

Unstable Inflation Targets*

William A. Branch

University of California, Irvine

George W. Evans

University of Oregon

University of St. Andrews

December 10, 2015

Abstract

This paper studies long-run inflation targets and stability in an imperfect information environment. When central banks set an inflation target that is not fully communicated, agents draw inferences about inflation from recent data and remain alert to structural change in their econometric model by forming expectations from a forecasting model that is estimated via discounted least squares. Inflation targets can lead agents' beliefs to depart from rational expectations through two channels. First, implementing a higher inflation target can lead to overshooting of the target. Second, there can be nearly self-fulfilling inflation, disinflation, or deflation that arises as an endogenous response to shocks. Policy implications for implementing a higher target without de-anchoring expectations are discussed.

JEL Classifications: E52; E32; D83; D84

Keywords: Inflation target, monetary policy, expectations, adaptive learning

*Financial support from National Science Foundation Grant no. SES-1025011 is gratefully acknowledged.

Policy makers have generally chosen a 2 % (inflation rate target). But there was no very good reason to use 2 % rather than 4 %. Two percent doesn't mean price stability. Between 2 % and 4 %, there isn't much cost from inflation. If I were to choose an inflation target today, I'd strongly argue for 4%.

– Interview with Olivier Blanchard in WSJ 2/11/2010

In this context, raising the inflation objective would likely entail much greater costs than benefits. Inflation would be higher and probably more volatile under such a policy, undermining confidence and the ability of firms and households to make longer-term plans, while squandering the Fed's hard-won inflation credibility. Inflation expectations would also likely become significantly less stable, and risk premiums in asset markets—including inflation risk premiums—would rise.

–Chairman Ben S. Bernanke, remarks at the 2010 Jackson Hole Symposium

Since the economic crisis erupted, statistical properties of UK core CPI have swung markedly. From 2008 onwards, UK core inflation cannot be modelled as a stationary process. This shift may indicate that the MPC has become more tolerant of inflation deviating from target as a pragmatic reaction to the recession. However, it would also be consistent with the tolerance of persistently higher inflation – effectively raising the inflation target. We believe the former is the more accurate description of reality, but the danger is that wage and price setters take the latter view, and higher inflation expectations become self-fulfilling.

– Barclays Capital, Global Economics Weekly 10/15/2010

1 Introduction

The recent global economic crisis has renewed debate over the appropriate long-run inflation rate that monetary policymakers should target. Some have argued (see Krugman (1998), Summers (1991), Ball (2013)) that in “normal” times monetary authorities should pursue a higher average inflation rate in order to provide a cushion that facilitates a lowering of nominal interest rates so as to avoid a liquidity trap due to a binding zero-lower-bound constraint on interest rates. Subsequent analyses, that are typically based in models with explicit micro-foundations, find little support for higher average inflation rates. For example,

Woodford (2003), Schmitt-Grohe and Uribe (2011), Eggertsson and Woodford (2003) and Coibion, Gorodnichenko, and Wieland (2012) find that the optimal inflation rate in New Keynesian models are close to zero, even after accounting for the protection that higher inflation rates provide in avoiding the zero lower bound. Ascari and Ropele (2007) find that higher target rates of inflation can destabilize economies. (See, also, Ascari and Sbordone (2014) for a recent review.) Nevertheless, among many economists there is the view that the added stability that might be achieved from a higher inflation target outweighs any of the distortionary losses associated with inflation.

Policymakers, though, are reticent to depart from a commitment to low and stable rates of inflation. Chairman Bernanke, in the quote above, expresses the view that higher average inflation rates will lead to more volatile inflation and “inflation expectations would also likely become significantly less stable.” Most research into optimal long-run inflation rates and optimal policy in the presence of a zero lower bound constraint do not have a channel through which higher average inflation leads to the de-anchoring of inflation expectations. Conventional models assume perfect information (and rational expectations) which have the feature that, away from the zero lower bound, inflation will be a stationary process around the long-run target.¹

This paper re-examines the stability of long-run inflation targets in an environment with imperfect information and adaptive learning. As a benchmark, the paper takes as its economic framework a simple dynamic Fisherian model in which nominal interest rates are adjusted, in accordance with a Taylor rule, whenever inflation deviates from its long-run target rate. Private-sector agents in the economy know the form of the Taylor rule. However, they do not know – or harbor some doubt about – the precise values for the response coefficient or the value and/or timing of the long-run inflation target. Instead, agents draw inferences about the inflation process from recent data by adopting an econometric forecasting model whose reduced-form nests the full class of rational expectations equilibria. These agents are Bayesian and, because of their uncertainty about the inflation target, they place a prior on structural change in their econometric model. This imperfect information framework implies that private sector agents adopt a simple AR(1) forecasting model, the parameters of which are updated in real time with a form of discounted least squares (“constant gain learning”). The priors of this model are specified in such a way that beliefs are, on average, close to rational expectations.

¹An important exception is Williams (2006) who examines whether higher inflation targets can make liquidity traps occur less often when rational expectations are replaced with a constant gain, or perpetual, learning rule.

The primary results of this paper are as follows. First, although over time beliefs tend to converge toward rational expectations, the combination of constant gain learning and a positive inflation target can lead agents in the economy to temporarily believe that inflation follows a random walk without drift. Under these beliefs, agents will interpret recent innovations to inflation as permanent shifts in the mean inflation rate. Random walk beliefs arise for a very intuitive reason. The long-run inflation target, and imperfect information about that target, lead agents to estimate the mean inflation rate from historical data. As a thought experiment, suppose there is a slight (temporary) upward drift in the inflation rate. Agents' econometric models will pick up that drift, leading to higher inflation expectations that feed back into higher inflation rates. This process is self-reinforcing and in some cases can lead agents to believe that inflation follows a random walk.

Second, random-walk beliefs, as we will show below, are nearly self-fulfilling, and consequently such beliefs tend to persist for a substantial period of time. Furthermore these beliefs tend to generate considerable economic volatility, characterized by significant bursts of inflation, disinflation, and even deflation. We find these dynamics arise across a broad range of model parameters, policy rules, and trend inflation rates. Third, implementing a higher target – say by moving the target from 2% to 4% – will introduce just the type of drift in inflation that can lead to random-walk beliefs and cause a substantial overshooting of the inflation target.

An important feature of our analysis is that imperfect information of inflation targets can generate instability in inflation rates even though the departure, on average, from rational expectations is small. The framework employed here is related to an extensive literature that employs adaptive learning in macroeconomics. Most closely related are papers that incorporate constant gain learning in studies of monetary policy and asset pricing (see, for example, Branch and Evans (2011); Sargent (1999); Orphanides and Williams (2005a); Cho, Williams, and Sargent (2002); Williams (2004); Cho and Kasa (2008); Eusepi and Preston (2011)). Branch and Evans (2011), in particular, find that when risk-averse agents in an asset pricing model forecast both the risk and return of stock prices using a forecasting model whose parameters are updated with constant gain least squares then traders may also come to temporarily believe that stocks follow a random walk. These nearly self-fulfilling random walk beliefs lead to recurrent bubbles and crashes in stock prices. The intuition for why inflation targets are destabilizing in an adaptive learning environment is similar to the existence of bubbles and crashes in stock markets.

Is it reasonable to assume that the private-sector might have imperfect information, or doubts, about the long-run inflation target? The answer is yes, for a variety of reasons. First,

the Federal Reserve Bank only recently adopted a stated inflation target and, in fact, it faces a dual mandate legislated by congress. In recent FOMC meetings, there has been debate about whether the Summary of Economic Projections (reported once each quarter by the FOMC) conveys the central bank's target. However, there is considerable uncertainty about that target value due to diverse views in the composition of the FOMC. In the Summary of Economic Projections, the central tendency ranges from 1.5-2% (though recently more concentrated at 2%). In the Survey of Professional Forecasters, there is considerable disagreement about average annual inflation over a 10 year period. SPF participants expect a 2.5% long-run inflation rate and the dispersion across forecasters ranges from 0.4% to 0.8% in each survey quarter.² Even in countries with an explicit target, such as England, the Barclays quote above shows that sophisticated market participants might hold doubts about the long-run target.

