

The dangers of policy experiments: initial beliefs under adaptive learning^{*}

Patrick A. Pintus[†] Jacek Suda[‡] M. Burak Turgut[§]

Abstract

The paper studies the implications of initial beliefs and associated confidence on the system's dynamics under adaptive learning in the context of policy considerations. We examine how the effects of discretionary experimentation with new macroeconomic policies depend on the expectations that agents have in relation to these policies. We show that a newly introduced macroprudential policy aimed at making leverage countercyclical can lead to a substantial increase in fluctuations under learning, when the economy is hit by financial shocks, if beliefs reflect imperfect information about the policy experiment. This is in stark contrast to the effects of such policy under rational expectations.

Keywords: Expectations, Learning, Prior beliefs, Collateral Constraint, Macroprudential Policy

JEL classification: E32, E44, D83, E61, G18

1 Introduction

Designing and implementing effective policy is a complex challenge that involves considering incentives and encouraging specific actions. Policymakers must navigate questions about how to respond to extreme events and address slow-moving trends, as well as determine the appropriate timing for deploying new policies. These questions arise both during crises—such as health, economic, financial, sovereign debt, and pension crises—and in more stable periods.

Researchers in political science, sociology, social policy, and economics have long sought answers to these questions. The rational expectations revolution that started over 50 years ago fundamentally reshaped macroeconomic theory and policy by emphasizing that individuals and firms form expectations about the future in a forward-looking, systematic manner. Central to this approach is the idea that economic agents do not passively react to policy changes based on historical patterns or naive forecasting methods. Instead, they incorporate all available information, including their understanding of economic policies and their likely

^{*}The authors gratefully thank Kirill Mansurov for the discussion and his comments as well as participants at FMND 2023, AEA 2022 conference, CEF 2021 conference, CFE 2020 meeting, SNDE 2018 conference and GRAPE seminar for comments and suggestions. We thank Artur Rutkowski for research assistance. The financial support of National Science Centre (grant UMO-2016/21/B/HS4/03017) is gratefully acknowledged.

[†]Institut de Neurosciences de la Timone and Aix-Marseille School of Economics, Aix-Marseille University. *Email:* papintus@gmail.com.

[‡]Corresponding author. SGH Warsaw School of Economics. *Email:* jsuda@sgh.waw.pl.

[§]University of Warsaw. *Email:* mturgut@wne.uw.edu.pl.

effects, into their decision-making. This approach recognizes that (i) expectations about the future influence today’s decisions, and (ii) individuals react to changes in their economic environment, such as tax changes. This insight challenged the efficacy of traditional stabilization policies, particularly those reliant on systematic monetary or fiscal interventions, by asserting that such policies would be anticipated and neutralized in their intended effects through adjustments in behavior.

However, the rational expectations hypothesis rests on a strong and, in some respects, controversial assumption: that agents possess a deep, almost encyclopedic knowledge of the economy and its underlying structure. They are assumed to understand the true model of the economy, including the complex relationships between policy instruments and macroeconomic outcomes, and to accurately and instantaneously incorporate this knowledge into their expectations. This presumes not only a high level of information but also a capacity for sophisticated economic reasoning that may be unrealistic in practice. This informational assumption becomes particularly pronounced in the context of policy changes. Households and firms are assumed not only to fully understand a new policy but also to comprehend how it will affect their current and future constraints. Moreover, they are expected to act immediately on this knowledge, adjusting their behavior and choices accordingly. Real-world agents face cognitive and informational constraints, uncertainty about the true workings of the economy, and the possibility of structural changes that render past relationships unreliable. While rational expectations provide an elegant benchmark for understanding how expectations shape economic outcomes, their strong informational assumptions raise questions about their applicability to real-world settings, where bounded rationality, adaptive learning, and institutional factors often play significant roles in shaping behavior.

The adaptive learning approach supports rational expectations as an equilibrium concept. It suggests that while agents may not initially have perfect knowledge, they can learn and update their beliefs to eventually align with rational expectations. Key aspects of this approach include (i) the learning process of agents, (ii) their initial beliefs, and (iii) their willingness to change their views. The speed and extent of belief revision depend on agents’ confidence in their initial assumptions. Those with skepticism may quickly update their beliefs, rendering initial assumptions temporary, while those with strong convictions may be less likely to significantly revise their views.

In this paper, we study how initial beliefs shape expectations and influence macroeconomic dynamics in the context of economic policy. Specifically, we examine how misperceptions about new macroprudential policy—or changes to the existing one—affect both macroeconomic outcomes and the effectiveness of these policies. Our goal is to assess whether the potentially appealing features of new policies hold up when confronted with the reality of agents’ misexpectations about their impact.

We find that whether agents incorporate a new policy into their expectations at once or learn about it only gradually has a first-order effect on both the transitional dynamics and the

ultimate success of the policy.

1.1 What we do and what we find

We consider a variant of an RBC model with a collateral constraint to show the mechanism through which initial beliefs and confidence affect macroeconomic dynamics. We replace rational expectations agents with econometricians and study how their estimates that describe the economy’s perceived law of motion are affected by the prior beliefs and their variance. We use a simplified version of a collateral constraint model from Pintus and Suda (2019), who show that the interaction of financial markets and learning can partially explain both the onset and the severity of the crisis. The last global financial crisis brought calls for policy solutions that were not considered or even available before the crisis. Fiscal policy, monetary policy, and macroprudential policy were used to respond to the repercussions of the crisis itself, but it is largely macroprudential policy that is sought to address some of its underlying factors. For that reason, we focus our attention on macroprudential policy.¹

We assume that agents use constant-gain adaptive learning instead of rational expectations. The underlying premise is that households do not have perfect knowledge about the model and its parameters, including policy parameters, and they use historical data to both learn about them and forecast. Agents behave as econometricians and form expectations about the future, treating current realizations of macroeconomic variables as model-consistent linear functions of their past realizations. Every period, as the new data become available, they update the coefficients of these linear functions. Given this structure, agents’ expectations depend on the time-varying coefficients that represent agents’ beliefs and their current perception of how the economy works.

This method has two important implications. First, the initial beliefs (priors) may have lingering effects both for the subsequent evolution of beliefs and for the resulting dynamics of the entire system. Second, the introduction of a new policy or a modification of the existing one alters the law of motion and needs to be learned. This might be a lengthy process—especially if the policy change is not properly announced and explained—that yields dynamics and outcomes different not only from those under rational expectations, but also from those under learning with alternative priors. It creates an entirely new problem for the policy design and for its implementation, namely, how to deal with beliefs that encapsulate imperfect information about the economy and about how a policy affects it.²

To assess the importance of these two implications, we analyze how the introduction of a new economic policy is affected by the initial perception of and learning about that policy.

¹Although the broad question we address also applies to unconventional monetary policy, we abstract from this important dimension due to the lack of a canonical model. In contrast, both the collateral channel and macroprudential policies targeting leverage are now part of the macroeconomic toolkit, see Ampudia et al. (2021) and Millard et al. (2024) for overviews.

²Although one may get the impression that the issue revolves simply around communication, it goes beyond that in a world where both policymakers and the private sector have knowledge about how the economy works that is far from perfect.

In particular, we argue that agents’ understanding of the economy and their attitudes toward policy are shaped by their experiences. In the baseline analysis, we examine the effects of a new macroprudential policy that is introduced to reduce macroeconomic volatility when it is credibly implemented and fully understood by agents. In this scenario, the policy is highly effective and generates significant gains—often exceeding those predicted under rational expectations—when agents learn over time. The overall volatility caused by financial shocks is greatly reduced, and negative shocks do not cause large recessions and swings in land prices. We then consider a contrasting case where agents do not initially incorporate the new policy into their expectations. Instead, they gradually learn about the new economic mechanism as they observe its effects. Under these conditions, the dynamics of the economy differ significantly, and in some circumstances, the effectiveness of the policy in volatility reduction is completely eroded. We show that it is the initial beliefs and confidence in them that determine the initial response of the economy and its subsequent dynamics. The less “trust” agents have in their initial beliefs, the bigger the revisions of agents’ beliefs and the larger the impact on endogenous variables. Importantly, however, we find that the extent to which the introduction of a new policy can be deemed successful does vary depending on whether agents account for that in their perceived law of motion.

The scope of our analysis is specifically tailored to the unique conditions of the Global Financial Crisis and the potential repercussions of introducing countercyclical macroprudential policy in such a high-volatility, high-persistence environment. While we recognize that these conditions are extreme, our primary goal is to investigate a “stress-test” scenario where policy is implemented during an era characterized by significant misperceptions and systemic fragility. Consequently, the results presented here should be interpreted as a study of policy effectiveness during crisis-like regimes rather than a general baseline for standard business or financial cycles. This targeted approach allows us to explore the limits of macroprudential intervention when beliefs have drifted significantly from the rational expectations benchmark.

Our results have important implications for policy design, as illustrated by our analysis of the macroprudential policy change. The introduction of a new policy or the modification of an existing one is likely to alter the economy’s existing law of motion. The extent to which agents’ perceived law of motion adjusts to this change determines their understanding of the economy under the new policy and, consequently, its ultimate success. Policymakers should take this into account and consider ways to inform agents about how the policy operates.

Given the importance of financial markets in the recent global financial crisis, we consider the case of macroprudential policy as a simple but important example. Even though we consider a specific case of macroprudential policy, our results and conclusions likely apply more generally to fiscal and monetary policies.³ This has an immediate implication for the viability and success of any conventional or unconventional policy changes or reforms.

³See the next section for the literature related to the effect of learning on fiscal and monetary policy.

1.2 Related literature

Our paper contributes to a relatively small but growing literature that examines the implementation of new policies or structural policy shifts in the context of imperfect information and adaptive learning. While this literature has established that the transition to a new policy regime can trigger dynamics vastly different from those predicted under Rational Expectations (RE), existing studies have focused almost exclusively on monetary and fiscal policy. To our knowledge, this is the first paper to analyze the implementation of macroprudential policy within this framework.

Within the domain of monetary policy, a primary focus has been the introduction of new frameworks and the critical role of central bank communication. Honkapohja and Mitra (2020) study the credibility of a newly introduced policy in a model with learning. They consider the case of the introduction of price level targeting, where the credibility varies endogenously over time in response to the relative performance of inflation forecasting. Their focus is, however, on the expectational stability of the system rather than on the dynamics. Similarly, Honkapohja and McClung (2026) show the significant risks associated with the implementation of average inflation targeting policy without clear communication to the public. They show that if agents fail to understand the new policy, either because it was not properly communicated or due to their inability to properly account for the new policy, the economy may experience economic instability leading to explosive dynamics. Cole (2021) also analyzes the effectiveness of policy communication (approximating forward guidance) as a new policy tool under both rational expectations and adaptive learning. While forward guidance does appear highly effective under rational expectations because agents immediately internalize the policy’s future path, its introduction under adaptive learning results in slower stabilization, higher volatility, and significant overshooting of inflation and output targets. Ultimately, the study suggests that the effectiveness of this new tool is significantly overstated in standard models, as learning-based agents do not fully grasp the policy’s structural implications during its initial implementation.

These theoretical insights are complemented by empirical and experimental evidence. Shukayev and Amano (2025) explore how individuals adjust their inflation expectations in response to a sudden transition from inflation targeting (IT) to price-level targeting (PLT) within a controlled laboratory experiment. The study focuses on understanding whether participants grasp the implications of this monetary policy shift. They find that while in a simpler experimental model individuals showed a clear shift in behavior, in a richer and more realistic setting results were no longer statistically significant. Using a daily survey of U.S. consumers, Coibion et al. (2023) examine the Federal Reserve’s introduction of its Average Inflation Targeting (AIT) strategy in August 2020 and its impact on household expectations. They find that the Fed’s new policy had little effect on public understanding or economic outlooks, with most households unaware of the policy or its implications. They also show that households’ inflation expectations, as well as their outlooks on income and economic growth, remained unchanged after the AIT announcement suggesting that the new strategy did not influence public

economic expectations in any significant way. Conversely, using a randomized controlled trial on the Bundesbank Online Panel Households, Hoffmann et al. (2022) find that households can adjust their expectations under a hypothetical shift to AIT when properly informed, with trust in the central bank playing a key moderating role.

The effects of policy changes on the stability and dynamics of economic systems have been studied extensively in the context of fiscal policy. Evans et al. (2009) analyze the effects of fiscal policy changes that will take place in the future but are (credibly) announced in advance by the policymaker in a very stylized model with adaptive learning. They find that while agents' behavior may eventually align with the rational expectations equilibrium, the convergence can be slow, and the dynamics along the learning path can be qualitatively different than their RE counterpart. Their setup, however, is significantly different from ours as agents need to forecast only one endogenous variable (interest rate) while the second one (government spending) is credibly announced.

Mitra et al. (2013) focus on the question of anticipated versus unanticipated changes in fiscal policy when agents are learning, in the case of lump-sum taxation. They find that when agents combine knowledge about future policy with econometric forecasts of future wages and interest rates, model dynamics feature large and hump-shaped responses. Gasteiger and Zhang (2014) introduce distortionary taxation in an RBC model with learning, where agents know the path of fiscal policy instruments. They find that pre-announced permanent tax changes can lead to oscillatory dynamic responses of consumption and capital.

