Discussion: A Monetary Policy Framework for a Volatile World, by Matteo Cacciatore and Stefano Gnocchi

Martin Eichenbaum

October 14, 2025

Introduction

- It's a pleasure to discuss this excellent and thoughtful paper by Cacciatori and Gnocchi.

- Challenge: add value to their discussion of how to rethink policy in a 'new' world where shocks have become more severe.

- Outline of my comments
  - The issue is not supply shocks per se - it's about the persistence of shocks.
  - If we're entering a 'new regime', how should we model private sector expectations?
  - What are effective forms of communication to help people learn about changes in the BoC's approach and the way the economy will evolve?

- What about exchange rates? No time today but this is an important aspect of policy not discussed in Cacciatore and Gnocchi.

Looking through supply shocks

- Premise of the Cacciatore and Gnocchi paper: we are (probably) entering a ‘new regime’.

- Supply shocks may become more frequent and severe.
  - Political Instability,
  - Unravelling of consensus on of free trade, globally interlinked economies,
  - Climate change.

Divine Coincidence: an intellectual curiosum

- Textbook NK model and many discussions of policy sharply distinguish between demand and supply shocks.
- NK model $\Rightarrow$ divine coincidence holds for 'demand shocks'.
  - Absent 'markup-shocks and an efficient steady state: stabilizing price inflation also stabilizes the output gap.
- So monetary should respond quickly and forcefully to demand shocks. Supply shocks are more complicated.
- It's often said that monetary policy should "look through" temporary supply-driven inflation as long as 'these shocks would fade without de-anchoring long-term inflation expectations'.
- Easy objections:
  - There's no reason to think that the flexible price, steady state level of output of inflation is efficient.
  - In real time, we don't know if inflation was caused by a demand or a supply shock (we're still debating how much of post-Covid inflation was driven by demand as opposed to supply shocks).

Demand and supply shocks

- Any reasonable empirical model of the business cycle *must* incorporate wage inertia.
  - True for RANK, TANK, HANK, Search labor markets in DSGE models.
  - See for example Christiano, Eichenbaum and Trabant, 2016, 2021).
- In most models, this means it's not particularly useful to distinguish between demand and supply shocks isn't so interesting.
- It also becomes clear that divine coincidence is, at best, an intellectual curiosum.

Demand and supply shocks with Calvo prices and wages

- There's two nominal distortions: price dispersion and wage dispersion.
  - Monetary policy can't simultaneously eliminate both.

- The price Phillips curve

$$\pi_t^p = \beta E_t \pi_{t+1}^p + \kappa_p x_t + \gamma(\mu_t^p - \mu_t^{p,*})$$

where x_t is the output gap and μ_t^p is ratio of price to nominal marginal cost.

- The wage Phillips curve ties the wage - markup gap to wage inflation

$$\pi_t^w = \beta E_t \pi_{t+1}^w + \kappa_w(\mu_t^w - \mu_t^{w,*}).$$

where μ_t^w is the gap between the real wage and the household's marginal rate of substitution.

- Since marginal costs depend on wages, sticky wages effectively introduce a cost-push term into the price Phillips curve.

Divine coincidence

- With sticky wages, sticky prices there's a policy trade-off between stabilizing price inflation and the output gap.
- μ_t^w moves differently than μ_t^p.
- Stabilizing prices doesn't ensure $\mu_t^w = \mu_t^{w,*}$.
- Real wage misalignment persists $\rightarrow$ output gap remains $\rightarrow$ divine coincidence fails.
- Conclude: divine coincidence is not a relevant concept for actual policy decisions. .

What causes movements in μ_t^w?

- μ_t^w reflects any shock that changes desired real wage (MRS) faster than nominal wages can adjust.
  - Productivity shocks (AI).
  - Demand (monetary, preference or exuberance) shocks.
  - Cost-push or price markup shocks (trade developments which affect competitiveness in product markets).
  - Labor supply shocks (changes in unemployment benefits).
  - Wage markup shocks (trade or legal developments that monopsony-like power of firms.

- Since many shocks affect both μ_t^p and μ_t^w, demand shocks don't have any obvious special status relative to supply shocks.

What deserves special status: the persistence of a shock

- Monetary operates with ‘long and variable logs’.
- Reacting to any transitory shock - demand or supply - will be de-stabilizing, even if long-run expectations are anchored (class Friedman critique).
- Transitory shocks are unlikely to de-anchor expectations if the central bank, has a modicum of credibility.
  - the Bank of Canada has enough credibility to look-through temporary shocks.
- A central task for policymakers is to assess whether a shock is temporary versus persistent.