The results in the current paper provide a caution to proposals for higher long-run inflation targets as a safeguard against hitting the zero lower bound. In an extension of our basic analysis, we show that the likelihood of destabilizing random-walk beliefs arising increases with the long-run inflation target in a New Keynesian model, closed with a standard Taylor rule and subject to imperfect information about the long-run target. Moreover, the benchmark New Keynesian model may exhibit endogenous crashes in the inflation rate that bump up against the zero lower bound in nominal interest rates. These self-fulfilling paths exhibit rapid deflation and large negative output gaps. These liquidity traps only last for a finite period of time as the global stability of the rational expectations equilibrium eventually prevails. Conversely, we show that if policy rules are more aggressive in responding to inflation deviations from target, or if the target is implemented gradually, then destabilizing random-walk beliefs are less likely to arise.³

There are important policy implications from these results. Only credibly and completely informing the private-sector about the long-run inflation target *and* the timing of when that target will be implemented can avoid the unstable dynamics associated with positive inflation targets. When agents know the mean inflation rate, and need only forecast its persistence, random-walk beliefs do not arise. This result is related to a finding by Eusepi and Preston (2010), who show that central bank communication about the policy-setting process can affect the stability of rational expectations equilibria, and Orphanides and Williams (2005b),

²Guraynak, Levin, and Swanson (2010) provide evidence that these long-run inflation expectations drift in response to surprise components of data announcements.

³This latter finding is in line with Cogley, Matthes, and Sbordone (2015) who show that the optimal coefficients in a Taylor rule should be set to higher values when implementing a disinflation in an economy where there is imperfect knowledge about the policy rule.

who emphasize the importance of a credible inflation target.⁴ The analysis below also considers other popular policy proposals, such as price-level targeting and optimal discretionary policy rules, and demonstrates that these policies are not immune to instability. The key message is that environments with imperfect information that cause agents to forecast *both* the mean and persistence of inflation can lead to unstable inflation dynamics.⁵

We also present results from policy experiments that implement a higher target without triggering unstable inflation dynamics. For instance, implementing the higher target gradually, or with sufficient inertia in the policy rule, can prevent the kind of unstable feedback loops that lead to random-walk beliefs. Similarly, we show that combining a higher target with a stronger inflation reaction coefficient in the policy rule can also prevent the onset of unstable dynamics. Thus, another take-away is that a policymaker may want to consider carefully other aspects of the policy rule (e.g. the policy coefficients) when implementing a higher inflation target. We leave, however, the optimal design of the policy rule and target, to future research.

This paper proceeds as follows. Section 2 presents the main results in a simple Fisherian model. Although, the basic result is far more general, and likely to arise in any linear, forward-looking model, the Fisherian model can illustrate the mechanics of the learning process in a transparent manner. Section 3 focuses on the New Keynesian model and demonstrates that the main qualitative results from Section 2 carry over to the New Keynesian environment. This section also discusses the range of policy and model parameters for which our results hold. Section 4 considers a number of extensions including policy communication, optimal discretionary policy and price-level targeting. Finally, Section 5 concludes.

2 A Simple Model of Inflation Targets and Imperfect Information

This paper begins by considering a simple Fisherian model, which can emerge as a special case from richer models that incorporate real and nominal frictions. The Fisherian model illustrates the main points and provides analytic results. As will be seen below, our results can arise in more practical models, such as the New Keynesian model.

⁴Neither of these papers looks at the impact under learning of a change in an imperfectly communicated inflation target.

⁵Also related to this paper is a literature on central bank transparency that has a long and venerable history (see Geraats (2002)). In contrast to this literature, the present paper focuses on econometric learning instead of strategic interactions between the central bank and the public sector.

2.1 Fisherian Model and Imperfect Information

The Fisherian model emerges from a flexible price economy that abstracts from frictions. The Fisher relation arises, e.g. in a constant endowment economy or a model with risk-neutral agents, from the household's first-order condition that prices one period nominal bonds. Monetary policy adjusts nominal interest rates according to a Taylor rule. In this simple (log-linearized) environment, interest rates and inflation are determined by the following two equations

$$\hat{i}_t = \hat{E}_t(\pi_{t+1} - \bar{\pi}) + r_t \quad (1)$$

$$\hat{i}_t = \alpha(\pi_t - \bar{\pi}) \quad (2)$$

where $\alpha > 0$, $\hat{i}_t = i_t - \bar{i}$ is the nominal interest rate in deviation from steady-state form, π_{t+1} is the inflation rate, $\bar{\pi}$ is the central bank's long-run inflation target, and $\bar{i} = \bar{\pi} + \bar{r}$ is the steady-state nominal interest rate and $\bar{r} > 0$ is the steady-state real interest rate.⁶ For simplicity the exogenous shock r_t is assumed to be white noise with variance σ_r^2 . Equation (1) is the Fisherian relation and (2) is the Taylor rule. The operator \hat{E} is the (possibly) non-rational expectations operator, highlighting that imperfect information can affect the economy through the self-referential nature of the asset pricing equation (1).

Combining (1) with (2) leads to an expectational difference equation that determines the path for inflation:

$$\pi_t = \frac{(\alpha - 1)}{\alpha} \bar{\pi} + \alpha^{-1} \hat{E}_t \pi_{t+1} + \alpha^{-1} r_t \quad (3)$$

A rational expectations equilibrium is any (non-explosive) solution to (3). Under rational expectations, $\hat{E} = E$, there are two classes of equilibria that satisfy (3), the minimum state variable (MSV) solution

$$\pi_t = \bar{\pi} + \alpha^{-1} r_t$$

and a non-fundamentals solution

$$\pi_t = (1 - \alpha) \bar{\pi} + \alpha \pi_{t-1} - r_{t-1} + \xi_t$$

where ξ_t is a martingale difference sequence, i.e. $E_{t-1} \xi_t = 0$. Provided that $\alpha > 1$, the non-fundamentals solution is explosive, and the MSV is the unique (non-explosive) rational expectations equilibrium. The condition $\alpha > 1$ is often referred to as the ‘‘Taylor principle’’ as it prescribes nominal interest rates to be adjusted more than one-for-one when inflation

⁶In deriving these equations, households are assumed to satisfy their Euler equations, and these equations provide an adequate (linear) approximation of economic behavior near the steady-state.

deviates from target, and in many models it is a key condition ensuring equilibrium determinacy. Throughout this paper, we focus on policy rules that satisfy the Taylor principle.

Households must have full information about the distribution of the endogenous variables in order to form rational expectations. This includes knowing the details of the policy rule, such as the long-run target $\bar{\pi}$ as well as the reaction coefficient α . An alternative to rational expectations is to assume that agents behave like econometricians who hold a (correctly) specified model of the economy but who must recover the parameters, in real time, from data. An extensive literature studies the conditions under which an economy with adaptive learning will converge to a rational expectations equilibrium (see Marcet and Sargent (1989) and Evans and Honkapohja (2001)). The imperfect information approach adopted in this paper builds on this approach.

When there is imperfect information about the economy, it is reasonable to expect households to form forecasts in the same way as an econometrician. Following this logic, agents are assumed to form expectations based on a perceived forecasting model that in effect nests the full class of rational expectations equilibria:

$$\pi_t = a + b\pi_{t-1} + \varepsilon_t \quad (4)$$

where ε_t is a (perceived) white noise error. Notice that for suitable parameter values (a, b) this perceived model of the economy coincides with either the REE steady state (in which case $b = 0$) or with a serially correlated non-fundamental equilibrium. Under real-time adaptive learning the agents' estimates of the forecast parameters (a, b) evolve slowly over time, e.g. via least-squares updating as discussed below in Section 2.2. Here we focus on the E-stability dynamics that are known to characterize the mean paths under least-squares learning.