In a life-cycle context, Cottle Hunt (2021) demonstrates that policy uncertainty regarding Social Security reform generates cycles of overoptimism and overpessimism, significantly amplifying welfare effects compared to rational models as agents fail to anticipate how demographic and policy shifts fundamentally alter the aggregate economy. While our setup differs significantly in its structural focus, we find a qualitatively similar “boom-bust” pattern following the policy change, driven by the delayed alignment of expectations.

A related strand of the literature studies the transitional dynamics that follow a policy change under learning, without agents anticipating the change in advance or facing uncertainty about it. Giannitsarou (2006) studies the effects of an unanticipated and permanent elimination of capital income taxation in a neoclassical growth model with adaptive learning. She finds that the transition is asymmetrically sensitive to the technology shock realized at the time of the reform: a negative shock leads learners to interpret the reform as a fluctuation around the old steady state, causing a slow, recession-like transition, while a positive shock produces dynamics close to those under rational expectations. Ascari et al. (2017) similarly study an unanticipated policy change under learning—a shift in the central bank's inflation target—in a New Keynesian model. They show that a higher inflation target both reduces the region over which expectations are E-stable and slows the speed of convergence to the new equilibrium. They also find central bank transparency—the case in which agents understand how the nominal interest rate is set—helps anchor expectations, though its benefit diminishes

as the target rises. Our setup is closer in spirit to these two papers than to the anticipated policy change and policy uncertainty literature discussed above: like Giannitsarou (2006) and Ascari et al. (2017), we study the transition following a policy change under learning without assuming that agents anticipate the change or are uncertain whether it will occur.

Our paper is also related to the news shock literature started by Beaudry and Portier (2004) and Jaimovich and Rebelo (2009), that explores how information about future economic fundamentals can affect current economic decisions and outcomes. Implications of news shocks in a model with bounded rationality and adaptive learning are studied by Dombeck (2022). The author finds that the introduction of new information in the form of a news shock does not lead to expectational instability coming from learning.

The impact of initial beliefs on the dynamics of a macroeconomic system under learning is analyzed in Bullard and Suda (2016), who study expectational stability in systems with Bayesian learners. While they find that priors and their variance do not alter the ultimate stability of the system, these factors significantly influence its transitional dynamics. Similarly, Cogley and Sargent (2008) and Suda (2018) show that the precision of initial beliefs determines the speed of belief updating, which in turn affects asset prices and real quantities such as output and capital.

While those studies often utilize Bayesian learning about Markov transition matrices, we analyze the importance of initial beliefs for the subsequent learning process using a more standard version of adaptive learning that assumes less sophistication on the part of agents. Furthermore, our contribution lies in extending these insights beyond purely mechanical dynamics to emphasize the critical implications of belief variance for economic policy design.

Finally, we contribute to a vast and growing literature on macroprudential regulations, see Cerutti et al. (2017), Araujo et al. (2024) and Millard et al. (2024) for recent surveys. To our knowledge, ours is the first paper that examines the impact of initial beliefs and adaptive learning on the effectiveness of this class of policies.

The rest of the paper is structured as follows. In Section 2, we present a model with a collateral constraint and endogenous leverage, and introduce the role for macroprudential policy. Section 3 shows how the introduction of an economic policy designed to make leverage countercyclical and to reduce the negative effects of financial shocks may actually lead to higher volatility and the amplification of such shocks. In Section 4 we discuss the implications for the design and the deployment of new macroeconomic policies in general, and Section 5 concludes.

2 Model

In this section we present a small open economy version of a real business cycle model augmented with a collateral constraint and endogenous leverage.⁴ The incorporation of financial frictions allows us to capture the interplay between real and financial sectors, with borrowing capacity tied to land values. In this framework, we introduce a macroprudential policy to examine its role in mitigating the amplification effects of financial shocks under rational expectations and adaptive learning.

2.1 Setup

Consider a small open economy with a representative agent that maximizes its expected lifetime utility

$$E_0^* \sum_{t=0}^{\infty} \beta^t \frac{\left[C_t - \psi \frac{N_t^{1+\chi}}{1+\chi} \right]^{1-\sigma} - 1}{1-\sigma}, \quad (1)$$

where C_t is consumption, N_t is hours worked, σ denotes the relative risk aversion, χ measures the elasticity of labor supply, β is the discount factor, and E_t^* denotes (possibly non-rational) expectations. The representative agent faces a budget constraint

$$C_t + K_{t+1} - (1 - \delta)K_t + Q_t(L_{t+1} - L_t) + (1 + R)B_t = B_{t+1} + AK_t^\alpha L_t^\gamma N_t^{1-\alpha-\gamma}, \quad (2)$$

where K_{t+1} is the capital stock, L_t is the stock of land, Q_t denotes the land price, B_t stands for bonds/debt, and A is (constant) total factor productivity. Additionally, the agent faces a collateral constraint

$$\tilde{\Theta}_t E_t^*[Q_{t+1}] L_{t+1} \geq (1 + R)B_{t+1}, \quad (3)$$

where R is the real interest rate and $\tilde{\Theta}_t$ denotes leverage—the fraction of the expected value of land that can be used as collateral. Given the small open economy assumption, the real interest rate is exogenous and assumed to be constant.

Following Pintus and Suda (2019), we assume that the leverage, $\tilde{\Theta}_t$, is given by

$$\tilde{\Theta}_t \equiv \Theta_t \left\{ \frac{E_t^*[Q_{t+1}]}{Q} \right\}^\varepsilon \quad (4)$$

and has exogenous and endogenous components. The exogenous part is captured by Θ_t and we assume that $\theta_t = \log(\Theta_t)$ follows an AR(1) process,

$$\theta_t = (1 - \rho_\theta)\bar{\theta} + \rho_\theta\theta_{t-1} + \xi_t, \quad (5)$$

where ξ_t is the leverage shock. The endogenous part is assumed to depend on the expected deviation of land price Q_{t+1} from its steady-state value, Q , and ε captures the direction and

⁴This is a simplified version of Pintus and Suda (2019).

the extent to which asset prices affect the leverage.

In our model, equation (4) encompasses both the structure of the financial sector and the associated policy stance. In Section 3.2 we treat ε as a sufficient statistic for macroprudential policy. Using the US microdata, Mian and Sufi (2011) show that, in the period preceding the crisis, the leverage was procyclical with respect to house prices, implying a positive value for ε . Our policy experiment involves changing the value of ε so that the associated movement of the leverage becomes countercyclical.

The representative agent chooses C_t , N_t , L_{t+1} , K_{t+1} and B_{t+1} to maximize the lifetime utility in (1) subject to the budget constraint (2) and the collateral constraint (3). The first-order conditions for this problem are

$$C_t : \quad \left[C_t - \psi \frac{N_t^{1+\chi}}{1+\chi} \right]^{-\sigma} = \Lambda_t \quad (6)$$

$$N_t : \quad \psi N_t^{\chi+\alpha+\gamma} = (1 - \alpha - \gamma) A K_t^\alpha L_t^\gamma \quad (7)$$

$$L_{t+1} : \quad Q_t \Lambda_t = \beta E_t^*[Q_{t+1} \Lambda_{t+1}] + \beta \gamma E_t^*[\Lambda_{t+1} Y_{t+1} / L_{t+1}] + \Phi_t \tilde{\Theta}_t E_t^*[Q_{t+1}] \quad (8)$$

$$K_{t+1} : \quad \Lambda_t = \beta E_t^*[\Lambda_{t+1} (\alpha Y_{t+1} / K_{t+1} + 1 - \delta)] \quad (9)$$

$$B_{t+1} : \quad \Lambda_t = \beta(1 + R) E_t^*[\Lambda_{t+1}] + (1 + R) \Phi_t, \quad (10)$$

where Λ_t and Φ_t denote the Lagrange multipliers of constraints (2) and (3), respectively.

Taking the first-order conditions and log-linearizing them around the steady state allows us to write down a linear expectational system

$$X_t = \mathbf{A} X_{t-1} + \mathbf{B} E_{t-1}^* X_t + \mathbf{C} E_t^* X_{t+1} + \mathbf{D} \xi_t, \quad (11)$$

where $X_t = \{\tilde{c}_t, \tilde{q}_t, \tilde{\lambda}_t, \tilde{b}_t, \tilde{k}_t, \tilde{\theta}_t\}'$ is the vector of endogenous variables expressed in log deviations from the steady state; matrices $\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$, and $\mathbf{D}$ are functions of structural parameters; ξ_t is the exogenous leverage shock, and E_t^* denotes the potentially non-rational expectations.⁵ Importantly, matrix $\mathbf{B}$ captures the relationship between land price and debt (and ε) given by the collateral constraint.

2.2 Rational expectations equilibrium

Under rational expectations $E_t^* = E_t$ and for $1 + R < \frac{1}{\beta}$ there exists a unique rational expectations solution with a binding collateral constraint (3).⁶ The minimal state variable (MSV) solution to that system has a VAR(1) representation

$$X_t = \mathbf{M}^{RE} X_{t-1} + \mathbf{G}^{RE} \xi_t, \quad (12)$$

⁵See Appendix A for the derivations and exact definitions of matrices $\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$, and $\mathbf{D}$.

⁶We assume $1 + R < \frac{1}{\beta}$ so that the borrowing constraint is always binding.

where $\mathbf{M}^{RE}$ is the solution to

$$\mathbf{M} = [\mathbb{I} - \mathbf{C}\mathbf{M}]^{-1}[\mathbf{A} + \mathbf{B}\mathbf{M}] \quad (13)$$

and $\mathbf{G}^{RE}$ is given by

$$\mathbf{G}^{RE} = [\mathbb{I} - \mathbf{C}\mathbf{M}^{RE}]^{-1} \mathbf{D}. \quad (14)$$

The rational expectations hypothesis implies that both perceived and actual laws of motion are described by (12): agents' perceived probability distributions of endogenous and exogenous variables agree with the true distributions governing the system's dynamics. In this setting, a policy change—whether it is a fiscal policy involving the introduction of lump-sum or distortionary taxes, or a macroprudential policy taking the form of financial regulations constraining the leverage—alters one (or more) of matrices $\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$, and $\mathbf{D}$. This change is reflected, in turn, in the equilibrium described by matrices $\mathbf{M}^{RE}$ and/or $\mathbf{G}^{RE}$. The rational expectations hypothesis implicitly assumes that households understand this, and that their forecasts following the policy change are consistent with these new matrices.

2.3 Adaptive learning equilibrium

Our key insight, however, relies on the notion that agents may not have rational expectations, i.e., $E_t^* \neq E_t$. We assume, instead, that agents do not know the full structure of the economy and, as a result, they do not know the exact equilibrium dynamics of the economy.⁷

We think of agents as econometricians, who routinely estimate a forecasting model (using historical data) that describes their view of the world, the perceived law of motion (PLM). We assume that agents' PLM is formulated in a way that is consistent with the RE solution (12) and they use this PLM to form their predictions about the future.

Agents use the perceived law of motion that coincides with the rational expectations equilibrium

$$X_t = \mathbf{M}X_{t-1} + \mathbf{G}\xi_t \quad (15)$$

for forecasting. The key element that differentiates equation (15) from its rational expectations counterpart in equation (12) is that we do not assume that $\mathbf{M} = \mathbf{M}^{RE}$ and $\mathbf{G} = \mathbf{G}^{RE}$, i.e., agents' perception may not match one-to-one with the rational expectations solution.

Agents use their most recent estimate of the perceived law of motion to forecast the future state:

$$E_t^* X_{t+1} = E_t^* (\mathbf{M}X_t + \mathbf{G}\xi_t) = \mathbf{M}_{t-1}X_t, \quad (16)$$

with $\mathbf{M}_{t-1}$ denoting the estimate of $\mathbf{M}$ obtained using data up to date $t - 1$. We follow the standard timing convention in the adaptive learning literature (Marcet and Sargent, 1989; Evans and Honkapohja, 2001): at period t , when forming expectations about $t + 1$, agents use $\mathbf{M}_{t-1}$ rather than an updated $\mathbf{M}_t$, since $\mathbf{M}_t$ depends on the current value of X_t , which

⁷We do assume, however, that agents “know” the steady state.

is itself partly determined by today's expectations. Agents update their estimate only at the end of period t , once the period's endogenous variables—prices and quantities—have been determined. Excluding $\mathbf{M}_t$ from the time- t information set avoids simultaneous determination of the parameters defining agents' forecast functions and the current values of X_t , and allows X_t to be solved for recursively.⁸

Substituting expectations $E_t^* X_{t+1} = \mathbf{M}_{t-1} X_t$ and $E_{t-1}^* X_t = \mathbf{M}_{t-2} X_{t-1}$ into the equation (11), the actual law of motion (ALM) under adaptive learning is given by

$$[\mathbb{I} - \mathbf{C}\mathbf{M}_{t-1}] X_t = [\mathbf{A} + \mathbf{B}\mathbf{M}_{t-2}] X_{t-1} + \mathbf{D}\xi_t. \quad (17)$$

This equation combines structural parameters of the model (matrices $\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}$) with the beliefs (matrices $\mathbf{M}_{t-2}, \mathbf{M}_{t-1}$, and $\mathbf{G}_{t-1}$) to govern the dynamics of X_t . Note that the rational expectations equilibrium is a fixed point of that equation and the point, where the perceived and actual laws of motion of the economy coincide.