On the reliability of models with ‘rational expectations’

- The Bank of Canada is correct that we’re probably entering a period of persistent changes to the economy.
- *Perhaps* rational expectations provides a reasonable approximation to market outcomes in tranquil times like the Great Moderation.
- Difficult to take REs seriously when people are confronted with novel environments.
- Central questions for policymakers:
  - What features of an economy determine which policy implications are sensitive to RE and how long does it take to converge to REs.

Three examples of the 'Learning Principle'

- Three examples
  - The Muth model
  - Prolonged period in which the ZLB is binding.
  - Predicting the effects of changes in fiscal policy.

- These examples suggest a simple *learning principle*:
  - When beliefs are partially self-fulfilling, learning equilibria converge slowly to rational expectations.

- The examples also provide a new role for policy as a communication tool.
  - e.g. different Taylor rule settings can greatly affect the rate of convergence to rational expectations.
  - CEJ (2025): the Taylor rule and the speed of convergence.

A general result for linearized solutions to the model

- A particular parameter which we denote by b, controls how beliefs about market outcomes affect actual market outcomes.
- We identify the analog of b in the linearized solution of our NK model.
  - It's the largest real part of the eigenvalues of the matrix that maps beliefs about the state of the economy into their realized values.
- This finding applies to a large class of linearized solutions of dynamic models which have a similar form.

Simple Example

- 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:

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

Simple Example: Learning

- Bayesian Learning about μ (assume people know the form of the REE process and value of σ^2)
  - In period 0, prior on μ 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 people observe $x_1, ..., x_t$and use Bayes' rule to form beliefs, $\mu \sim N\left(\mu_t, \frac{\sigma^2}{\lambda_0+t}\right)$,

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

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

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

Simple Example: Convergence Questions

- Does $\mu_t \to \mu = a/\left(1-b\right)$?
  - Yes for $b < 1$.
  - This result is known at least since Bray and Savin (1986).
- How fast does convergence occur?
  - potentially, *very* slowly.
- Rate of convergence is *non-linearly decreasing* in b.

A Feedback Loop and Speed of Convergence

- Data-generating process under learning:

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

  - There's a *feedback loop* $\mu_{t-1} \to x_t \to \mu_t \to x_{t+1}\ldots$

- If $0 < b < 1$: feedback loop is positive and expectations are (partially) self-fulfilling.
  - The higher is b, the more *self-fulfilling* is μ_{t-1}, that is, $dx_t/d\left(\mu_{t-1}\right) > 0$.
    - ★ The higher is x_t, the higher is μ_t, that is, $d\mu_t/dx_t > 0$.
    - ★ So, the strength of the positive feedback loop from μ_0 to μ_t is stronger the higher is b and the slower is the speed of convergence.

- If $b < 0$ expectations are self-defeating.
  - People are quick to shift away from μ_0 implying fast convergence.

Rate of Convergence

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

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

- $1 - z_t$: fraction of the initial gap, $\mu_0 - a/\left(1-b\right)$, closed by period t.

- 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	3	11	113	2201	5.2 billion

NK Model with Learning

- Learning principle: *When expectations of a variable are partially self-fulfilling, learning converges slowly to REE.*
- The larger is b the longer it takes to converge.

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

- For practical purposes . . . 'slow' convergence may mean no convergence.
- CEJ (2025) generalize this idea for a large class of linear systems (log linearized DSGE models.
- Analog of b in the simple NK model: the largest real part of the eigenvalues of the matrix that maps beliefs about $x = \begin{bmatrix} C & \pi \end{bmatrix}$ into the realized values of x.

Learning in the ZLB

- CEJ (2024): convergence is *very* slow.
- ZLB would almost certainly be over long before learning has come close to converging.
- *Policy analysis* - increases in government spending and forward guidance - based on REEs is *very* misleading.

New Keynesian Model: Key Results

- When the ZLB is binding, NK model corresponds to a *high-b economy*.
  - By learning principle, convergence to REE is slow.
  - Taylor rule inoperative $\Longrightarrow$ there's a strong, positive feedback loop between expected future inflation and actual future inflation.
  - Implies convergence to REE is *very* slow.

- Policy analysis based on RE is very misleading.