Using the forecasting model (4), with given parameters (a, b) , conditional expectations are formed as

$$\hat{E}_t \pi_{t+1} = a(1 + b) + b^2 \pi_{t-1} \quad (5)$$

To calculate these expectations it is assumed that time t -dated variables are not observed contemporaneously.⁷ With expectations in hand, inflation is determined by plugging expectations (5) into the temporary equilibrium equation (3), yielding

$$\pi_t = \alpha^{-1}(\alpha - 1)\bar{\pi} + \alpha^{-1}a(1 + b) + \alpha^{-1}b^2\pi_{t-1} + \alpha^{-1}r_t \quad (6)$$

$$\equiv T(a, b)'X_{t-1} + \alpha^{-1}r_t \quad (7)$$

⁷A frequently-used timing convention in adaptive learning models is that agents do not observe contemporaneous endogenous variables. This timing protocol eliminates the simultaneity of inflation and expected inflation.

where $T(a, b)' = \left(\frac{(\alpha-1)\bar{\pi}}{\alpha} + \alpha^{-1}a(1+b), \alpha^{-1}b^2 \right)$ and $X' = (1, \pi)$. It is straightforward to verify that a fixed point to $T(a, b)$ is a rational expectations equilibrium.

The map T , which takes perceived coefficients (a, b) to the actual coefficients in (6), plays a prominent role in analyses of the expectational stability (“E-stability”) of rational expectations equilibria. The T-map can be interpreted in the following way. If agents held beliefs in the form of the perceived law of motion (4), with parameters (a, b) held constant over time at (possibly) non-RE values, then their forecast rule would be (5). The stochastic process for inflation (6) takes the same form as the perceived process, but with the coefficients $T(a, b)$ replacing (a, b) in (4). Since a rational expectations equilibrium aligns perceptions with outcomes, it is not surprising that a rational expectations equilibrium is a fixed point of the T-map.

Under real-time learning the parameters (a, b) are updated over time, e.g. with least squares, in response to new data. Evans and Honkapohja (2001) has shown, first, that one can easily compute from the T-map a stability condition, E-stability, that governs whether the rational expectations equilibrium is locally stable under learning; and second, that the ordinary differential equation, used to define E-stability, also provides information on the global learning dynamics. More formally, the E-stability Principle states that Lyapunov stable rest points of the E-stability ordinary differential equation

$$\frac{d(a, b)'}{d\tau} = (T(a, b) - (a, b))' \quad (8)$$

are locally stable under least squares learning and other closely related learning algorithms.⁸ Here τ denotes “notional” time, which can, however, be linked to real time t .

That the E-stability condition governs stability of an equilibrium under learning is intuitive, since (8) states that the estimated coefficients (a, b) should be adjusted in the direction of the actual law of motion parameters that generate the data. Local stability of (8) thus addresses whether a small perturbation in the perceived coefficients (a, b) will tend to return to their rational expectations equilibrium values.

In the Fisherian model, it is fairly simple to compute the E-stability of the rational expectations equilibrium. A rational expectations equilibrium will be E-stable provided the roots of $DT(a, b)$, evaluated at their equilibrium values, have real parts less than one. In the current case, it is straightforward to verify that provided the Taylor principle is satisfied, i.e. $\alpha > 1$, the unique rational expectations equilibrium is E-stable. Figure 1 illustrates

⁸Other examples include constant gain (discounted least squares) learning, suitable generalized stochastic gradient learning schemes and related Bayesian updating methods. See Evans and Honkapohja (2001) and Evans, Honkapohja, and Williams (2010) for details.

the intuition by plotting the rest points of the E-stability ODE and the associated vector field. The solid lines indicate the values for which $\dot{a} = 0, \dot{b} = 0$ and the arrows indicate the direction of adjustment in (8). The figure illustrates the two rational expectations equilibria. The $b = \alpha$ equilibrium is explosive and is also unstable under the E-stability dynamics. In contrast, the unique non-explosive rational expectations equilibrium is a sink under learning.

Figure 1 illustrates two further features. First, the fundamentals equilibrium at $\bar{\pi}$ is E-stable and its basin of attraction includes all initial conditions with $b < \alpha$. Second, most analyses of policy under learning focus on the E-stability properties of a particular equilibrium. The figure also suggests that the transitional dynamics might be of independent interest. The vector field indicates that some transitional paths may include non-linear paths to the rational expectations equilibrium.

2.2 Inflation Targets and the Dynamics of Imperfect Information

The E-stability dynamics govern the stability of the rational expectations equilibrium. However, they do not give the full picture of global learning dynamics. This subsection details the learning dynamics and illustrates how long-run inflation targets can alter the qualitative nature of learning dynamics. The central idea is the following: private sector agents are aware of the form of the policy rule but the specifics, such as the size of the reaction coefficient and/or the value and timing of the long-run inflation target, are unknown. An agent in this setting would be wise to remain alert to potential changes in the size and timing of the implementation of the long-run inflation target. Such an agent will then place a prior probability on drifting coefficients in their forecasting model. There are two central ingredients to the results that come below: a positive long-run inflation target that is imperfectly known by agents, and a prior belief of possible structural change.

Let $\theta' = (a, b), X' = (1, \pi)$. Agents are assumed to update their parameter estimates according to the following recursive algorithm

$$\theta_t = \theta_{t-1} + \gamma S_{t-1} X_{t-1} (\pi_t - \theta'_{t-1} X_{t-1})' \quad (9)$$

$$S_t = S_{t-1} + \gamma (X_t X_t' - S_{t-1}) \quad (10)$$

The equations in (9)-(10) are the updating equations for recursive least squares where past data are discounted by a constant “gain” γ . Here S_t is an estimate of $EX_t X_t'$, the second moment matrix of the regressors. Least-squares updating arises when the constant gain γ is replaced by a decreasing sequence $\gamma_t = t^{-1}$.

Sargent and Williams (2005) demonstrate that constant gain least squares can arise from an (approximate) Kalman Filter when agents believe that the process for the drift-

ing coefficients θ_t follows a random walk, a standard assumption in applied econometric work. Specifically, the constant gain learning equations (9)-(10) arise from an approximate Bayesian learning process in which the prior on parameter drift is proportional to the ratio of observation noise variance to the second moment matrix of the regressors, with the speed of drift controlled by the constant gain γ . An alternative interpretation of (9)-(10) is that agents use least squares modified to discount past data due to a concern about possible structural change of an unknown form.⁹

The asymptotic behavior of θ_t is a non-trivial issue because the model is self-referential. It turns out, though, that for small gains γ it is possible to obtain results on the asymptotics by studying a continuous time approximation to the recursive algorithm. More specifically, results from stochastic approximation theory show that asymptotically the dynamics are governed by the “mean dynamics” ordinary differential equation (ODE)

$$\frac{d\theta}{d\tau} = S^{-1}M(\theta)(T(\theta) - \theta) \quad (11)$$

$$\frac{dS}{d\tau} = M(\theta) - S \quad (12)$$

where $\tau = \gamma t$, and $M(\theta)$ is the unconditional covariance matrix of the regressors holding θ fixed. The ordinary differential equation governing the evolution of θ is identical to the E-stability differential equation with the exception that it includes weighting terms that depend on estimates of the regressors covariance matrix. See the Appendix for further details on the derivation of the ODE. It is straightforward to see that the fundamentals REE $(\bar{a}, \bar{b}) = (\bar{\pi}, 0)$ is a locally stable rest point of (11)-(12) provided $\alpha > 1$, in accordance with the earlier E-stability results.