As econometricians, agents update their beliefs once new data becomes available using the recursive constant gain algorithm:

$$\mathbf{M}_t = \mathbf{M}_{t-1} + \nu \mathbf{R}_t^{-1} X_{t-1} (X_t - \mathbf{M}_{t-1} X_{t-1}) \quad (18)$$

$$\mathbf{R}_t = \mathbf{R}_{t-1} + \nu (X_{t-1} X_{t-1}' - \mathbf{R}_{t-1}), \quad (19)$$

where ν is a small constant.⁹ $\mathbf{M}_t$ is a time-varying matrix of coefficients and $\mathbf{R}_t$ is the associated variance-covariance matrix. Equation (18) describes the path of beliefs (represented by $\mathbf{M}_t$) given some initial beliefs $\mathbf{M}_0$. Similarly, equation (19) presents the evolution of the variance-covariance matrix of estimates $\mathbf{M}_t$ given the confidence in initial beliefs, $\mathbf{R}_0$.

Under rational expectations, the dynamics of all endogenous variables are determined entirely by the equation (12). The dynamics under adaptive learning are, in turn, jointly determined by equations (17)-(19) and conditional on $\mathbf{M}_0$ and $\mathbf{R}_0$.

Having set the stage, we can now analyze the effect of a new macroeconomic (macroprudential) policy aimed at reducing the volatility induced by financial shocks.

3 The dangers of macroprudential policy experiments

In this section, we show that while the introduction of macroprudential policy considerably reduces the volatility of endogenous variables under rational expectations, it may produce the opposite effect under learning. To illustrate this result, we first compare the economy's impulse responses to a financial shock before and after the introduction of such a policy, and then compute the standard deviations of macroeconomic variables following the policy change.

⁸Empirically, this assumption does not make a significant difference for the IRFs or for the dynamics of $\mathbf{M}_t$.

⁹Under recursive least squares, gain is a decreasing function of t with $\nu = \frac{1}{t}$. In section 3.6 we illustrate the effect of ν on macroeconomic volatility.

3.1 Calibration

For the numerical exercise we follow the calibration of Pintus and Suda (2019), see Table 1. Such calibration delivers average values for leverage ($\bar{\Theta} \approx 0.88$), debt-to-GDP ($\frac{B}{Y} \approx 0.52$) and land value-to-GDP ($\frac{QL}{Y} \approx 0.59$) ratios observed during the run-up in house prices in 1996Q1–2008Q4, that is, preceding the Global Financial Crisis. We set $\mu = 0.99$ to reflect the annual real interest rate of 4%, the time preference parameter to $\beta = 0.96\mu$, the inverse of labor elasticity to $\chi = 1/3$, the capital share to $\alpha = 0.33$, and the land share to $\gamma = 0.0093$ to deliver these ratios.

The persistence of leverage shocks is estimated using data spanning 1975Q1–2010Q1. To isolate the exogenous component of the leverage, we remove the portion explained by land prices as defined in Equation (4). We then estimate AR(1) processes on the log of the resulting series, θ_t , using two distinct methods to capture the beliefs of different agents. Under rational expectations, ρ_θ is given by the OLS estimate from a univariate regression over the full 1975Q1–2010Q1 sample. This approach aligns with the assumption that rational expectations (RE) agents know the true process governing leverage.

In contrast, the adaptive learning (AL) estimate is obtained from a constant gain learning of equation (5), in which agents forecast and update their beliefs in real time. We consider two specifications of learning. In the AL_1 case, we utilize the full-sample OLS estimate of the autocorrelation parameter, $\rho_\theta = 0.9756$. This matches the value used in the data-generating process and under rational expectations. In the AL_2 case, agents learn from the actual leverage data and overestimate the autocorrelation parameter, which reaches approximately $\hat{\rho}_\theta = 0.9904$ by 2008Q3. The higher persistence in AL_2 is therefore a result of agents’ real-time reactions to the data available during the peak of the crisis. These respective values are used to initialize the agents’ beliefs in the $\mathbf{M}_0$ matrix.

To initialize the learning process, we define $\mathbf{R}_0$ as the variance-covariance matrix derived from the rational expectations equilibrium. We compute this by simulating the RE model ($\mathbf{M}^{RE}$) one million times using the baseline calibration from Table 1. This estimated matrix provides the initial condition for $\mathbf{R}_0$ in our impulse response simulations under adaptive learning. Because the choice of $\mathbf{R}_0$ determines the speed of convergence toward the REE and the dynamics of the system under learning, we maintain a consistent $\mathbf{R}_0$ across all simulations and calibrations to ensure comparability of the results.

Table 1: Parameter values

Parameter		Value	Source/Target
(World) discount factor	μ	0.99	$R = 4\%$
Discount factor	β	0.96μ	Iacoviello (2005)
Depreciation rate	δ	0.025	
Capital share	α	0.33	Gertler et al. (2012)
Land share	γ	0.0093	$\frac{QL}{Y} \approx 0.59$
Leverage (steady state)	$\bar{\Theta}$	0.88	Pintus and Suda (2019)
Inv. labor elasticity	χ	$\frac{1}{3}$	Gertler et al. (2012)
(Constant) gain	ν	0.014	Pintus et al. (2021)
s.d. leverage shock	σ_{ξ}	0.033	Pintus et al. (2021)
RE , AL_1 persistence of leverage shock	ρ_{θ}	0.9756	Pintus and Suda (2019)
AL_2 persistence of leverage shock	$\hat{\rho}_{\theta}$	0.9904	Pintus and Suda (2019)

3.2 Macroprudential policy

The key parameter in our exercise is the land price elasticity of leverage ε . Using the data on 2002–2006 changes in house prices and debt-to-income, Mian and Sufi (2011) find evidence of a mildly procyclical (in housing prices) leverage. We set $\varepsilon = 0.5$ to match their results on the impact of housing price changes on the debt-to-income ratio. To examine the effects of the introduction of macroprudential policy, we vary the value of ε . Specifically, we set ε to negative values to approximate countercyclical measures.

While a more granular model incorporating a full financial sector would undoubtedly offer further insights into the structural transmission of such policies, we justify our approach by drawing on the results of simpler models featuring ex-post moral hazard and costly monitoring. Specifically, the case of procyclical leverage ($\varepsilon > 0$) is consistent with the framework of Aghion et al. (1999), while a countercyclical regime can be theoretically mapped to the implementation of procyclical taxes on lenders.

We adopt this ε parameterization for two primary reasons. First, it allows us to avoid the additional layers of expectational complexity and stability issues inherent in modeling explicit fiscal instruments under adaptive learning (see, for example, Evans et al., 2009; Gasteiger and Zhang, 2014; Cottle Hunt, 2021). Importantly, our treatment of expectations of ε would correspond to an unanticipated fiscal policy shock. This treatment is consistent with a further feature of our household sector. Households are always at their borrowing constraint and, in our calibration, they always want to borrow more than the constraint permits. Consequently, even complete knowledge of a future change in ε would likely not alter household behavior in advance of the change, since there is no unconstrained margin along which anticipation could operate. In our setting, the distinction between anticipated and unanticipated policy changes

is therefore unlikely to be quantitatively material. This need not hold in a richer model. With an explicit financial sector, where macroprudential policy operates through the behavior of banks or other financial intermediaries, and the borrowing constraint is an equilibrium outcome rather than a policy parameter, households or lenders could have a genuine margin to respond to anticipated changes. In that case, the anticipation effects emphasized by Evans et al. (2009) or Mitra et al. (2013) could become material.

Second, because ε does not influence the model's steady state, this setup remains consistent under both rational expectations and adaptive learning. This ensures that our results represent a “minimum departure” from the RE benchmark, allowing us to focus strictly on how different expectation formation schemes—rather than changes to the long-run equilibrium—determine the transitional impact of the policy.

3.3 Leverage shocks when macroprudential policy is known

Pintus and Suda (2019) show that learning can significantly amplify leverage shocks when agents' beliefs about the model parameters differ from rational expectations, see also Figure 1. In particular, if households overestimate the persistence of the financial shock process, ρ_θ , and the leverage is in fact mildly procyclical (represented by $\varepsilon = 0.5$), the financial shock ξ_t causes responses of output, capital, and consumption over 2.5 times larger under learning than under rational expectations. Not only does adaptive learning amplify economic shocks, but also the actual effect is quantitatively large: a large negative shock to leverage of about -5% reduces output by around 3.2% under learning but only by about 1.3% under rational expectations.

The foundation of the economy's large response to a leverage shock under learning lies in the interaction of the forecast of land prices with the borrowing constraint in equation (3). If a negative shock to leverage would not translate into a fall in land prices and, in turn, would not lower the value of collateral, resulting in less borrowing, the learning economy would behave like a RE economy. Moreover, eliminating the effect of land price swings on the borrowing constraint would not only bring the dynamics under learning closer to the ones under RE but would also reduce the response under rational expectations.


Consider now a macroprudential policy that makes leverage countercyclical in housing/land prices. In our model this assumption is governed by the negative value of ε in equation

$$\tilde{\Theta}_t \equiv \Theta_t \left\{ \frac{E_t^*[Q_{t+1}]}{Q} \right\}^\varepsilon, \quad (4)$$

which implies that an increase in expected housing/land prices leads to a decrease of the leverage.

If such a policy implies $\varepsilon = -0.75$ and adaptive learning agents have a correct understanding of how this new ε affects $\mathbf{M}_\tau$, the economy's dynamic responses under learning and RE are greatly reduced, and the path of the economy with adaptive learning is considerably closer to the one under rational expectations. Figure 1 depicts the impulse responses under

Figure 1: Responses under procyclical ($\varepsilon = 0.5$) leverage for learning (solid red) and RE (dotted blue) and countercyclical ($\varepsilon = -0.75$) leverage for learning (dashed purple) and RE (dashed-dotted black) in case of the overestimation of the persistence of leverage shocks ρ_θ , i.e. AL_2 learning.



Notes: The AL_2 learning case refers to constant gain learning, where the estimated persistence of the leverage shock $\hat{\rho}_\theta = 0.9904$ exceeds the data-generating rational expectations value, $\rho_\theta = 0.9756$.

adaptive learning and rational expectations before and after the introduction of such a policy. Learning amplifies the effects of leverage shocks on output and debt by a factor of 2.5 for procyclical leverage, but the difference almost disappears if the leverage is countercyclical. The introduction of the macroprudential policy reduces the debt response to the leverage shock from -35% to -6% under learning and from -16% to -6% under RE. This result seems to provide unequivocal support for the macroprudential policy if one wants to reduce fluctuations following financial shocks, especially in the case of imperfect information and learning.

The effects of the policy can also be observed by analyzing the volatility of macroeconomic variables under alternative values of ε for both the rational expectations and adaptive learning cases. Table 2 presents the standard deviations of four key variables—consumption, output, borrowing, and land price—under different informational assumptions, as well as for various values of ε , which captures the degree and nature (procyclical, acyclical, or countercyclical) of cyclicity in the borrowing constraint, as specified in equation (4). The results provide insights into how learning dynamics and macroprudential policy design influence economic stability. Similar to the case illustrated by the impulse response functions, the volatility of all variables is affected by whether agents operate under rational expectations (column 1) or adaptive learning (columns 2 and 3).

Column (2) presents results for the adaptive learning specification, AL_1 , in which agents' beliefs about the law of motion coincide with the rational expectations solution. Differences in standard deviations between column (1) and column (2) are therefore solely due to constant gain learning. Interestingly, the standard deviation of output is slightly lower under learning than under rational expectations, which suggests that constant gain learning may introduce a stabilizing effect in this particular case.