- When the ZLB is not binding, then our NK model looks like a *lower-b economy*.
  - Taylor principle reduces feedback between expected inflation and actual future inflation.
  - The strength of this channel depends on the parameter values of the Taylor rule

NK Model with Learning

- Simple closed economy, NK model without capital, flexible wages, Rotemberg-sticky prices.
- Household and firm problems expressed in recursive form.
- Recursive form convenient for thinking about learning.
  - people begin each period with beliefs as a state variable.
  - when they see that period's prices and income, they update their beliefs and make a decision.
  - their decision fully internalizes that new data will arrive in the future and that they will change their beliefs.

NK Model with Learning

- 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 economy jumps back to unique steady state REE with $R > 1$.

Model

- What people in the model don't know:
  - how the economy will evolve during the ZLB when $r = r_\ell$.
  - the dynamic general equilibrium effects of government policies.
- People learn about these things as data come in a Bayesian way.
- Circular process: learning influenced by the data and data influenced by learning.

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:

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

- CEK (2025) also consider perturbations on this policy, including forward guidance.

Initial beliefs

- People think both elements of $log(x) = ln[C, \pi]$ are independently drawn from Normal distributions.
  - They're uncertain about the mean and variance of each Normal.
- Their joint prior over the means and variances of C and π are (truncated) Normal inverse Wishart.
- The vector Θ contains the parameters that characterize the prior distributions.
  - location, shape, scale, and precision.Θ is a free parameter in the first period of the simulation.

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

- For today, I will set Θ so that
  - the means of C and π are centered on their REE steady state values
  - the standard deviation of C and π are relatively large: on the order of 5% and 10%, respectively (diffuse priors).

Experiment #1: Slow Learning in the ZLB

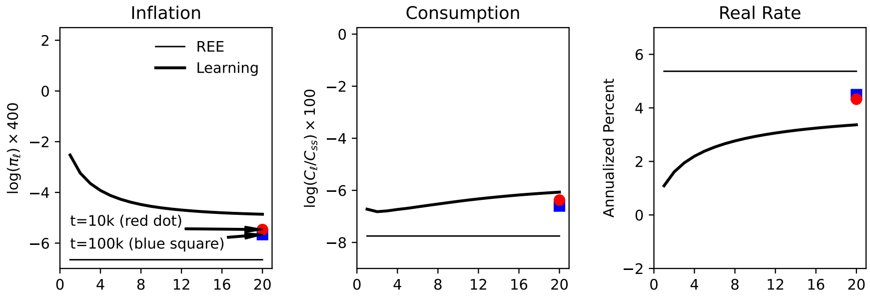

- Key results:
  - Economic impact of the shock under learning is small compared with REE.
  - Learning is **extremely** slow.
  - Learning moves the model in the 'right' empirical direction:
    - addresses 'missing deflation puzzle'. ▸ nzlb

Intuition

- Suppose that firms and households expect lower inflation in the future.
  - Because of price-setting adjustment costs, firms are incentivized to cut prices today.
  - In the ZLB, low inflation expectations mean households believe the real interest rate is high.
  - So, households reduce their demand for consumption, which leads to a fall in the marginal cost of production.
  - So, the actions of both households and firms lead to lower current inflation.
  - With learning, low π_t shifts expected inflation down in the next period, partially fulfilling their beliefs about next period's inflation rate.

- Previous mechanism repeats itself in the next period so that actual $\pi_{t+1}$is also low.

- Conclude that in the ZLB, deflation expectations are partially self-fulfilling, and the NK model behaves like a high *b* economy.

Role of Binding ZLB in Slow Convergence

- Same shock as above, but ignore lower bound on R.

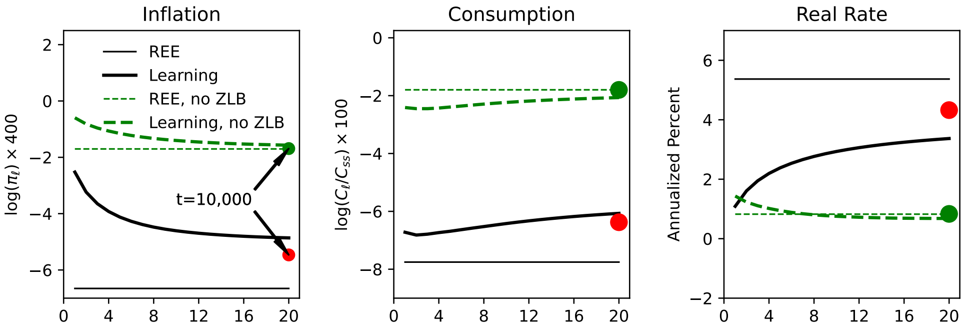

- Learning equilibrium converges relatively quickly to REE with Taylor Principle (ZLB ignored).
- Reflects learning principle: Central Bank "Does what it Takes" to keep inflation on target, independent of prior expectations of the public: 'low b economy' with Taylor Principle.