Under decreasing gain learning γ is replaced with $1/t$ and, provided $\alpha > 1$ so that the fundamentals REE is E-stable, then locally the learning dynamics converge to it with probability one as $t \rightarrow \infty$. This paper focuses on constant gain learning, in which parameter estimates weight recent data more heavily than past. We next summarize the analytical results for constant gain learning.

2.3 Analytic Results

The first result establishes that for a sufficiently small constant gain the perceived coefficients θ_t will be an approximately normal random variable with a mean equal to its rational

⁹The possibility of a changing or drifting inflation target has implications for the mean and persistence of inflation (Cogley and Sbordone (2008)). Thus, agents should use a model that can track structural change in all coefficients.

expectations values and a variance that depends on both the constant gain and the long-run inflation target $\bar{\pi}$.¹⁰ The second result shows that from a given initial condition (θ_0, S_0) the solution to the “mean dynamics” of the ODE (11)-(12) gives the expected transition path to the rational expectations values.

Proposition 1 *Let $\alpha > 1$. The belief parameters θ_t are approximately distributed as $\theta_t \sim N(\bar{\theta}, \gamma V)$ for small $\gamma > 0$ and large t , where $\bar{\theta} = (\bar{\pi}, 0)'$ and*

$$V = \begin{pmatrix} (\alpha - 1)\bar{\pi}^2 + \alpha\sigma_r^2 & -\bar{\pi}/2 \\ -\bar{\pi}/2 & 1/2 \end{pmatrix}$$

Proposition 2 *Let $\alpha > 1$ and define $\phi'_t = (\theta'_t, \text{vec}(S_t)')$. For any ϕ_0 within a suitable neighborhood of the unique, non-explosive rational expectations equilibrium, define $\tilde{\phi}(\tau, \phi_0)$ as the solution to the differential equation (11)-(12), with initial condition ϕ_0 . Fix $T > 0$. The mean dynamics of (9)-(10) satisfy $E\phi_t \approx \tilde{\phi}(\gamma t, \phi_0)$ for γ sufficiently small and $0 \leq t < T/\gamma$.*

There are several important consequences from Proposition 1. First, the rational expectations equilibrium provides a benchmark solution in the sense that the coefficients for the forecast rule under learning are centered on the rational expectations values. Second, for a constant gain $\gamma \rightarrow 0$, the learning dynamics are close to the rational expectations equilibrium with high probability. Third, to gain insight into the global learning dynamics, for finite periods of time, one can study the solution paths to the mean dynamics differential equation, given initial conditions. The remainder of the paper uses these tools to study the implications for learning dynamics of inflation targets.

2.4 Learning Dynamics and Random-Walk Beliefs

Under constant gain learning there can be significant, temporary departures from RE. These departures can arise either as a result of a (imperfectly announced) change in the long-run target or as an endogenous response to exogenous shocks. This section illustrates these possibilities in the Fisherian model with constant gain learning.

Proposition 1 shows that the real-time estimates θ_t are approximately normal with a mean equal to the rational expectations equilibrium and a standard deviation that is increasing in the constant gain γ . To illustrate the implications this has for learning dynamics, Figure 2 plots the 95% confidence ellipses around the REE of the constant gain learning coefficients

¹⁰As discussed in the Appendix, Proposition 1 requires that the learning algorithm be supplemented with a suitable “projection facility” that restricts estimates to an appropriate compact set.

for various values of the long-run inflation target $\bar{\pi}$. This figure was generated by setting $\alpha = 1.05, \sigma_r^2 = 0.02$ and a constant gain $\gamma = 0.05$, though the qualitative results hold for alternative parameterizations. In particular, in the New Keynesian example, the qualitative results arise for a broad range of policy coefficients and rules.

Figure 2 demonstrates the finding in Proposition 1 that the constant gain parameter estimates are distributed around the rational expectations equilibrium, which is $(\bar{a}, \bar{b}) = (\bar{\pi}, 0)$. For small long-run inflation targets, the principal axis of the confidence ellipse is close to horizontal. For higher inflation targets, the confidence ellipses feature a decreasing principal axis. The slope of the principal axis is important since one can expect many trajectories moving in the direction of the axis. Note that even for high inflation targets, the ellipses are pointed in the direction of a random walk without drift, with larger values of b associated with smaller values of a . The relative size of these ellipses depends on the sizes of the constant gain. The confidence ellipses pointing toward a random walk without drift does not imply that actual learning dynamics will converge to a random walk model. However, the slope of the principal axes suggest that one can expect many trajectories moving in the direction of a random walk. Then the mean dynamics can help illustrate what happens subsequently for trajectories that move along the principal axis. As shown by Proposition 1, the variance of inflation in the Fisherian model is independent of $\bar{\pi}$ which suggests an endogenous route to random walk beliefs that is independent of the inflation target. We show later that this is not true in the New Keynesian model where, like Ascari and Ropele (2007), we find that long-run inflation targets can alter the qualitative nature of economic dynamics.

Proposition 2 shows that for any initial condition, and finite period of time, the expected transition path to the rational expectations equilibrium will be the solution path to the mean dynamics. One can think of constant gain learning dynamics as re-initializing the mean dynamics. The mean dynamics have the interpretation that, for given initial values, the average path for beliefs will follow the mean dynamics trajectory. Figure 3 plots representative mean dynamics where the initial values $a = 3.2, b = 0.33$ are on the principal axis of the 95% confidence ellipse in Figure 2 for a 4% inflation target.¹¹ Thus, along a stochastic equilibrium path we would expect to observe many finite stretches of time with properties similar to Figure 3. These initial values correspond to an increase in the perceived mean and perceived serial correlation in inflation. The top panel plots the perceived value for the mean of inflation, a , while the bottom panel plots the perceived lag coefficient b .

The fundamental rational expectations equilibrium is a stable rest point of the mean

¹¹The sample second moment matrix of the regressors, R , are initialized at their REE values.

dynamics, implying that along a learning path the mean dynamics will converge to the rational expectations equilibrium. Additionally, as anticipated in Figure 2, the transition path for the mean dynamics is interesting in its own right. At first the estimate for b moves toward the rational expectations equilibrium, slightly overshooting $b = 0$, but then reverses course and increases to a value of $b = 0.993$, where it remains for some time before returning to its rational expectations equilibrium value. At the same time, the value of a increases before abruptly decreasing to near zero (0.03) and then converging to its equilibrium value as b converges to zero. Note, in particular, that $a \approx 0$ at the same time that $b \approx 1$. Therefore, the mean dynamics show that private-sector agents come to believe temporarily that the inflation process is approximately a random walk. Paths to temporary random walks also exist for a lower 2% inflation target. However, for the New Keynesian model presented in the next section, we find that trend inflation (a positive target) changes the linearized inflation equation in a way that makes the system less stable under learning for higher targets in the sense that random-walk beliefs are more likely to arise.