Column (3) shows results for the adaptive learning specification, AL_2 , in which agents overestimate the persistence of leverage shocks, a scenario corresponding to the conditions

Table 2: Volatility of macroeconomic variables under alternative ε under adaptive learning and rational expectations (median standard deviations)

	rational expectations (<i>RE</i>) $\rho_\theta = 0.9756$ (1)	adaptive learning (<i>AL</i> ₁) with $\hat{\rho}_\theta = \rho_\theta$ (2)	adaptive learning (<i>AL</i> ₂) with $\hat{\rho}_\theta > \rho_\theta$ (3)
$\varepsilon = 0.5$			
consumption, $\tilde{c}_t$	2.13	2.15	6.12
output, $\tilde{y}_t$	2.04	1.98	5.28
borrowing, $\tilde{b}_t$	35.69	38.07	119.32
land price, $\tilde{q}_t$	17.35	20.67	76.31
$\varepsilon = 0$			
consumption	1.18	1.19	1.71
output	1.13	1.07	1.43
borrowing	19.95	20.92	33.25
land price	9.76	11.57	21.86
$\varepsilon = -0.75$			
consumption	0.71	0.70	0.80
output	0.67	0.62	0.65
borrowing	11.94	12.09	15.30
land price	5.87	6.90	10.48
$\varepsilon = -1.5$			
consumption	0.50	0.50	0.52
output	0.48	0.43	0.41
borrowing	8.50	8.33	9.65
land price	4.19	4.89	6.85

Notes. The standard deviation represents the variability of a macroeconomic variable over 100 periods (equivalent to 25 years), driven solely by leverage shocks. Matrices describing the VAR(1) dynamics of series are initialized at the corresponding (to the parameter values) RE equilibria. The median is computed across 1000 simulations. *AL*₁ refers to constant gain learning, where the perceived persistence of the leverage shock matches the true value (i.e., the rational expectations value). *AL*₂ refers to constant gain learning, where the estimated persistence of the leverage shock exceeds the rational expectations value.

observed in 2008Q3.¹⁰ In this case, the higher volatility arises from both the endogenous response to shocks and the process of learning the “true” value of ρ_θ (understood as the rational expectations or full-sample estimation value). Under the AL_2 specification, the volatility of all variables is noticeably higher compared to the rational expectations framework.

From a modeling perspective, AL_1 is designed to isolate the “pure” effect of adaptive learning; in this specification, initial beliefs are perfectly consistent with the rational expectations (RE) benchmark, meaning any subsequent divergence is driven solely by the agents’ recursive updating process. The AL_2 specification, in turn, introduces an additional layer of data-driven misperception in the form of an initial overestimation of the leverage shock persistence. By adding this misperception on top of the AL_1 framework, we are able to analyze how the learning mechanism interacts with the specific type of heightened persistence observed during the Global Financial Crisis.

Examining specific variables, the results indicate that consumption and output experience relatively smaller increases in volatility compared to borrowing and land prices. Nonetheless, under the adaptive learning specifications, standard deviations of these variables still increase significantly, suggesting that agents’ misperceptions can propagate shocks through the real economy. Borrowing and land prices, on the other hand, are particularly sensitive to both ε and the learning framework. Under adaptive learning with a high ε , these variables experience much larger fluctuations, underscoring their vulnerability to financial shocks and misaligned expectations.

The degree of cyclicity in the borrowing constraint, as captured by ε , plays a critical role in determining the magnitude of volatility. A positive value of ε , which represents procyclical borrowing, leads to a substantial increase in the standard deviations of all variables, particularly borrowing and land prices. Procyclical leverage amplifies financial shocks, causing larger fluctuations in these variables. Conversely, when ε is negative, representing countercyclical borrowing behavior, volatility decreases. This stabilization effect is evident under both rational expectations, where agents fully internalize the implications of countercyclical policies, and adaptive learning.

Notably, under adaptive learning with an overestimated degree of persistence of leverage shocks (as was the case in 2008Q3), the effectiveness of countercyclical policy—measured in terms of volatility reduction—is significantly greater. For example, under $\varepsilon = -0.75$, standard deviations are about one-third the size, compared to those under $\varepsilon = 0.5$ for both the RE and AL_1 specifications across all four variables. However, under AL_2 , these standard deviations are eight times smaller for countercyclical leverage compared to procyclical leverage. This highlights the potential impact of introducing such a policy at the onset of a global financial crisis.

¹⁰See Pintus and Suda (2019).

3.4 The policy surprise

The previous results highlight that countercyclical macroprudential policies are highly effective in stabilizing the economy, provided that agents’ information sets are updated to reflect the new policy regime. Under both rational expectations and adaptive learning—particularly the AL_2 framework—policies that are correctly understood succeed in dampening borrowing during booms and encouraging it during downturns. In these scenarios, the policy aligns the dynamics of financial shocks closely with the rational expectations benchmark.

However, this stabilization fails to materialize if households are unaware of how a new policy alters the cyclicity of the collateral constraint. Such a disconnect occurs if policymakers fail to communicate the change effectively or implement measures in a discretionary, unannounced fashion. To examine this, we consider a scenario where policymakers introduce procyclical taxes—intended to result in a countercyclical leverage—but learning agents initially lack the information to account for it because the policy shift was not advertised. In this context, the agents’ perceived law of motion (PLM) fails to incorporate the updated negative ε , instead relying on the positive value from before the policy change. The resulting “policy surprise” is shown in Figures 2 and 3, which present the impulse responses for both the “true” (AL_1) and (over)estimated (AL_2) persistence of the financial shock.¹¹

Following the negative leverage shock, we observe a considerably larger fall in debt under learning than under rational expectations, but this negative effect drops considerably faster than under the case of a procyclical leverage. However, under adaptive learning the response of macroeconomic variables to the financial shock is significantly different.


Figure 2 examines the effects of a 5% negative leverage shock under two belief frameworks: rational expectations and adaptive learning AL_1 , where agents misunderstand the countercyclical nature of macroprudential policy. Under AL_1 , agents fail to recognize the policy’s countercyclical stance, leading to significant deviations in macroeconomic responses compared to RE . The analysis considers two levels of policy countercyclicity: mildly countercyclical ($\varepsilon = -0.5$) and strongly countercyclical ($\varepsilon = -1.5$).

The response of output (Panel a) shows that, under AL_1 , output declines more sharply and recovers more slowly than under RE . For mildly countercyclical policy, the initial drop in output is deeper under AL_1 , and the recovery remains sluggish, reflecting agents’ slower adjustment to the policy environment. When the policy becomes strongly countercyclical, output recovers faster in both frameworks, but initial volatility under adaptive learning is heightened due to the gradual updating of agents’ beliefs. Consumption dynamics (Panel b) mirror those of output. The immediate drop in consumption is larger under adaptive learning for both levels of ε , driven by heightened uncertainty and slower belief adjustments. Stronger countercyclicity amplifies this volatility, delaying consumption stabilization under AL .


Land prices (Panel c) highlight further destabilizing effects of adaptive learning. After

¹¹It is important to remember that this “overestimated” value of the leverage shock persistence corresponds to the actual data estimate of the leverage process at 2008Q3.


Figure 2: Responses to a negative leverage shock for mildly countercyclical ($\varepsilon = -0.5$) leverage under AL_1 learning (solid red) and RE (dotted blue) and strongly countercyclical ($\varepsilon = -1.5$) leverage under learning (dashed purple) and RE (dashed-dotted black) given the incorrect beliefs regarding the macroprudential policy.




(a) Output after 5% leverage shock.



(b) Consumption after 5% leverage shock.




(c) Land price after 5% leverage shock.



(d) Debt after 5% leverage shock.

Notes: The AL_1 learning case refers to constant gain learning, where the estimated persistence of the leverage shock $\hat{\rho}_\theta$ in period 0 of IRF equals the data-generating rational expectations value, $\rho_\theta = 0.9756$.

Figure 3: Responses to a negative leverage shock for mildly countercyclical ($\varepsilon = -0.5$) leverage under AL_2 learning (solid red) and RE (dotted blue) and strongly countercyclical ($\varepsilon = -1.5$) leverage under AL_2 learning (dashed purple) and RE (dashed-dotted black) given the incorrect beliefs regarding the macroprudential policy.



Notes: The AL_2 learning case refers to constant gain learning, where the estimated persistence of the leverage shock $\hat{\rho}_\theta = 0.9904$ exceeds the data-generating rational expectations value, $\rho_\theta = 0.9756$.

the shock, land prices decline more sharply and recover more slowly under AL_1 than under RE , particularly when the policy is strongly countercyclical. This reflects the uncertainty introduced by agents' misperceptions of the policy's nature. Finally, debt (Panel d) exhibits the most pronounced differences. The decline in debt is steeper under AL_1 , reflecting greater sensitivity of borrowing to leverage shocks. While debt begins to recover under both frameworks, the recovery under AL_1 is significantly slower, especially with strongly countercyclical policies.

These results suggest that incorrect beliefs regarding the macroprudential policy exacerbate the negative effects of leverage shocks. Adaptive learning, in particular, introduces greater volatility and slower recovery across key macroeconomic variables compared to rational expectations. This implies that the stabilizing effects of countercyclical macroprudential policies are undermined when agents do not fully understand the policy's cyclicity.

Figure 3 explores the impact of an additional layer of incorrect beliefs, where agents overestimate the persistence of a leverage shock. This misperception interacts with their misunderstanding of the countercyclical macroprudential policy, amplifying the differences between

adaptive learning (AL_2) and rational expectations (RE). Under AL_2 , these compounded incorrect beliefs lead to pronounced destabilizing effects across macroeconomic variables.

The response of output and consumption (Panels a and b) reveals striking dynamics. Initially, the decline in output is smaller under AL than under RE . However, approximately 12 periods after the shock, output under AL_2 unexpectedly surges into a “boom,” despite the negative leverage shock. Similarly, the consumption dynamics feature overshooting. These responses reflect the destabilizing influence of incorrect beliefs, as agents misjudge both the shock’s persistence and the policy’s stance. In contrast, output and consumption under RE follow smoother and more stable recovery trajectories.

Land prices and debt (Panel c and Panel d, respectively) exhibit different dynamics. The overestimation of shock persistence leads to a deeper and more prolonged decline in land prices compared to RE for mildly countercyclical ε . However, for strongly countercyclical policy, the initial response and the dynamics after the shocks are much closer to RE and to the ones observed in Figure 2.

Overall, this scenario illustrates how compounding incorrect beliefs—both about the policy stance and the shock’s persistence—intensify volatility and destabilize macroeconomic dynamics under adaptive learning for the case of output and consumption, in stark contrast to the smoother responses observed under rational expectations.

Taken together, the results from Figures 2 and 3 demonstrate that the effectiveness of macroprudential policies depends critically on the expectations framework and the accuracy of agents’ beliefs. When agents operate under adaptive learning and hold incorrect beliefs about the policy stance and the persistence of shocks, the intended stabilizing effects of countercyclical leverage policies are significantly weakened. Instead of mitigating volatility, such policies can inadvertently amplify fluctuations in output, consumption, land prices, and debt.

These results are also evident when examining the overall variability of macroeconomic variables. To quantify the effects of policy changes on volatility, we use simulations. Specifically, we generate 1000 series of leverage shocks, each with an estimated standard deviation of 3.3%, and compute the resulting dynamics of all endogenous variables under different informational assumptions. The standard deviations reported in Table 3 represent the variability of each macroeconomic variable across different subperiods, driven solely by leverage shocks. The reported values correspond to the median across all 1000 simulations.

Table 3 presents the effects of changes in the macroprudential policy stance, represented by a shift from ε_0 (columns (1)–(3)) to ε_1 , on the standard deviations of key macroeconomic variables under different expectations frameworks: rational expectations in columns (4), (7), (10), adaptive learning with correct beliefs about shock persistence, AL_1 , in columns (5), (8), and (11), and adaptive learning with overestimated persistence of shocks, AL_2 , in columns (6), (9), and (12). We track the changes of volatility by looking at standard deviations over the first 16 quarters after the change (columns (4)–(6)), next 40 quarters (columns (7)–(9)), and another 40 quarters starting 60 years after the change (columns (10)–(12)). The simulation

Table 3: Effects of policy change on volatility (median standard deviation, %)

	before change			16 periods after			periods 17–56			periods 241–280		
	<i>RE</i>	<i>AL</i> ₁	<i>AL</i> ₂	<i>RE</i>	<i>AL</i> ₁	<i>AL</i> ₂	<i>RE</i>	<i>AL</i> ₁	<i>AL</i> ₂	<i>RE</i>	<i>AL</i> ₁	<i>AL</i> ₂
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$\varepsilon_0 = 0.5 \rightarrow \varepsilon_1 = 0$												
<i>consumption</i>	2.08	2.08	5.98	1.12	1.60	4.82	1.18	1.77	5.92	1.14	1.71	4.68
<i>output</i>	1.99	1.92	5.18	0.99	1.20	3.46	1.11	1.32	4.46	1.09	1.32	3.05
<i>borrowing</i>	37.23	39.50	127.77	20.00	28.28	84.52	20.83	29.44	74.27	20.40	28.12	61.59
<i>land price</i>	18.06	21.11	81.19	9.64	19.01	70.46	9.99	19.36	62.67	9.88	18.25	51.95
$\varepsilon_0 = 0.5 \rightarrow \varepsilon_1 = -0.5$												
<i>consumption</i>	2.08	2.08	5.98	0.93	1.45	5.67	0.85	1.92	7.59	0.78	1.63	4.65
<i>output</i>	1.99	1.92	5.18	0.77	1.03	5.39	0.79	1.46	6.44	0.75	1.15	3.51
<i>borrowing</i>	37.23	39.50	127.77	13.82	19.25	48.77	14.43	19.99	41.89	14.13	18.59	33.69
<i>land price</i>	18.06	21.11	81.19	6.67	18.16	68.97	6.93	18.29	59.01	6.85	16.43	46.13
$\varepsilon_0 = 0.5 \rightarrow \varepsilon_1 = -0.75$												
<i>consumption</i>	2.08	2.08	5.98	0.89	1.44	6.56	0.76	2.11	8.55	0.68	1.65	4.71
<i>output</i>	1.99	1.92	5.18	0.73	1.21	6.88	0.69	1.72	7.64	0.65	1.21	3.94
<i>borrowing</i>	37.23	39.50	127.77	11.96	14.98	31.16	12.51	15.43	26.71	12.23	14.39	21.62
<i>land price</i>	18.06	21.11	81.19	5.77	17.83	67.56	6.01	17.83	57.24	5.93	15.55	43.12
$\varepsilon_0 = 0.5 \rightarrow \varepsilon_1 = -1.5$												
<i>consumption</i>	2.08	2.08	5.98	0.88	1.89	9.70	0.58	2.83	11.35	0.48	1.83	5.27
<i>output</i>	1.99	1.92	5.18	0.67	2.15	11.45	0.52	2.73	11.26	0.46	1.68	5.30
<i>borrowing</i>	37.23	39.50	127.77	8.57	2.12	20.69	8.91	2.55	15.73	8.70	4.21	8.43
<i>land price</i>	18.06	21.11	81.19	4.10	17.08	63.73	4.29	16.75	51.40	4.23	13.22	35.75
$\varepsilon_0 = 0 \rightarrow \varepsilon_1 = -0.75$												
<i>consumption</i>	1.16	1.16	1.67	0.66	0.87	1.52	0.70	1.04	2.01	0.69	0.97	1.39
<i>output</i>	1.11	1.05	1.40	0.58	0.61	1.32	0.66	0.77	1.64	0.65	0.70	0.99
<i>borrowing</i>	20.81	21.69	35.54	11.98	13.09	18.73	12.51	13.54	16.28	12.23	13.04	14.28
<i>land price</i>	10.16	11.85	23.25	5.81	10.57	20.51	6.02	10.69	17.67	5.93	9.84	14.61

