{m} - \overline{q}}{1+\pi}}^{a} + \overbrace{\frac{\pi}{1+\pi}}^{b} \mathbb{E}_{t-1} p_t + \overbrace{\frac{u_t + \upsilon_t - \zeta_t}{1+\pi}}^{\varepsilon_t}$$

# Rational Expectations Equilibrium

- Reduced form model:

$$x_t = a + b\mathbb{E}_{t-1}x_t + \varepsilon_t,\ \varepsilon_t \sim E\varepsilon_t = 0, E\varepsilon_t^2, E\varepsilon_t\varepsilon_{t-j} = 0, j \neq 0.$$

- In rational expectations equilibrium, $\mathbb{E}_{t-1}x_t = E_{t-1}x_t$, so

$$x_t = \frac{a}{1-b} + \varepsilon_t$$

- To verify this, note:

$$\begin{aligned} x_t = a + bE_{t-1}x_t + \varepsilon_t &\overbrace{=}^{REE} a + b\overbrace{\frac{a}{1-b}}^{E_{t-1}x_t} + \varepsilon_t \\ &= \frac{a}{1-b} + \varepsilon_t. \end{aligned}$$

Go Back

# Constant-gain learning

- Assume people update their view of $\mu_{t-1}$by constant-gain learning:

$$\mu_t = \mu_{t-1} + \gamma \left( x_t - \mu_{t-1} \right), \tag{1}$$

for $0 < \gamma < 1$.

- Now

$$\mu_t - \frac{a}{1-b} = \sum_{j=0}^{t-1} (1-\gamma_b)^j \left( \frac{\varepsilon_{t-j}}{1-b} \right) \gamma_b + (1-\gamma_b)^t \left( \mu_0 - \frac{a}{1-b} \right),$$

where $\gamma_b = (1-b)\,\gamma$,

$$z_t = E\left( \frac{\mu_t - \frac{a}{1-b}}{\mu_0 - \frac{a}{1-b}} \right) = (1-\gamma_b)^t\,.$$

# Learning principle again

- Again calculate how long it takes to close 2/3 of the initial gap, i.e., calculate, $T$, the value of $t$ such that $z_T \simeq 1/3$.

- Suppose $\gamma = 0.5$ and $b = 0, 0.5, 0.75, 0.85, .95$.

| $b$ | 0 | 0.5 | 0.75 | 0.85 | 0.95 |
|---|---|---|---|---|---|
| $T$ | 1.6 | 3.8 | 8.23 | 14.1 | 4 billion |

- Note: speed of convergence is quicker for ‘small’ values of $b$ than under Bayesian learning.

- But again speed of convergence increases nonlinearly with $b$. Go Back