Random-walk beliefs play a key role in the learning dynamics. In essence, agents come to believe that recent innovations in inflation are permanent shifts and not mean-reverting fluctuations. These random-walk beliefs are nearly self-fulfilling. A detailed argument is presented in Branch and Evans (2011), but an overview of the argument is useful. Suppose that agents hold random walk beliefs in terms of a forecasting model of the form

$$\pi_t = \pi_{t-1} + \varepsilon_t$$

which will arise in the learning model when $a = 0, b = 1$. Given these beliefs, actual inflation outcomes will be

$$\pi_t = \alpha^{-1}(\alpha - 1)\bar{\pi} + \alpha^{-1}\pi_{t-1} + \alpha^{-1}r_t$$

If $\alpha > 1$ is close to $\alpha = 1$, then the actual law of motion for inflation is a stationary but highly persistent process that is difficult to distinguish from a random walk.¹² Also, the mean inflation rate is the same under random walk beliefs as it is in the unique REE. That random-walk beliefs are nearly self-fulfilling has been pointed out in other settings by Sargent (1999) and Lansing (2009). Random walk beliefs introduce serial correlation into a model that is not serially correlated under rational expectations, as the random walk model uses

¹²For larger values of α , the random-walk model provides a progressively worse approximation to actual inflation dynamics. Thus, for α sufficiently large random-walk beliefs will not emerge from the learning dynamics. A somewhat related point has been made by Orphanides and Williams (2005b). However, in the New Keynesian model below, inflation persistence is increasing in the target rate of inflation and random-walk beliefs will arise for a much greater range of policy coefficients.

higher order moments to track low frequency drift in inflation. An increase in $\bar{\pi}$, combined with adaptive learning, provides the low frequency drift that is approximated well by a random walk model.

The mean dynamics show that random walk beliefs only can last for finite stretches of time. However, because the random walk beliefs are nearly self-fulfilling, it is difficult to detect the misspecification except using a long history of data. Most importantly, the random walk model provides a robust way to capture a time-varying conditional mean. When this drift is large enough then random-walk beliefs will fit the data well. Thus, random-walk beliefs can be long-lasting and, as will be seen below, they have important implications for the dynamics of inflation.

2.5 Implications of Inflation Targets

Having established the possibility of random-walk beliefs emerging under learning, we turn briefly to real-time simulations.

Consider the following experiment. The central bank is going to implement an increase in its long-run annual target from 2% to 4%. Assume that the economy is initially in a rational expectations equilibrium, but the private-sector has imperfect information about when the central bank will implement its new target and is unsure about the central bank's commitment to the new target, i.e. the central bank provides imperfect information about the new target so that private sector agents must infer its value from observable data. Figure 4 plots the resulting dynamics from constant gain learning.¹³ At time 0, the central bank's target $\bar{\pi}$ increases, leading to an increase in inflation without a corresponding increase in inflation expectations (which are determined by the adaptive learning rule). Initially inflation is below target and the central bank begins reducing nominal interest rates in order to bring inflation up to target. The increase in the inflation rate is tracked by agents' econometric model as an increase in the persistence of inflation. As the mean dynamics predict, eventually agents' beliefs are that inflation follows a random walk. At this point, there is a burst as inflation increases to nearly 8% before returning to its new long-run value. Thus, implementing a higher target, as many observers have recommended, can lead to an overshooting of the new target. This overshooting arises because the initial upward drift in inflation, as the central bank implements its new target, leads to a nearly self-fulfilling belief that inflation follows a random walk.¹⁴

¹³This figure was generated as the average time-path across 1000 stochastic simulations of length 1000. Beliefs are initialized at their REE values for when the target is 2%.

¹⁴These results are reminiscent of McGough (2006) who examines changes to the natural rate of unem-

Even without a change in the long-run target, inflation may deviate substantially from its rational expectations equilibrium value as an endogenous response to fundamental shocks. For example, Figure 5 plots two real-time simulations of inflation dynamics in the Fisherian model with a 4% inflation target.¹⁵ As before, the figure is computed setting $\alpha = 1.1$, $\sigma_r^2 = .1$, and as in Figure 2 the constant gain is set $\gamma = .05$. To generate this figure the model is initialized at the REE, expectations are generated according to (5) with parameters updated via constant gain least-squares, and inflation is determined by (3). Figure 5 plots the results for two different typical simulations. The top two panels plots a_t, b_t , respectively, and the bottom panel plots inflation. The left panels are for the case where the deviation from equilibrium results in a burst of inflation, while the right panels show a rapid disinflation.

Under constant gain learning the economy hovers around its rational expectations equilibrium value. Then there is an abrupt qualitative change in the dynamics with bursts of inflation or disinflation before returning to a neighborhood of the rational expectations equilibrium. The pattern of beliefs correspond with what was observed in Figure 3, and Proposition 2, in that for finite stretches of time private-sector agents believe inflation follows a random walk. In simulations, these large deviations from rational expectations are recurrent.

Unlike Figure 4, the deviations away from the rational expectations equilibrium in Figure 5 are an endogenous response to fundamentals rather than to a change in the long-run inflation target of the central bank. Using techniques employed by Cho, Williams, and Sargent (2002), it is possible to examine which “escape paths” are most likely to drive the system away from the REE and to generate random walk beliefs by looking for the “most likely unlikely sequences” of shocks that move the system a given distance away from the equilibrium. In principle one can compute these escape paths analytically in special cases, but more typically it is necessary to resort to simulations.

However, it is intuitive that random walk beliefs can arise for the right sequence of shocks. Take the case of positive inflationary shocks. These shocks place inflation on an upward path leading agents’ econometric model to pick up this trend with higher estimated values of b_t and lower values of a_t . In turn, inflation expectations will increase leading to a further upward drift in inflation, higher estimated values of b_t , until the estimated coefficients arrive at a random walk model which, as argued above, is nearly self-confirming. Moreover, the mean dynamics predict that even if beliefs are, on average, close to rational expectations this

ployment in the model of policymaker learning developed in Sargent (1999).

¹⁵Similar qualitative results arise from stochastic simulations with a 2% target in the Fisherian model. However, in the New Keynesian model these unstable dynamics do depend on the size of the target.

is the expected transition path following a series of these “most likely unlikely” sequences of shocks.

There is one strong conclusion to draw from Figures 4-5: in the Fisherian model, when the central bank implements a long-run inflation target with imperfect information, then inflation will deviate significantly from its equilibrium values (i) when the target is first implemented and (ii) as an endogenous response to certain sequences of shocks. The remainder of the paper demonstrates that these results are found in standard New Keynesian models under a wide range of policy rules. In particular, it is demonstrated that higher inflation targets increase the likelihood of observing these large, temporary deviations in inflation.

2.6 Further Discussion

The results presented in this section illustrate that, in a flexible price model, with imperfect knowledge about the inflation target there can be unstable inflation dynamics that arise from changes in the target or as an endogenous response to the exogenous shocks. In subsequent sections, it is shown that in a New Keynesian model with nominal rigidities, a higher target also makes it more likely that destabilizing random-walk beliefs can arise.

Another natural question is whether there are alternative ways for the central bank to implement a higher target without destabilizing expectations. To address this issue, we briefly consider two alternative experiments. The first implements the target gradually over a longer horizon than just one time period. The second experiment considers a central bank that initially targets an inflation rate greater than the desired 4% target and then gradually lowers the target to its desired level. Figure 6 illustrates how these alternatives for implementing a higher target affect the over-shooting result that was seen in Figure 4.¹⁶

The left panel of Figure 6 plots the dynamics when the inflation target is increased gradually over a 1,8,16,32,64,128, and 256 period horizon. A very gradual implementation of the target can avoid the destabilizing random-walk beliefs by not creating considerable positive autocorrelation in the actual inflation process. However, for this parameterization, the implementation of the target needs to take place over a very long horizon. The right panel considers an experiment where the central bank initially over-shoots its 4% target and then gradually lowers its target to the desired 4%. This implementation induces a smaller over-shooting of the 4% target since over-shooting and then bringing inflation down avoids the strong degree of positive autocorrelation that can lead to destabilizing random-walk beliefs.

¹⁶Note that Figure 6 plots the mean paths across 1000 stochastic simulations.

3 Application to the New Keynesian Model

The previous section adopted a simple Fisherian model of inflation to illustrate that setting policy to implement a long-run inflation target in an imperfect information environment can lead to substantial deviations of inflation from its rational expectations equilibrium value. We now show that similar results obtain in the New Keynesian model with trend inflation (see Ascari and Ropele (2007)). The richer setting of the New Keynesian model facilitates a wider exploration of the generality of the results as well as policy implications.