Notes: The standard deviation represents the variability of a macroeconomic variable over 40 periods (equivalent to 10 years) in columns (1)-(3), (7)-(12) and 16 periods (equivalent to 4 years) in columns (4)-(6), driven solely by leverage shocks. Matrices describing the VAR(1) dynamics of series are initialized at the corresponding (to the parameter values) RE equilibria. Series are simulated over 400 periods with 60 first observations being discarded. The median is computed across 1000 simulations. *AL*₁ refers to constant gain learning, where the perceived persistence of the leverage shock at the beginning of simulations (period 0) matches the true value (i.e., the rational expectations value ρ_θ). *AL*₂ refers to constant gain learning, where the estimated persistence of the leverage shock ($\hat{\rho}_\theta$) at the beginning of simulations (period 0) exceeds the rational expectations value. Agents engage in the learning process before the policy change occurs.

incorporates a pre-change initialization period of 100 periods with the first 60 observations being discarded. During this window, agents engage in active learning, and the perceived persistence of leverage shocks, $\hat{\rho}_\theta$, is updated recursively based on realized data. This ensures that at the moment of the policy change, the belief structure is not arbitrary but is grounded in the historical dynamics of the economy. Following the policy shift, $\hat{\rho}_\theta$ continues to evolve freely for both the *AL*₁ and *AL*₂ specifications. This allows agents to gradually learn from the realization of shocks if they have initially over- or underestimated the persistence of the new environment.¹²

¹²While robustness checks—imposing the policy change as early as period 3 or 10—yielded qualitatively

The table allows for a comparative analysis of the stabilizing or destabilizing effects of the policy change across these alternative specifications.

Under the rational expectations framework, the results indicate that a shift from procyclical ε_0 to a countercyclical ε_1 reduces the standard deviation of most macroeconomic variables, including output, consumption, and debt. This reflects the intended stabilizing effects of countercyclical macroprudential policies. By tightening leverage constraints during booms and relaxing them during downturns, the policy dampens fluctuations, particularly in financial variables such as land prices and debt. However, the magnitude of these reductions depends on the extent of the policy change, with larger shifts in ε leading to greater stabilization. Importantly, under RE, agents fully internalize the countercyclical nature of the policy, allowing the policy to achieve its objectives without significant unintended consequences.

In contrast, under the adaptive learning framework with correct beliefs about shock persistence (AL_1), the stabilizing effects of the policy are generally weaker. While the standard deviations of output, consumption, and debt decline following the policy change, these reductions are smaller compared to the RE framework. This occurs because agents under learning frameworks update their expectations gradually, leading to delayed adjustments in behavior. Furthermore, the results indicate that for a change to strongly countercyclical policy ($\varepsilon_1 = -1.5$) in certain variables, such as output and consumption, the policy change may even result in a temporarily higher volatility under AL_1 than under RE. Nevertheless, the overall direction of the effects under learning remains consistent with the RE framework, albeit with muted magnitude.

The adaptive learning framework with overestimated persistence of shocks (AL_2) produces the most distinct results. In this case, the policy change from ε_0 to ε_1 often leads to an increase in the standard deviation of key variables, contrary to the intended stabilizing effects of the policy. For example, right after the policy change to $\varepsilon_1 = -0.75$ and $\varepsilon_1 = -1.5$, the standard deviations of output and consumption rise significantly under AL_2 , highlighting the destabilizing effects of incorrect beliefs about the duration of shocks. These results reflect the interaction between agents' misperceptions about shock persistence and their slower adaptation to the countercyclical nature of the policy. Under AL_2 , agents may overreact to the policy-induced changes, amplifying fluctuations instead of dampening them. This underscores the risks associated with implementing countercyclical macroprudential policies in environments where agents hold incorrect beliefs about the economic environment. However, financial variables such as land prices and debt do not exhibit such heightened volatility following the policy change.

The heightened volatility observed in the AL_2 specification is highly specific to the environment of the Global Financial Crisis, where both the persistence and the perceived persistence of shocks were exceptionally high. While these conditions delay the return to the steady state, the discrepancies between the AL_1 and AL_2 specifications are primarily a byproduct of slow

identical patterns of volatility and amplification, the 40-period window is maintained to clearly illustrate the transition from an established learning state to the post-change regime.

belief updating that diminishes over longer horizons. Under conditions of lower shock persistence or when the policy response is delayed, the initial overreaction is dampened and the two learning specifications align more rapidly regardless of the gain parameter. As the simulation horizon is extended, these differences narrow consistently, confirming the models' asymptotic equivalence despite the medium-run persistence of learning effects during crisis episodes.¹³

Overall, the results from Table 3 demonstrate that the effectiveness of macroprudential policy in reducing economic volatility depends critically on the expectations framework. While the policy change yields clear stabilizing effects under rational expectations, these effects are weaker under adaptive learning with correct beliefs and may even reverse under adaptive learning with incorrect beliefs. This highlights the importance of ensuring that agents have accurate information about the policy regime and its implications. In particular, the destabilizing effects observed under AL_2 emphasize the need for careful communication and transparency in policy implementation to mitigate the risks associated with learning dynamics and belief misalignment.

Table 3 shows that, for some variables, the volatility following the policy change can temporarily be higher than before the policy shift. A key question raised by this analysis is how long it takes for the policy change to affect volatility and achieve its goal of reducing it. Figure 4 illustrates the dynamics of volatility following a transition from a procyclical $\varepsilon = 0.5$ to a countercyclical $\varepsilon = -0.75$ leverage under adaptive learning, AL_2 , where agents hold incorrect beliefs about the macroprudential policy.


We compute the rolling standard deviations over 40-period (10-year) windows starting from the policy change in period 100. The figure reveals that for output and consumption, the policy change initially raises volatility, as reflected by the increase in rolling standard deviations shortly after the policy shift. This suggests a transitional adjustment period during which agents adapt to the new policy framework, leading to temporary disruptions. Over time, however, the volatility of output and consumption begins to decline, eventually falling below the baseline levels (dashed blue line) observed prior to the policy change. In contrast, for land prices and debt, the volatility decreases immediately following the policy shift, with rolling standard deviations dropping below their pre-policy change levels almost instantly. This indicates that the stabilizing effects of the countercyclical leverage policy are more immediate and pronounced for these variables.

3.5 Confidence in policy

To draw the impulse response functions in Figures 2 and 3 we assume that at the moment of the deployment of the countercyclical macroprudential policy households are completely oblivious to that change. From the perspective of the model, this implies that at that very moment not only matrix $\mathbf{M}_{policy\ change}$ corresponds to the RE case matrix, $\mathbf{M}_{\varepsilon>0}^{RE}$, with a

¹³Detailed results confirming these convergence properties and the sensitivity analysis regarding shock persistence are provided in Tables OA.1 and OA.2 of the Online Appendix.

Figure 4: Standard deviation of endogenous variables following the change from procyclical ($\varepsilon = 0.5$) leverage to countercyclical ($\varepsilon = -0.75$) leverage under AL_2 learning (solid red) given the incorrect beliefs regarding the macroprudential policy.



Notes: The AL_2 learning case refers to constant gain learning, where the estimated persistence of the leverage shock $\hat{\rho}_\theta = 0.9904$ exceeds the data-generating rational expectations value, $\rho_\theta = 0.9756$.

procyclical leverage ($\varepsilon > 0$), but also the variance-covariance matrix $\mathbf{R}_{policy\ change}$ corresponds to the RE dynamics.

Could some form of policy announcements that make agents more open to change—and potentially learn faster (corresponding to a decrease in confidence in matrix $\mathbf{M}_{policy\ change}$)—reduce the boom or speed up the convergence to the new $\mathbf{M}_{\varepsilon < 0}^{RE}$? Figure 5 illustrates the case where $\mathbf{M}_{policy\ change} = \mathbf{M}_{\varepsilon > 0}^{RE}$, but households remain uncertain about the new relationships between land prices, debt, and the rest of the economy. For this figure, we scaled up the variance-covariance matrix, $\mathbf{R}_{policy\ change}$, thereby increasing the impact of forecast errors on $\mathbf{M}$ (and enhancing the speed of learning) without altering the rate at which confidence is built (and matrix $\mathbf{R}_t$ is updated).¹⁴ If the policymaker manages to announce its new policy in a way that makes agents more open and more likely to update their beliefs with the new data (less confidence approximated by larger variance), the dynamics can be even more surprising with an even larger economic boom and an increase of debt. It is clear from Figures 2, 3 and 5 that the introduction of new macroprudential regulations or conducting the announced policy experiments without properly addressing how they affect the economy and, therefore, guiding agents' beliefs, can result in significantly higher volatility than in the absence of such a change.¹⁵

3.6 Robustness

To measure the extent to which these effects depend on the specification of the gain parameter, we evaluate alternative specifications. Figure 6 presents the (rolling) standard deviation for alternative specifications of the gain parameter.


The speed of learning and the duration of incorrect beliefs are influenced by the gain parameter in the learning specification. To evaluate the extent to which heightened volatility and its duration depend on the gain, we analyze alternative gain specifications. Specifically, we consider both fast learning (with gain parameters of 0.035 and 0.025) as well as our baseline learning (with a gain parameter of 0.014). Additionally, we examine a recursive least squares (RLS) learning specification, where the gain evolves as $\nu = 1/t$. For the RLS case, we analyze two scenarios: one where estimation begins at period 1 (and $\nu = 1/t$) and another where estimation starts shortly before the policy change (in which $\nu = 1/(t - 90)$), with the policy change occurring at period 100.

Figure 6 presents the results of these alternative specifications. Consistent with the baseline case, the volatility of output increases substantially following the policy change. However, the gain parameter has an important effect on the dynamics: it influences both the level of pre-policy change volatility and the speed at which post-policy change volatility declines. While the pre-policy change level of volatility varies across gain specifications, a common pattern emerges. At some point after the policy change, volatility decreases below its initial pre-policy


¹⁴Numerically, this corresponds to increasing the gain parameter ν in equation (18) while keeping the gain parameter for equation (19) unchanged.

¹⁵In the Appendix C we consider the impulse responses and volatility under partially correct beliefs.


Figure 5: Responses to a negative leverage shock for mildly countercyclical ($\varepsilon = -0.5$) leverage under AL_2 learning (solid red) and RE (dotted blue) and strongly countercyclical ($\varepsilon = -1.5$) leverage under AL_2 learning (dashed purple) and RE (dashed-dotted black) given the incorrect beliefs regarding the macroprudential policy but with less confidence.




(a) Output after 5% leverage shock.



(b) Consumption after 5% leverage shock.



(c) Land price after 5% leverage shock.




(d) Debt after 5% leverage shock.

Notes: The AL_2 learning case refers to constant gain learning, where the estimated persistence of the leverage shock $\hat{\rho}_\theta = 0.9904$ exceeds the data-generating rational expectations value, $\rho_\theta = 0.9756$.

level, indicating that the policy eventually achieves its intended stabilizing effects. While this analysis highlights the important role of the gain parameter in determining the persistence of heightened volatility and the time horizon over which the policy's benefits are realized, it does not change the overall conclusion: regardless of the gain specification, the policy change is associated with an increase in volatility.