Experiment #2: Increase in G During ZLB

- Standard result in rational expectations (REE) literature:
  - multiplier on government spending can be very large in the ZLB.
- Large multiplier in REE happens chiefly by raising expected inflation.
  - If learning is backward-looking, then this inflation expectation channel broken.
- Our finding:
  - The multiplier under learning is very small, compared to REE.
  - Rational expectations generates *very* misleading prediction about the effects of government spending.
- For similar reasons: *forward guidance*: much less powerful under learning.

Experiment #2: Impact of an Increase in Government Spending

- In the REE ZLB, government purchases have a big effect, by raising expected inflation expectations.
- Multiplier, $\frac{dY}{dG}$.

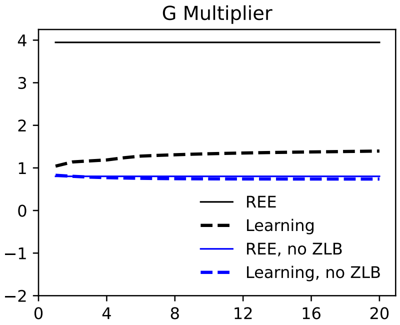

- A huge difference between REE and the learning equilibrium.
- Government purchases do very little in the learning equilibrium.

ZLB application is *not* unique

- See for example model of the Phillips curve in Beaudry et. al. 2024 in which people estimate persistent component of inflation from data on observed prices.
- Shares key feature of NK ZLB model: partially self-fulfilling beliefs.
- CEJ (2025): the Taylor rule and the speed of convergence.

The Taylor rule and communication in normal time

- Central banks are very interested in communicating their strategies.
- In practice they communicate primarily to economists, Wall Street, Bay Street and large firms.
- There are many survey-based reasons to think that households don't pay attention to central bank communication.
- 'Normal' people learn from their personal experiences.
- The 'settings' of the Taylor rule can greatly influence how quickly people and the rate of converge to REs.
- This role for the Taylor rule is very different from preventing 'sun spot' equilibria
  - See Nakamura et. al. 2025 for a summary of why we shouldn't take non-uniqueness seriously.

The Taylor rule, learning and communication

- CEJ (2025) examine speed of convergence in a full scale DGSE model (CEE).
- Again, the analog of 'b' determines asymptotic rate of convergence of the learning equilibrium.
- We already saw how the coefficient on inflation (α) in the Taylor rule affects the speed of learning (lowers'b') when we discussed the ZLB.
  - If people expect high inflation, they begin to spend more today.
  - If actual inflation is high today, the central bank raise the nominal and real interest rate today.
  - That lowers aggregate demand and actual inflation today.
  - So tomorrow people will learn that their expectations were wrong and they quickly revise their expectations.
- A high α communicates information about the structure of the data by changing the data that people see.

The Taylor rule, learning and communication

- Using similar logic, CEJ (2025) show that the more inertia there is in the Taylor rule, the slower is convergence,

- In addition, the more weight the policymaker gives to fighting the variance of the output gap, the slower is convergence to an REE.

- This logic amplifies the logic of looking through transitory shocks.

- But it highlights a potential tension between the communication channel given anchored long-term expectations and communication about the central bank's commitment to long- term inflation targets.

Conclusion

- It's likely we're entering an era in which the structure of the economy is changing.
- In deciding how to react to shocks, it is critical for the central bank to assess how persistent a shock is.
- We should look through transitory shocks - to minimize the variability of output - the long and variable argument may trump the communication argument.

Conclusion

- I hope my examples convey the importance of people's expectations when confronted with persistent 'regime changes'.
- Data about those expectations are critical to good policy design even when long-term expectations are anchored.
- That data will also provide critical information about how well anchored expectations are if we let transitory shocks pass-through.
- We must devote resources to measuring firm and household expectations in real time and guide policy with that information.