3.1 A New Keynesian Model with Imperfect Information

Ascari and Ropele (2007) take a standard New Keynesian setting log-linearized around a non-zero steady-state rate of inflation.¹⁷ (See Appendix for details.) They show that this leads to the following equations that determine aggregate output and inflation

$$\hat{x}_t = \hat{E}_t \hat{x}_{t+1} - \sigma^{-1} \left(\hat{i}_t - \hat{E}_t \hat{\pi}_{t+1} - r_t^n \right) \quad (13)$$

$$\hat{\pi}_t = \theta_1 \hat{x}_t + \theta_2 \hat{E}_t \hat{\pi}_{t+1} + \sum_{j \geq 0} \xi_2^j \left(\theta_3 \hat{E}_t \hat{x}_{t+1+j} + \theta_4 \hat{E}_t \hat{\pi}_{t+2+j} \right) + u_t \quad (14)$$

where the reduced-form parameters $\theta_k, k = 1, \dots, 4$ and ξ_2 are complicated expressions that depend on the underlying structural parameters. (See the Appendix for details). Here \hat{i}_t is the deviation of interest rate from its steady state value and $\hat{x}_t, \hat{\pi}_t$ are log deviations of the output gap and the inflation rate, respectively, from steady-state. The shocks r_t^n, u_t are assumed for simplicity to be zero-mean iid with variance σ_r^2, σ_u^2 . Here, as in Section 2, \hat{E} denotes the subjective expectation of the agents, which need not be identical to “rational expectations.”

The model is closed by assuming that monetary policy sets nominal interest rates according to the Taylor rule, $\hat{i}_t = \alpha_\pi (\pi_t - \bar{\pi}) + \alpha_x (x_t - \bar{x})$, i.e.

$$\hat{i}_t = \alpha_\pi \hat{\pi}_t + \alpha_x \hat{x}_t$$

Stacking these equations, it is possible to represent the economy in vector form

$$Y_t = H + F \hat{E}_t Y_{t+1} + M \sum_{j=1}^{\infty} \xi_2^j \hat{E}_t Y_{t+1+j} + G z_t$$

¹⁷This paper adopts the version of Ascari-Ropele without price indexation. We omit indexation for two reasons. First, indexation does not seem to be empirically relevant in a model that includes trend inflation (see Ascari and Sbordone (2014)). Second, including indexation in an environment with imperfect knowledge about the trend inflation rate raises interesting theoretical issues that are beyond the scope of the present paper.

with $Y = (x, \pi)'$, now written in (log) levels, and $z = (r^n, u)'$. When the inflation factor $\Pi = 1$, i.e. when there is zero steady-state inflation, these equations reduce to the benchmark New Keynesian model:

$$\hat{x}_t = \hat{E}_t \hat{x}_{t+1} - \sigma^{-1} \left(\hat{i}_t - \hat{E}_t \hat{\pi}_{t+1} - r_t^n \right) \quad (15)$$

$$\hat{\pi}_t = \beta \hat{E}_t \hat{\pi}_{t+1} + \kappa \hat{x}_t + u_t \quad (16)$$

Ascari and Ropele demonstrate that, under most parameterizations, steady-state inflation lowers steady-state output and it alters the determinacy conditions when monetary policy follows a Taylor rule. Kobayashi and Muto (2013) show that E-stability properties of the model can also differ from the New Keynesian model linearized around a zero steady-state inflation rate. Throughout, monetary policy will be assumed to guarantee the model is determinate and E-stable.

As in the previous section, it is assumed that there is imperfect information about the economy, and so private-sector agents base forecasts on a perceived law of motion whose reduced-form is consistent with a rational expectations equilibrium. In the bivariate New Keynesian model, agents are assumed to forecast based on a simple VAR(1) model:

$$Y_t = A + BY_{t-1} + \varepsilon_t \quad (17)$$

where ε_t is a perceived white noise error. From this perceived law of motion, conditional expectations can be computed. As before, expectations conditional on the forecast model (17) are

$$\hat{E}_t Y_{t+1} = (I + B)A + B^2 Y_{t-1}. \quad (18)$$

As in the Fisherian model, the actual law of motion can be found by substituting the forecasts (18) into the expectational difference equation:

$$\begin{aligned} Y_t &= H + F(I + B)A + (1 - \xi_2)^{-1} M [I + (I - \xi_2 B)^{-1} B] A \\ &+ [F + M(I - \xi_2 B)^{-1}] B^2 Y_{t-1} + Gz_t \\ &\equiv T(A_{t-1}, B_{t-1})' \begin{bmatrix} 1 \\ Y_{t-1} \end{bmatrix} + Gz_t \end{aligned}$$

Is this learning rule reasonable in this economic environment? We think the answer is clearly yes. The possibility of an imperfectly understood change in the inflation target can lead forecasters to be uncertain about both the mean and persistence of inflation (e.g. Cogley and Sbordone (2008)). An econometric model that did not allow for structural change in both

the constant and slope coefficients would not track a change in the inflation target. Finally, the learning rule assumes agents forecast based on a parsimonious VAR(1), a commonly employed practice by professional forecasters.

Analytic results on E-stability, and convergence of constant gain learning, are unavailable in the New Keynesian model with trend inflation. Instead, the following sections focus on a calibrated version of the model and present numerical results that illustrate the theoretical possibilities. The model’s parameters are calibrated in Table 1. The parameter values are chosen so that a time period corresponds to a quarter. The value of ζ implies a mark-up of 11%, the policy coefficients are equivalent to the Taylor (1993) original policy prescription, the frequency of price adjustment α is in line with most empirical studies, while the values of σ_r^2, σ_u^2 come from Smets and Wouters (2007). The gain parameter is calibrated to 0.05, within the range found to be consistent with professional forecasters in Branch and Evans (2006). The basic qualitative results do not hinge on the specifics of the calibration. The necessary ingredients are a non-zero inflation target, imperfect information, and a sufficiently strong feedback from expectations.¹⁸

3.2 Inflation Dynamics in the New Keynesian Model with Learning

The “mean dynamics” for the New Keynesian model take the same form as in the Fisherian model with the exception that the multivariate model complicates the expressions. Nevertheless, as in the Fisherian model a great deal can be learned about learning dynamics by examining the mean dynamics for the calibrated New Keynesian model. The mean dynamics are the solution path to the ordinary differential equation (11)-(12) where $\theta' = (A, B)$.

Figure 8 plots the mean dynamics, with the parameters calibrated as in Table 1, and a 4% inflation target. The initial values for all coefficients, except the constant and the own lag coefficient in the inflation component of the forecasting model, are set to their REE values. The remaining initial conditions were chosen to lie on the principal axis from the 95% confidence ellipse, see Figure 7.¹⁹ There are 6 coefficients in θ , the two constants and

¹⁸In the Fisherian model, for sufficiently large values of α , a random-walk forecasting model provides a poor approximation to the implied inflation dynamics. In the New Keynesian model with trend inflation, similar results obtain for sufficiently large values of α_π and small inflation targets. Random-walk beliefs still arise, however, for empirically plausible values of α_π .

¹⁹Specifically, Figure 8 sets the coefficients in the output gap forecasting equation at their REE values. The results are robust to most positive initial values for the constant and when the own lag coefficient in the inflation equation were between 0.6 and .71. The 95% confidence ellipse were generated from monte-carlo

the four coefficients in the lag matrix B . Each panel plots a different component of θ , while the bottom two panels plot the roots of the matrix B along the mean learning path. The key result to notice in Figure 8 are the first two panels on the right hand side of the figure. These plots are equivalent to the mean dynamics plots in the Fisherian model. Notice that agents come to believe that the process for inflation, but not the output gap, is a random-walk without drift. Private-sector agents can hold these beliefs temporarily before they converge to the rational expectations equilibrium. The bottom panel on the right shows that along the mean path there is, in fact, a unit root in the perceived coefficient matrix B .