Figure 6: The median standard deviation of output following the change from procyclical ($\varepsilon = 0.5$) leverage to countercyclical ($\varepsilon = -0.75$) leverage (in period 100) under AL_2 learning under alternative gain parameter: $\nu = 1/t$ (blue solid line); $\nu = 1/(t - 90)$ (black thin line); $\nu = 0.014$ (red dashed line); $\nu = 0.025$ (dotted-dashed grey line); and $\nu = 0.035$ (dark yellow thin line).



Note: The standard deviation is computed over a rolling window of 40 periods (10 years), starting from the policy change in period 100. The displayed results represent the median across 1000 simulations, while the dotted lines indicate the median values of the standard deviations prior to the policy change. AL_2 refers to constant gain learning, where the estimated persistence of the leverage shock exceeds the rational expectations value.

4 The implication for the design and introduction of new macroeconomic policy

The evolution of agents' beliefs—and, consequently, the dynamics of endogenous variables—depends critically on the confidence agents have in their initial beliefs. When agents have high confidence and low variance in their initial beliefs (represented by the initial precision of the variance-covariance matrix), there is less revision observed in the data. Conversely, if agents are uncertain about the specific values in matrix $\mathbf{M}$, their belief revision can be significant.

This result has important implications for policy design.

As our findings demonstrate, the effectiveness of the proposed policy is heavily contingent on the underlying state of the economy and the agents' belief structures. Accordingly, we view this exercise as being specifically tailored to the unique conditions of the Global Financial Crisis and the potential repercussions of introducing countercyclical macroprudential policy in such a high-volatility, high-persistence environment. While a lower persistence value—as explored in our sensitivity analysis in Table OA.2—would lead to a faster alignment of learning specifications and lower volatility, our primary goal is to investigate the “stress-test” scenario where policy is implemented during an era of significant misperceptions and systemic fragility. By anchoring our calibration to this period of heightened persistence, we highlight how the interaction between learning and crisis-level shocks can create substantial, albeit transitory, volatility.

More generally, any change in economic policy—or the introduction of a new parameter that alters the existing perceived law of motion—requires agents to form new perceptions. If agents have no prior data on which to base these beliefs, policymakers must prioritize clear communication regarding how the policy will operate. Furthermore, strict adherence to the newly introduced policy is essential to reinforce agents' expectations about its effects on economic outcomes, particularly during periods of profound economic turbulence and of deep parameter uncertainty.

Consider the case of a central bank announcing that it will incorporate additional information or dynamics when setting the nominal interest rate. Agents will only confirm the credibility of the announcement if the central bank consistently adheres to the new policy. If the policy is not followed—either due to deliberate deviation by policymakers or implementation challenges—agents will revise their perceptions of the policy. This is particularly relevant when policymakers announce a new policy but are unable to implement it effectively. For instance, imagine a central bank promising to keep interest rates low even when a traditional Taylor rule would suggest raising them. Until such a deviation from the rule is observed in practice, households and firms cannot update their beliefs about the credibility of this policy. Thus, for successful implementation, policymakers must “walk the walk” after “talking the talk.”

Our findings emphasize a key implication for any new macroeconomic policy: the choices of households and firms are based on both current and expected future economic conditions, which influence their objectives and constraints. Relaxing the assumption of rational expectations and adopting an adaptive learning framework reveals that these expectations are formed based on subjective probability distributions and the agents' perceived law of motion. Importantly, these may differ from the actual law of motion and the true probability distributions of endogenous and exogenous variables.

As agents update their PLM, their forecasts evolve over time, shaped not only by changes in the economic environment but also by revisions in their perception of the linkages within that

environment. This makes the implementation of any policy contingent on agents’ understanding of the policy and how it unfolds. In some cases, this process poses minimal challenges—for example, when a government announces and enacts a change in the tax rate. Agents typically have prior knowledge of how taxes work, and variations in tax rates are unlikely to introduce significant uncertainty regarding individual or aggregate constraints.

However, the introduction of entirely new policies—such as unconventional monetary policy measures, macroprudential regulations, or changes to the Taylor rule—can generate much greater uncertainty. This uncertainty pertains to (i) how the new policy operates, (ii) the channels through which it affects the economy, and (iii) the details of its implementation. These factors are critical to the expectation formation process. Our results indicate that agents’ perceptions, along with the uncertainty associated with those perceptions, significantly influence how an economy responds to new policies at the time of their announcement, introduction, and, most importantly, implementation.

At the same time, the findings on the potential failure of macroprudential policies (discussed in Section 3) highlight the importance of properly accounting for agents’ beliefs. Clear communication and transparency are essential for the effective implementation of macroprudential policies.¹⁶ If agents are unaware of a policy’s countercyclical nature or misjudge its implications, the policy may fail to achieve its intended objectives. Policymakers must ensure that agents have accurate information about the policy regime and its expected effects. This may involve explicitly announcing and publicizing policy changes, as well as providing guidance to help agents form correct expectations.

Moreover, caution is advised when implementing strongly countercyclical leverage policies, as such policies can amplify volatility under learning—particularly when agents hold incorrect beliefs. In a way, we extend Lucas’ critique by adding a new dimension: it is not enough to account for the possibility that agents may alter their behavior in response to a policy change. Policymakers must also recognize that adaptation to the new policy takes time, as agents must learn how the policy operates. Properly accounting for expectations and initial beliefs is therefore paramount for the successful implementation of any policy.

5 Conclusions

In an environment characterized by imperfect information and learning, initial perceptions can critically shape macroeconomic dynamics. Using calibrated models, we show that the degree of confidence that households have in their perception of the law of motion has a large impact on how this perception changes and on the economy itself. We then show that the deployment of even the best policies can be very costly.

The Lucas critique put rational expectations to the forefront of macroeconomics and revolutionized how economists think about expectations and economic policy. It led to important

¹⁶This point is further illustrated in Appendix C.

results, from rules versus discretion, to policy-dependent determinacy of equilibria, to the importance of public versus private signals, with far-reaching implications for policy design and implementation.

Our analysis centers on the challenges of policy design under conditions of significant systemic fragility. While our calibration reflects extreme circumstances, it is essential for constructing policies that are robust to agent misperceptions. By focusing on these crisis-like regimes rather than a standard business cycle baseline, we highlight how the design of countercyclical tools must account for the heightened persistence of beliefs. Consequently, our findings offer a framework for designing interventions that remain effective even when the economy departs significantly from the rational expectations benchmark.

Our analysis abstracts from an explicit banking or financial sector, so that macroprudential policy operates directly on households' borrowing constraint rather than through bank behavior. It does mean, however, that the model is silent on how lenders would respond to policy, as well as on the anticipation and amplification effects that could arise in a richer setting with an explicit financial sector. In such a setting—whether modeled through banks or other financial intermediaries—the supply side of credit, not just the demand side, could adapt under learning. Modeling this explicitly would require a different model structure than the one developed here, and we leave it for future work.

We should not forget, however, that even temporary deviations from rational expectations can change what one would consider good or desirable policy. Designing macroeconomic policy that is robust to such deviations may be difficult, but it is a worthwhile step toward better policy design.

References

- Aghion, P., A. Banerjee, and T. Piketty (1999). Dualism and macroeconomic volatility. *The Quarterly Journal of Economics* 114(4), 1359–1397.
- Ampudia, M., M. Lo Duca, M. Farkas, G. Perez-Quiros, M. Pirovano, G. Rünstler, and E. Tereanu (2021, May). On the effectiveness of macroprudential policy. ECB Working Paper Series 2559, European Central Bank.
- Araujo, J., M. Patnam, A. Popescu, F. Valencia, and W. Yao (2024). Effects of macroprudential policy: Evidence from over 6000 estimates. *Journal of Banking & Finance* 169, 107273.
- Ascari, G., A. Florio, and A. Gobbi (2017). Transparency, expectations anchoring and inflation target. *European Economic Review* 91, 261–273.
- Beaudry, P. and F. Portier (2004). An exploration into Pigou's theory of cycles. *Journal of Monetary Economics* 51, 1183–1216.

- Bullard, J. and J. Suda (2016). The stability of macroeconomic systems with Bayesian learners. *Journal of Economic Dynamics and Control* 62(C), 1–16.
- Cerutti, E., S. Claessens, and L. Laeven (2017). The use and effectiveness of macroprudential policies: New evidence. *Journal of Financial Stability* 28, 203–224.
- Cogley, T. and T. J. Sargent (2008). The market price of risk and the equity premium: A legacy of the Great Depression? *Journal of Monetary Economics* 55(3), 454–476.
- Coibion, O., Y. Gorodnichenko, E. S. Knotek, and R. Schoenle (2023). Average inflation targeting and household expectations. *Journal of Political Economy Macroeconomics* 1(2), 403–446.
- Cole, S. J. (2021). Learning and the effectiveness of central bank forward guidance. *Journal of Money, Credit and Banking* 53(1), 157–199.
- Cottle Hunt, E. (2021). Adaptive learning, social security reform, and policy uncertainty. *Journal of Money, Credit and Banking* 53(4), 677–714.
- Dombeck, B. (2022). On the expectational stability of rational expectations equilibria in news-shock dsge models. *Macroeconomic Dynamics* 26, 102–126.
- Evans, G. W. and S. Honkapohja (2001). *Learning and Expectations in Macroeconomics*. Princeton University Press.
- Evans, G. W., S. Honkapohja, and K. Mitra (2009). Anticipated fiscal policy and adaptive learning. *Journal of Monetary Economics* 56, 930–953.
- Gasteiger, E. and S. Zhang (2014). Optimal fiscal policy under learning. *Journal of Economic Dynamics and Control* 39, 113–126.
- Gertler, M., N. Kiyotaki, and A. Queralto (2012). Financial crises, bank risk exposure and government financial policy. *Journal of Monetary Economics* 59(S), 17–34.
- Giannitsarou, C. (2006). Supply-side reforms and learning dynamics. *Journal of Monetary Economics* 53, 291–309.
- Hoffmann, M., E. Moench, L. Pavlova, and G. Schulte Frankenfeld (2022). Would households understand average inflation targeting? *Journal of Monetary Economics* 129, S52–S66.
- Honkapohja, S. and N. McClung (2026). On robustness of average inflation targeting. *European Economic Review* 183, 105217.
- Honkapohja, S. and K. Mitra (2020). Price level targeting with evolving credibility. *Journal of Monetary Economics* 116(C), 88–103.

- Iacoviello, M. (2005). House prices, borrowing constraints, and monetary policy in the business cycle. *American Economic Review* 95(3), 739–764.
- Jaimovich, N. and S. Rebelo (2009). Can news about the future drive the business cycle? *American Economic Review* 99, 1097–1118.
- Marcet, A. and T. J. Sargent (1989). Convergence of least squares learning mechanisms in self-referential linear stochastic models. *Journal of Economic Theory* 48(2), 337–368.
- Mian, A. and A. Sufi (2011). House prices, home equity-based borrowing, and the US household leverage crisis. *American Economic Review* 101(5), 2132–2156.
- Millard, S., M. Rubio, and A. Varadi (2024). The macroprudential toolkit: effectiveness and interactions. *Oxford Bulletin of Economics and Statistics* 86(2), 335–384.
- Mitra, K., G. W. Evans, and S. Honkapohja (2013). Policy change and learning in the RBC model. *Journal of Economic Dynamics and Control* 37, 1947–1971.
- Pintus, P. and J. Suda (2019). Learning financial shocks and the Great Recession. *Review of Economic Dynamics* 31, 123–146.
- Pintus, P., J. Suda, and B. Turgut (2021). Estimating financial frictions under adaptive learning. Technical report, GRAPE.
- Shukayev, M. and R. Amano (2025). Price-level targeting and inflation expectations: experimental evidence. *International Journal of Central Banking* 21(3), 185–227.
- Suda, J. (2018). Belief-twisting shocks and the macroeconomy. *Macroeconomic Dynamics* 22(7), 1844–1858.

Appendices

A Log-Linearized Model

We now derive the log-linearized version of the set of equations (2)–(10) which define the intertemporal equilibrium near steady state.

In all equations below, lowercase letters denote logs and $\tilde{x}_t$ denotes the log of X_t/X , where X is the steady-state value of X_t . For example, $\tilde{k}_t \equiv k_t - k$, with $k_t = \log(K_t)$ and $k = \log(K)$, so that lowercase variables without time subscript are steady-state levels in log.

Eliminating N_t by using log-linearized (7) and Φ_t by using log-linearized (10), we get the following set of linearized equations:

$$\tilde{b}_t = (1 + \varepsilon)E_{t-1}[\tilde{q}_t] + \tilde{\theta}_{t-1} \quad (20)$$

$$\frac{K}{Y}\tilde{k}_t - \frac{B}{Y}\tilde{b}_t = -\frac{C}{Y}\tilde{c}_{t-1} - \frac{(1+R)B}{Y}\tilde{b}_{t-1} + \left(\alpha + \alpha\frac{1-\alpha-\gamma}{\chi+\alpha+\gamma} + (1-\delta)\frac{K}{Y}\right)\tilde{k}_{t-1} \quad (21)$$

$$\frac{C/Y}{\Lambda/Y}\tilde{c}_t + \frac{1}{\sigma}\tilde{\lambda}_t = \frac{1-\alpha-\gamma}{\chi+\alpha+\gamma}\frac{\alpha}{\Lambda/Y}\tilde{k}_t \quad (22)$$

$$\begin{aligned} \tilde{q}_t + \tilde{\lambda}_t(1 - \mu\bar{\Theta}) &= E_t[\tilde{q}_{t+1}] (\beta + \bar{\Theta}(1 + \varepsilon)(\mu - \beta)) + \tilde{\theta}_t\bar{\Theta}(\mu - \beta) \\ &\quad + E_t[\tilde{\lambda}_{t+1}] \left(\beta(1 - \bar{\Theta}) + \gamma\beta\frac{Y}{Q} \right) + \alpha\gamma\beta\frac{Y}{Q} \left(1 + \frac{1-\alpha-\gamma}{\chi+\alpha+\gamma} \right) E_t[\tilde{k}_{t+1}] \end{aligned} \quad (23)$$

$$\tilde{\lambda}_t = E_t[\tilde{\lambda}_{t+1}] (\beta(1 - \delta) + \alpha\beta\frac{Y}{K}) + \alpha\beta\frac{Y}{K}(\alpha - 1 + \alpha\frac{1-\alpha-\gamma}{\chi+\alpha+\gamma})E_t[\tilde{k}_{t+1}] \quad (24)$$

$$\tilde{\theta}_t = \rho_\theta\tilde{\theta}_{t-1} + \xi_t. \quad (25)$$

Note that (20) and (21) are the linearized around the steady state, lagged versions of collateral (3) and budget (2) constraints, respectively.

Define $P'_t \equiv (b_t, k_t, \theta_t)$ and $S'_t \equiv (c_t, q_t, \lambda_t)$ the vectors of predetermined and jump variables in log, respectively. Then equations (21)–(25) can be decomposed into two subsystems, each pertaining to P_t and S_t . The first block composed of (21), (20), and (25) can be written:

$$M_0P_t = M_1S_{t-1} + M_2E_{t-1}[S_t] + M_3P_{t-1} + V\xi_t, \quad (26)$$

where:

$$M_0 = \begin{pmatrix} 1 & 0 & 0 \\ -\frac{B}{Y} & \frac{K}{Y} & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad M_1 = \begin{pmatrix} 0 & 0 & 0 \\ -\frac{C}{Y} & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad M_2 = \begin{pmatrix} 0 & 1 + \varepsilon & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix},$$

$$M_3 = \begin{pmatrix} 0 & 0 & 1 \\ -(1+R)\frac{B}{Y} & \alpha + \alpha\frac{1-\alpha-\gamma}{\chi+\alpha+\gamma} + (1-\delta)\frac{K}{Y} & 0 \\ 0 & 0 & \rho_\theta \end{pmatrix}$$

and $V' = (0, 0, 1)$. Note that matrix M_2 captures the direct effect of ε in the leverage constraint (20). The second block (22)-(24) can be written:

$$M_4 S_t = M_5 E_t[S_{t+1}] + M_6 P_t + M_7 E_t[P_{t+1}], \quad (27)$$

where:

$$M_4 = \begin{pmatrix} 0 & 1 & 1 - \mu\bar{\Theta} \\ 0 & 0 & 1 \\ \frac{C/Y}{\Lambda/Y} & 0 & 1/\sigma \end{pmatrix}, \quad M_5 = \begin{pmatrix} 0 & \beta + \bar{\Theta}(1+\varepsilon)(\mu - \beta) & \beta(1 - \bar{\Theta}) + \gamma\beta\frac{Y}{Q} \\ 0 & 0 & \beta(1 - \delta) + \alpha\beta\frac{Y}{K} \\ 0 & 0 & 0 \end{pmatrix},$$

$$M_6 = \begin{pmatrix} 0 & 0 & \bar{\Theta}(\mu - \beta) \\ 0 & 0 & 0 \\ 0 & \frac{\alpha(1-\alpha-\gamma)}{(\chi+\alpha+\gamma)\Lambda/Y} & 0 \end{pmatrix}, \quad M_7 = \begin{pmatrix} 0 & \alpha\gamma\beta\frac{Y}{Q}(1 + \alpha\frac{1-\alpha-\gamma}{\chi+\alpha+\gamma}) & 0 \\ 0 & \alpha\beta\frac{Y}{K}(\alpha - 1 + \alpha\frac{1-\alpha-\gamma}{\chi+\alpha+\gamma}) & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

Finally, substituting the expression of P_t from (26) in (27) and combining two blocks of equations allows one to rewrite the system as:

$$X_t = \mathbf{A}X_{t-1} + \mathbf{B}E_{t-1}[X_t] + \mathbf{C}E_t[X_{t+1}] + \mathbf{D}\xi_t, \quad (28)$$

where $X'_t = \text{vec}(S'_t, P'_t)$ and

$$\begin{aligned} \mathbf{A} &= \begin{pmatrix} M_4^{-1}M_6M_0^{-1}M_1 & M_4^{-1}M_6M_0^{-1}M_3 \\ M_0^{-1}M_1 & M_0^{-1}M_3 \end{pmatrix}, \\ \mathbf{B} &= \begin{pmatrix} M_4^{-1}M_6M_0^{-1}M_2 & 0_3 \\ M_0^{-1}M_2 & 0_3 \end{pmatrix}, \\ \mathbf{C} &= \begin{pmatrix} M_4^{-1}M_5 & M_4^{-1}M_7 \\ 0_3 & 0_3 \end{pmatrix}, \\ \mathbf{D} &= \begin{pmatrix} M_4^{-1}M_6M_0^{-1}V_1 \\ M_0^{-1}V_1 \end{pmatrix}, \end{aligned}$$

where 0_3 is a 3-by-3 zero matrix.

Note that matrix $\mathbf{B}$, via M_2 , captures the direct effect of ε on X_t , while matrix $\mathbf{C}$ captures the effect ε associated with Lagrange multiplier associated with the leverage constraint (20).

B Additional dynamics

Figure 7: Responses under learning (solid) and RE (dotted) incorrect beliefs regarding the macroprudential policy with stronger confidence







Figure 8: Responses under learning (solid red) and RE (dotted blue) incorrect beliefs regarding the macroprudential policy with $\varepsilon = -1$ (stochastic debt limit)




(a) Output following 5% leverage shock.




(b) Capital following 5% leverage shock.



(c) Consumption following 5% leverage shock.



(d) Land price following 5% leverage shock.



(e) Debt following 5% leverage shock.

C “Partially” understanding the economy

In this section we consider the case, where agents to some extent “understand” the economy and the impact of the change of macroprudential policy. To model it we consider equation (17), in which matrices $\mathbf{B}$ and $\mathbf{C}$ capture the impact of expectations on the dynamics of vector $\mathbf{X}_t$. In section A we show that matrix $\mathbf{B}$ and $\mathbf{C}$ are related to the linearized borrowing constraint in equation (4) and Lagrange multiplier on that constraint, respectively, capturing the direct impact of macroprudential policy (measured by ε).

We analyze two case in which agents in the model are still assumed to employ adaptive learning but they partially understand how the change of ε changes the structure of the economy. First, in the AL^B , we consider the case in which upon changing ε (in period 10) the $M_{policy\ change}^B$ corresponds to the rational expectations solution of the model in which matrix $\mathbf{B}$ uses the ε_1 but matrix $\mathbf{C}$ uses ε_0 . Next, in the AL^C case, we consider the reverse situation in which $M_{policy\ change}^C$ corresponds to the rational expectations solution of the model in which matrix $\mathbf{C}$ uses the ε_1 but matrix $\mathbf{B}$ uses ε_0 .

Figures 9 and 10 illustrate the impulse response functions of output, consumption, land prices, and debt following a shift in a macroprudential leverage policy—from procyclical ($\varepsilon_0 = 0.5$) to countercyclical ($\varepsilon_1 = -0.75$)—under three different settings: rational expectations, adaptive learning, and adaptive learning with partially correct beliefs. The partially correct beliefs assume that agents have knowledge of either matrix $\mathbf{B}$ or matrix $\mathbf{C}$.


Figure 9 shows impulse response functions under $\tilde{\rho}_\theta = \rho_\theta$, while Figure 10 shows impulse response functions for $\tilde{\rho}_\theta = 0.9904 > \rho_\theta = 0.9756$.

Consistent with the patterns observed in Figure 2, under pure adaptive learning, the financial shock results in a larger decline in both output and consumption compared to RE. However, once the macroprudential policy becomes countercyclical in period 10, the dynamics under adaptive learning with partial knowledge begin to align more closely with those under RE. Notably, this adjustment reduces the overshooting effects observed under pure adaptive learning.


Furthermore, the responses of output and consumption under AL with $\mathbf{B}$ known are closer to RE than under AL with $\mathbf{C}$ known. This suggests that understanding the direct impact of macroprudential policy on borrowing constraints (captured by matrix $\mathbf{B}$) plays a more significant role in stabilizing these variables than alternative specifications.

For land prices and borrowing, all scenarios exhibit a sharp initial decline. However, the policy shift produces distinct dynamics across the models. When agents possess partial knowledge of $\mathbf{C}$, both land prices and debt exhibit a sharp reversal, returning to levels consistent with RE. In contrast, the dynamics under adaptive learning with $\mathbf{B}$ known more closely resemble those under pure adaptive learning, indicating slower convergence to RE-like behavior in these variables.


Figure 9: Impulse response functions of endogenous variables following the change from procyclical ($\varepsilon = 0.5$) leverage to countercyclical ($\varepsilon = -0.75$) leverage under learning (solid red) with partially correct beliefs regarding the macroprudential policy.




(a) Response of output.



(b) Response of consumption.




(c) Response of land price.




(d) Response of debt.


Figure 10: Impulse response functions of endogenous variables following the change from procyclical ($\varepsilon = 0.5$) leverage to countercyclical ($\varepsilon = -0.75$) leverage under learning (solid red) with partially correct beliefs regarding the macroprudential policy but incorrect beliefs about the persistence of leverage shock.




(a) Response of output.



(b) Response of consumption.



(c) Response of land price.



(d) Response of debt.

To quantitatively evaluate the impact of partial knowledge about the effect of changes in ε on volatility, we again rely on simulations. Table 4 reports the median standard deviation (volatility) of key variables under different informational assumptions. Consistent with the results in Table 3, volatility under adaptive learning is higher than under rational expectations. However, effective communication—where agents partially “understand” the impact of the policy change on the economy—leads to significantly lower volatility. Notably, when agents understand the effect of ε on the economy through both matrices $\mathbf{B}$ and $\mathbf{C}$ (denoted as AL^{BC}), the overall volatility becomes very close to the levels observed under RE, as shown in column (5).

Table 4: Effects of policy change on volatility under partially correct beliefs (median standard deviation, %)

	RE (1)	AL_1 (2)	AL_1^B (3)	AL_1^C (4)	AL_1^{BC} (5)	AL_2 (6)	AL_2^B (7)	AL_2^C (8)	AL_2^{BC} (9)
consumption									
<i>before change</i>	2.13	2.15	2.15	2.15	2.15	6.12	6.12	6.12	6.12
<i>periods 1-16</i>	0.89	1.44	1.05	1.13	1.04	6.56	5.10	5.69	5.61
<i>periods 17-56</i>	0.76	2.11	0.90	1.06	0.77	8.55	2.33	2.72	2.04
<i>periods 261-300</i>	0.68	1.65	0.82	0.88	0.69	4.71	1.37	1.11	0.74
output									
<i>before change</i>	2.04	1.98	1.98	1.98	1.98	5.28	5.28	5.28	5.28
<i>periods 1-16</i>	0.73	1.21	0.83	0.84	0.76	6.88	3.75	4.50	4.25
<i>periods 17-56</i>	0.69	1.72	0.79	0.77	0.66	7.64	1.82	2.13	1.53
<i>periods 261-300</i>	0.65	1.21	0.73	0.63	0.61	3.94	1.10	0.80	0.61
borrowing									
<i>before change</i>	35.69	38.07	38.07	38.07	38.07	119.32	119.32	119.32	119.32
<i>periods 1-16</i>	11.96	14.98	14.78	12.04	12.02	31.16	30.26	15.89	15.86
<i>periods 17-56</i>	12.51	15.43	15.37	12.53	12.53	26.71	26.18	13.96	13.96
<i>periods 261-300</i>	12.23	14.39	14.53	12.29	12.23	21.62	21.96	12.68	12.67
land price									
<i>before change</i>	17.35	20.67	20.67	20.67	20.67	76.31	76.31	76.31	76.31
<i>periods 1-16</i>	5.77	17.83	16.88	6.57	6.51	67.56	63.83	9.52	9.45
<i>periods 17-56</i>	6.01	17.83	17.36	6.81	6.81	57.24	55.46	8.89	8.85
<i>periods 261-300</i>	5.93	15.55	15.55	6.96	6.93	43.12	44.08	8.30	8.26

The standard deviation represents the variability of a macroeconomic variable driven solely by leverage shocks. Matrices describing the VAR(1) dynamics of series are initialized at the corresponding (to the parameter values) RE equilibria. Series are simulated over 400 periods with 60 first observations being discarded. The median is computed across 1000 simulations. AL_1 refers to a constant gain learning, where the perceived persistence of the leverage shock at the beginning of simulations (period 0) matches the true value (i.e., the rational expectations value ρ_θ). AL_2 refers to constant gain learning, where the estimated persistence of the leverage shock ($\hat{\rho}_\theta$) at the beginning simulations (period 0) exceeds the rational expectations value. AL^B and AL^C denote the adaptive learning with known B and C , respectively. Agents engage in the learning process of relevant matrices before the policy change occurs.