Figure 9 solves the mean dynamics just as in Figure 8 for long-run targets of 0%, 1%, 2% and 4%.²⁰ The figure plots just the constant and the own lag coefficient from the inflation forecasting equation, while suppressing the other coefficient paths. As can be seen in the figure, when there is a zero inflation target the beliefs converge monotonically and rapidly to the rational expectations equilibrium. For successively larger values of the inflation target, the estimates for $B(2, 2)$ begin to move towards the REE value and then reverses course and approaches $B(2, 2) \approx 1$ when the target $\bar{\pi} = 4\%$. At the same time, the estimated $A(2)$ moves towards zero before converging to the mean value. When the inflation target is 4% random walk beliefs emerge.

In Section 2.4, random-walk beliefs could arise when $\alpha \approx 1$ so that a random-walk provides a good approximation to the actual inflation process. A similar intuition holds in the New Keynesian model, though, now the coefficient on expectations in equation (14) – the analogue to α^{-1} in the Fisherian model – depends on the long-run inflation target, the policy rule coefficients, and other structural parameters. Figure 10 demonstrates that in the New Keynesian model higher inflation targets make random-walk beliefs more likely as the largest root of B approaches a unit root as the inflation target nears 4%. Thus, in the NK model higher targets, combined with adaptive learning, give the low frequency drift that is well-approximated by a random walk model. This result is related to Ascari and Ropele (2007) who show that higher targets make indeterminacy more likely. The random-walk beliefs arising in cases of higher targets are a different form of (nearly) self-confirming expectations. The next subsection shows how altering the policy coefficients affect these results.

The US Federal Reserve Bank has an explicit inflation target of 2%. One popular argu-

stochastic simulations of the model with a constant gain of 0.05, assuming a Gaussian asymptotic distribution for the coefficients.

²⁰As before, in Figures 9 and 10, the coefficients A, B, S were initialized at their REE values except the coefficients $A(2)$ and $B(2, 2)$ are drawn from the 95% confidence interval in Figure 7.

ment to avoid the possibility that in the future the zero lower bound will again bind, is to increase the target to 4%. Figure 11 conducts this experiment in the New Keynesian model with imperfect information. The model is calibrated according to Table 1, initialized in the unique REE when $\bar{\pi} = 2\%$. The central bank then increases its target to 4% and the private-sector must learn in real time about the new higher average rate of inflation. The figure plots the time-series averaged across 1000 stochastic simulations. The figure demonstrates that, as in the Fisherian model, there is a significant overshooting of inflation as agents come to believe that inflation follows a random walk. In this experiment inflation increases to above 10% per annum before returning to the new targeted value of 4%.

Many of those who advocate higher inflation targets do so in order to avoid abrupt disinflationary episodes. The results from the Fisherian model suggest that this may not be the case with imperfect information. Figure 12 plots two separate simulations of the real-time learning of the calibrated New Keynesian model with a 4% inflation target. In each panel, the model is initialized in a rational expectations equilibrium, within which inflation will always remain bounded in a small neighborhood of the inflation target. Under constant gain learning, however, there is an abrupt qualitative change as agents come to believe that inflation follows a random walk and there is a rapid disinflation, in the left panel, and an abrupt inflationary episode in the right panel. These types of destabilizing dynamics are typical and recurrent in the model when the central bank sets a sufficiently large inflation target. In this case, sufficiently large is 4%. The onset of destabilizing dynamics occurs, on average, once every 92 years. Larger constant gains will, however, increase the frequency.

One might also wonder whether these inflation dynamics are consistent with the Volcker era, a period where the Federal Reserve lowered the inflation target in an imperfectly communicated manner. Since the inflation dynamics for a drop in inflation in the model are symmetric, the model predicts a disinflation will lead to inflation and inflation expectations below the target value. While a serious study of this episode is beyond the scope of the present paper, it is worth noting that inflation expectations in the Michigan survey and the Survey of Professional Forecasters fell below the 10 year expected inflation rate in the Blue Chip survey and the Survey of Professional Forecasters for much of the mid to late 1980's, consistent with the model's predictions.²¹

²¹Nunes (2009) shows that a heterogeneous expectations model, where a fraction have rational expectations and the remaining form beliefs with an adaptive learning rule of the MSV form, generates time paths for inflation that fit well the Volker disinflation.

3.3 Robustness

One reasonable question is to what extent the benchmark results presented so far depend on relatively weak responses to inflation and the output gap in the policy coefficients α_π, α_x . Though the values $\alpha_\pi = 1.5, \alpha_x = .125$ are the values recommended by Taylor (1993), some empirical and optimal Taylor rules have values for α_π higher, for example Rudebusch (2001) estimates $\alpha_\pi = 1.78, \alpha_x = .82$. Moreover, the discussion in Section 2.4 seems to suggest that random-walk beliefs only arise for values of α_π sufficiently close to 1. This subsection demonstrates that random-walk beliefs can arise for a broad range of policy coefficients and other parameters thereby demonstrating that the main insights are not fragile to a particular model calibration.

We first illustrate the robustness to alternative policy coefficients in Figure 13 by computing the pairs of (α_x, α_π) for which random-walk beliefs will arise in the mean dynamics with a target of 4% and the parameters in Table 1.²² The figure demonstrates that random-walk beliefs can arise for many empirically plausible policy coefficients and initial coefficients picked from plausible points based on the 95% confidence ellipse.

Figure 13 suggests that a policymaker seeking to implement a higher inflation target should also adjust the coefficients in the policy rule so that the new policy rule lies above the frontier. We consider these issues in Section 4.2 which discusses ways in which policymakers can alter the policy rule to implement a higher target without triggering unstable inflation expectations.

The robustness also extends to alternative parameterizations. In particular, random-walk beliefs were found to arise for the following range of parameter values: $7 \leq \zeta \leq 17$; $.6 \leq \alpha \leq 1$; $.01 \leq \sigma_r^2 \leq 1$; $.001 \leq \sigma_u^2 \leq 0.1$. Finally, one might prefer a medium-scale New Keynesian model known to fit the data better, under rational expectations, than the purely forward-looking model. However, random-walk beliefs are even more likely to arise in a model with inertia as this paper shows they can arise in a model with no intrinsic serial correlation. The self-fulfilling serial correlation that arises from the learning dynamics is even more prevalent in models with inertia or other forms of intrinsic serial correlation.

²²Figure 13 was generated by setting all initial learning coefficients, including the constant, to their equilibrium value except $B(2, 2)$, the own inflation lag coefficient, is initialized at values in $[0.6, 0.8]$. As mentioned earlier, the results are not sensitive to the initial values for the constant. If any value of the own inflation lag coefficient crosses 0.98 – for at least one initial value in $[0.6, 0.8]$ – then the pair (α_x, α_π) is recorded as leading to random-walk beliefs.

3.4 Zero Lower Bound

One of the primary reasons that some economists and policymakers advocate a higher long-run inflation target is the belief that this would reduce the likelihood of a binding zero lower bound on interest rates. In New Keynesian models, e.g. Eggertsson and Woodford (2003), Eggertsson (2008), the zero lower bound on nominal interest rates can bind when there are sufficiently large and persistent negative shocks. A liquidity trap with deflation can arise, under rational expectations, since the persistent shocks generate expectations of deflation. Under learning, Williams (2006) and Williams and Reifschneider (2000), shows that the zero lower bound may be reached even more often. In both cases, increasing the inflation target makes it less likely that the zero lower bound will bind.