Online Appendix

(Available at <https://jaceksuda.com/research/>)

OA.1 Additional results

Table OA.1 illustrates the evolution of the difference in simulated standard deviations between the AL_1 and AL_2 specifications as the post-policy change simulation horizon is extended from 300 to 700 periods.

As the simulation length increases from 300 to 700 periods, the influence of the initial overreaction to change of ε and the slow updating of beliefs—governed by the gain parameter—gradually diminishes. The data in Table OA.1 clearly show a consistent narrowing of the gap in standard deviations across the extended sample, confirming that the two specifications converge as agents accumulate more information. However, the actual pace of this convergence is sensitive to the specific nature of the policy change (governed by both ε_0 and ε_1) and the macroeconomic variables under consideration. There is a notable difference in both the overall reduction of volatility—for example, the median standard deviations of output and consumption decrease over time more significantly than those of land prices and borrowing—and the specific convergence of AL_2 to AL_1 . In the latter case, the volatility of borrowing and land prices appears closer across specifications than that of output and consumption. This suggests that variables more heavily influenced by the persistence of shocks, or those exhibiting higher degrees of “learning inertia”, tend to retain differences for a longer duration. This heterogeneity highlights that while the models are asymptotically equivalent, the medium-run persistence of learning effects is strictly conditional upon the underlying economic environment and the specific variable being analyzed.

Table OA.1: Longer sample effects of policy change on volatility (median standard deviation, %)

	before change			periods 1–40			periods 261–300			periods 661–700		
	<i>RE</i>	<i>AL</i> ₁	<i>AL</i> ₂	<i>RE</i>	<i>AL</i> ₁	<i>AL</i> ₂	<i>RE</i>	<i>AL</i> ₁	<i>AL</i> ₂	<i>RE</i>	<i>AL</i> ₁	<i>AL</i> ₂
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
$\varepsilon_0 = 0.5 \rightarrow \varepsilon_1 = 0$												
<i>consumption</i>	2.08	2.08	5.98	1.12	1.71	5.70	1.14	1.71	4.68	1.17	1.60	3.70
<i>output</i>	1.99	1.92	5.18	1.09	1.30	4.25	1.09	1.32	3.05	1.12	1.28	2.75
<i>borrowing</i>	37.23	39.50	127.77	20.09	28.69	80.67	20.40	28.12	61.59	20.87	28.19	55.51
<i>land price</i>	18.06	21.11	81.19	9.70	18.99	67.69	9.88	18.25	51.95	10.10	18.08	45.06
$\varepsilon_0 = 0.5 \rightarrow \varepsilon_1 = -0.5$												
<i>consumption</i>	2.08	2.08	5.98	0.95	1.77	7.03	0.78	1.63	4.65	0.93	1.45	2.58
<i>output</i>	1.99	1.92	5.18	0.81	1.30	6.21	0.75	1.15	3.51	0.77	1.03	1.88
<i>borrowing</i>	37.23	39.50	127.77	13.95	19.62	46.10	14.13	18.59	33.69	13.82	19.25	27.64
<i>land price</i>	18.06	21.11	81.19	6.71	18.07	65.11	6.85	16.43	46.13	6.67	18.16	34.31
$\varepsilon_0 = 0.5 \rightarrow \varepsilon_1 = -0.75$												
<i>consumption</i>	2.08	2.08	5.98	0.88	1.88	7.90	0.69	1.65	4.71	0.70	1.12	2.06
<i>output</i>	1.99	1.92	5.18	0.75	1.46	7.47	0.65	1.21	3.94	0.67	0.80	1.53
<i>borrowing</i>	37.23	39.50	127.77	12.10	15.08	29.24	12.23	14.39	21.62	12.46	14.19	17.89
<i>land price</i>	18.06	21.11	81.19	5.81	17.76	63.54	5.93	15.55	43.12	6.05	13.25	26.56
$\varepsilon_0 = 0.5 \rightarrow \varepsilon_1 = -1.5$												
<i>consumption</i>	2.08	2.08	5.98	0.79	2.45	10.87	0.48	1.83	5.27	0.50	0.81	1.84
<i>output</i>	1.99	1.92	5.18	0.65	2.43	11.80	0.46	1.68	5.30	0.48	0.66	1.87
<i>borrowing</i>	37.23	39.50	127.77	8.65	2.31	18.73	8.70	4.21	8.43	8.90	6.71	7.46
<i>land price</i>	18.06	21.11	81.19	4.13	16.86	58.95	4.23	13.22	35.75	4.33	9.17	13.46
$\varepsilon_0 = 0 \rightarrow \varepsilon_1 = -0.75$												
<i>consumption</i>	1.16	1.16	1.67	0.70	0.99	2.01	0.69	0.97	1.39	0.70	0.82	1.01
<i>output</i>	1.11	1.05	1.40	0.65	0.73	1.88	0.65	0.70	0.99	0.67	0.62	0.73
<i>borrowing</i>	20.81	21.69	35.54	12.11	13.15	17.57	12.23	13.04	14.28	12.46	13.20	14.12
<i>land price</i>	10.16	11.85	23.25	5.87	10.54	19.48	5.93	9.84	14.61	6.05	9.41	12.71

Notes: The standard deviation represents the variability of a macroeconomic variable over 40 periods (equivalent to 10 years) driven solely by leverage shocks. Matrices describing the VAR(1) dynamics of series are initialized at the corresponding (to the parameter values) RE equilibria. Series are simulated over 800 periods with 60 first observations being discarded. The median is computed across 1000 simulations. *AL*₁ refers to a constant gain learning, where the perceived persistence of the leverage shock at the beginning of simulations (period 0) matches the true value (i.e., the rational expectations value $\rho_\theta = 0.9756$). *AL*₂ refers to constant gain learning, where the estimated persistence of the leverage shock ($\hat{\rho}_\theta = 0.9904$) at the beginning simulations (period 0) exceeds the rational expectations value. Agents engage in the learning process before the policy change occurs.

Table OA.2 explores the sensitivity of learning dynamics to two critical factors: the persistence (and the perceived persistence) of leverage shocks, ρ_θ and $\hat{\rho}_\theta$, and the value of the gain parameter, ν . The results demonstrate that the discrepancy between Rational Expectations (RE) and *AL*₁ is significantly attenuated when shock persistence is lower. Specifically, as shown in the upper panel of the table—which considers lower baseline values for ρ_θ and $\hat{\rho}_\theta$ —the initial overreaction of agents is reduced. This leads to standard deviations that are much closer to the RE benchmark across all variables and horizons, a result that holds regardless of the specific value of the gain parameter ν .

Furthermore, the results highlight that the convergence of *AL*₂ to *AL*₁ is much more pronounced under conditions of lower persistence. When shocks are less persistent, “learning

inertia” is diminished, allowing the two learning specifications to align more rapidly. This suggests that the persistence of exogenous shocks acts as a multiplier for the effects of the gain parameter; even with a low gain, the model returns to a state of alignment more quickly if the underlying shocks do not continuously displace the system from its steady state. Ultimately, this table confirms that the persistent gaps discussed in the main text are not a general feature of learning, but are specifically intensified by the high-persistence environment required to match empirical macroeconomic data.

The minor discrepancies in volatilities between Table OA.1 and Table OA.2 arise from the implementation of the projection facility. While the model is stable under learning—meaning small deviations from Rational Expectations (RE) do not cause the model to explode—certain large shock realizations may occasionally push the learning algorithm outside the convergence space. When such a draw was encountered, it was excluded from the sample and replaced with a new realization. This resampling process results in slight variations in the actual sample used to compute the median standard deviations across different simulation lengths.

Table OA.2: Effects of gain and persistence on volatility following policy change (median standard deviation, %)

$\rho_\theta = 0.9, \hat{\rho}_\theta = 0.95$	$\nu = 0.008$			$\nu = 0.014$		$\nu = 0.024$	
	<i>RE</i> (1)	<i>AL</i> ₁ (2)	<i>AL</i> ₂ (3)	<i>AL</i> ₁ (4)	<i>AL</i> ₂ (5)	<i>AL</i> ₁ (6)	<i>AL</i> ₂ (7)
consumption							
<i>before change</i>	0.61	0.60	1.05	0.60	1.06	0.61	1.07
<i>periods 1-40</i>	0.43	0.47	0.72	0.48	0.74	0.49	0.76
<i>periods 261-300</i>	0.43	0.46	0.43	0.46	0.60	0.46	0.56
<i>periods 461-500</i>	0.43	0.45	0.58	0.45	0.53	0.48	0.53
output							
<i>before change</i>	0.70	0.67	1.05	0.67	1.06	0.68	1.07
<i>periods 1-40</i>	0.50	0.45	0.52	0.46	0.55	0.47	0.59
<i>periods 261-300</i>	0.50	0.46	0.50	0.47	0.50	0.48	0.51
<i>periods 461-500</i>	0.50	0.47	0.47	0.47	0.48	0.51	0.51
borrowing							
<i>before change</i>	9.46	9.69	17.72	9.75	17.87	9.88	18.17
<i>periods 1-40</i>	6.78	6.95	7.97	6.95	9.97	6.91	7.95
<i>periods 261-300</i>	6.72	6.85	7.37	6.81	7.23	6.82	7.12
<i>periods 461-500</i>	6.78	6.88	7.31	6.83	7.11	7.15	7.15
land price							
<i>before change</i>	2.37	3.10	7.68	3.17	7.73	3.23	7.95
<i>periods 1-40</i>	1.70	2.72	5.31	2.77	5.39	2.86	5.59
<i>periods 261-300</i>	1.69	2.63	4.65	2.56	4.25	2.59	3.89
<i>periods 461-500</i>	1.70	2.54	4.24	2.46	3.66	2.74	3.51
$\rho_\theta = 0.9756, \hat{\rho}_\theta = 0.9904$	$\nu = 0.008$			$\nu = 0.014$		$\nu = 0.024$	
	<i>RE</i>	<i>AL</i> ₁	<i>AL</i> ₂	<i>AL</i> ₁	<i>AL</i> ₂	<i>AL</i> ₁	<i>AL</i> ₂
consumption							
<i>before change</i>	2.03	2.00	5.73	2.02	5.80	2.06	5.97
<i>periods 1-40</i>	0.87	1.91	8.10	1.94	8.23	2.00	8.11
<i>periods 261-300</i>	0.68	1.84	5.63	1.66	4.86	1.37	3.44
<i>periods 461-500</i>	0.71	1.63	4.94	1.36	3.26	1.14	2.08
output							
<i>before change</i>	1.97	1.85	4.92	1.86	5.04	1.91	5.18
<i>periods 1-40</i>	0.76	1.51	7.68	1.55	7.78	1.60	7.66
<i>periods 261-300</i>	0.65	1.41	4.83	1.23	4.05	0.98	2.69
<i>periods 461-500</i>	0.67	1.22	4.18	0.95	2.46	0.82	1.57
borrowing							
<i>before change</i>	36.81	37.74	127.44	38.53	129.80	39.70	135.96
<i>periods 1-40</i>	13.03	16.06	28.98	16.13	30.25	16.21	30.53
<i>periods 261-300</i>	12.52	15.00	23.34	14.94	22.44	14.58	20.64
<i>periods 461-500</i>	12.44	14.79	22.38	14.40	19.82	14.24	18.37
land price							
<i>before change</i>	17.74	20.40	81.25	20.75	82.69	121.42	84.76
<i>periods 1-40</i>	6.26	18.57	64.39	18.78	65.23	19.49	66.41
<i>periods 261-300</i>	6.07	16.56	47.83	15.82	44.34	14.97	37.66
<i>periods 461-500</i>	6.04	15.66	44.03	14.22	35.58	13.77	28.17

Notes: Notes: The standard deviation represents the variability of a macroeconomic variable over 40 periods (equivalent to 10 years) driven solely by leverage shocks. Matrices describing the VAR(1) dynamics of series are initialized at the corresponding (to the parameter values) RE equilibria. Series are simulated over 600 periods with 60 first observations being discarded. The median is computed across 1000 simulations. *AL*₁ refers to a constant gain learning, where the perceived persistence of the leverage shock at the beginning of simulations (period 0) matches the true value (i.e., the rational expectations value $\rho_\theta = 0.90$ in upper part and $\rho_\theta = 0.9756$ in lower part of the table). *AL*₂ refers to constant gain learning, where the estimated persistence of the leverage shock ($\hat{\rho}_\theta = 0.95$ in upper part and $\hat{\rho}_\theta = 0.9904$ in the lower part of the table) at the beginning simulations (period 0) exceeds the rational expectations value. Agents engage in the learning process before the policy change occurs.