Within the New Keynesian model, we have so far not analyzed the issue of the zero lower bound, since the main message of this paper is the destabilizing learning dynamics that can arise under long-run inflation targets. However, it is straightforward to extend the framework to incorporate a zero lower bound and examine whether a higher inflation target can rule out a binding zero lower bound on nominal interest rates. To a certain extent, this is not a fair comparison: in our model, under rational expectations, the zero lower bound will almost never bind since all shocks are iid with small variances. However, if there can be deflationary spirals under learning when there is no liquidity trap under rational expectations, there is even greater reason to suspect that inflation targets can lead to destabilizing economic dynamics.

To study this issue it is convenient to use the “benchmark” version of the New Keynesian model:

$$\begin{aligned}\hat{x}_t &= \hat{E}_t \hat{x}_{t+1} - \sigma^{-1} \left(i_t - \bar{r} - \hat{E}_t \pi_{t+1} - r_t^n \right) \\ \pi_t &= \beta \hat{E}_t \pi_{t+1} + \kappa \hat{x}_t + u_t,\end{aligned}$$

where the central bank employs the Taylor rule subject to the zero lower bound:

$$i_t = \max\{\bar{i} + \alpha_\pi (\pi_t - \bar{\pi}) + \alpha_x (\hat{x}_t - \bar{x}), 0\}$$

and where $\bar{x} = (1 - \beta)\bar{\pi}/\kappa$ ensures the steady-state inflation rate equals the central bank’s target. This is the framework adopted in Eggertsson and Woodford (2003), though here it is assumed r_t^n are iid shocks, and can be derived from the New Keynesian model above by setting the steady-state inflation equal to zero. Woodford (2003) argues that these equations are valid approximations in a low inflation environment so long as the central bank’s inflation target is not too large.

Calibrating the model as above which implies a value for $\kappa \approx 0.16$, with a 4% inflation target, and simulating it under constant gain learning leads to occasional departures from a neighborhood of the rational expectations equilibrium and a deflationary spiral: see Figure 14.²³ At first, the economy fluctuates around the long-run inflation target. Then a qualitative change in the dynamics arises as agents come to believe inflation follows a random walk and inflation (and expected inflation) rapidly disinflates as was seen earlier. However now, when the zero lower bound binds, rapid deflation sets in, resulting in a severe recession, with output gaps approaching -10% . The spiral does not persist forever as the mean dynamics take over and the economy returns to a neighborhood of the steady-state.²⁴

The deflation episode in Figure 14 suggests that a higher inflation target cannot rule out paths that collapse to the zero lower bound. In the benchmark New Keynesian model, where the coefficients in the IS and AS equations are independent of the inflation target, a higher target will decrease the frequency of collapses to the zero lower bound. Nevertheless, a higher target cannot rule out destabilizing dynamics that hit the target under learning.

4 Extensions

The paper concludes with several important extensions to the New Keynesian model with learning.

4.1 Implications for Policy Communication

The previous results assumed that monetary policy is set to satisfy the Taylor rule. The results do not hinge on this assumption: the key is that the target is not perfectly communicated. This section considers implications for policy communication of the present learning framework with long-run inflation targets.

The assumption maintained in this paper is that the central bank cannot completely communicate the long-run inflation target. The private-sector has imperfect information about the economic environment and the policy target. It suffices that they do not know the timing of when the central bank will implement its target. If there is imperfect information about the economy, but perfect information about the long-run inflation target then the

²³These deflationary spirals are similar to the paths in Evans, Guse, and Honkapohja (2008).

²⁴To prevent deflationary paths that fall without bound, we impose the following restrictions on the learning and economic dynamics: (i) agents' only update their parameter estimates of B if the roots lie inside the unit circle, and (ii) the deflation rate and the output gap have ceilings of 8% and 10% per annum, respectively. The latter restriction is consistent with Evans (2013) and Bullard and Cho (2005).

random-walk beliefs will not arise and the instability witnessed in the previous sections will not exist.

To illustrate this point, assume that agents know the mean values for inflation and the output gap. Thus, their perceived law of motion does not require them to estimate the mean rates of inflation and the output gap, but just the lag coefficients. That is, assume a perceived law of motion of the form:

$$\hat{Y}_t = B\hat{Y}_{t-1} + \varepsilon_t$$

where the perceived law of motion is written in deviation from mean form, \hat{Y} . It is straightforward to verify that the actual law of motion will be of a similar form, except written in deviation from mean form. Thus, the B component of the T-map is unchanged. Figure 15 plots the mean dynamics path for the calibrated model and various long-run inflation targets when the long-run average inflation rate is credibly and perfectly communicated. The figure clearly demonstrates that in this case the long-run inflation target does not have much impact on the learning dynamics and that random-walk beliefs do not arise.

4.2 Altering the Policy Rule

The results just presented clearly demonstrate that a perfectly communicated inflation target can avoid the destabilizing learning dynamics associated with imperfect information about the target. But, Figure 13 also suggests that if the central bank responds more aggressively to inflation deviations from target, via the response coefficient α_π , it may be possible to rule out the destabilizing random-walk beliefs.

This subsection now demonstrates that a central bank can minimize the possibility of destabilizing inflation expectations by simultaneously increasing the inflation response coefficient. Figure 16 again considers the experiment of a central bank increasing its long-run inflation target from 2% to 4% while, at the same time, increasing its inflation response coefficient to $\alpha_\pi = 3$ from $\alpha_\pi = 1.5$. Figure 16 shows that such a policy can prevent random-walk beliefs and bring the economy to the long-run target relatively quickly. Although, such a policy can be effective at implementing a higher target it is worth remarking that $\alpha_\pi = 3$ is outside of the range most studies consider reasonable.

Moreover, a stronger inflation reaction coefficient α_π will also dampen the extent to which inflation overshoots the higher target. For instance, when the central bank raises the target from 2% to 4% and simultaneously increases α_π from 1.5 to 2.0, inflation overshoots the target with a peak inflation rate of 7%, less than when the target is changed alone. Thus, when the policy coefficients are chosen to lie in the “stable” region of Figure 12, a central

bank can implement the target without the potential for instability. Thus, one take-away from the findings in this paper is that the optimal choice of inflation target can not be isolated from the optimal choice of policy instrument.²⁵

Similarly, a policy rule that responds aggressively to changes in the inflation rate, as well as contemporaneous inflation and the output gap, can also help stabilize the economy while implementing a higher target. Such a policy rule would behave like a Taylor rule in ‘normal’ times and then respond aggressively in the beginning of one of the destabilizing spirals described in this paper.

4.3 Optimal Discretionary Policy

The interest-rate rules so far considered in this Section dictate that nominal interest rates should be adjusted to deviations of inflation and output, from their long-run target values, using rule-of-thumb adjustment parameters. One might instead imagine a central bank facing a dual mandate of stabilizing both inflation and output, with the costs of deviations from long-run targets made explicit according to an optimal policy problem. This subsection demonstrates that random-walk beliefs also arise in such a setting under optimal discretionary policy, i.e. optimal policy without commitment. The next subsection will show that even with commitment it is possible for random-walk beliefs to arise.

For simplicity, assume the economy can be represented by the benchmark New Keynesian equations (15)-(16). The objective function of the central bank is

$$\max_{\pi_t, x_t} -(1/2)E_0 \sum_{t \geq 0} \beta^t [\lambda(x_t - \bar{x})^2 + (\pi_t - \bar{\pi})^2]$$

where \bar{x} is the long-run output gap consistent with the inflation target $\bar{\pi}$. The central bank takes the New Keynesian Phillips Curve (16) as its constraint. Without commitment, the central bank will set policy to satisfy its first order condition

$$\pi_t - \bar{\pi} = -\frac{\lambda}{\kappa}(x_t - \bar{x}) \tag{19}$$

Combining (19) with (16) gives an expectational difference equation that generates the stochastic process for inflation

$$\pi_t = \alpha_0 + \alpha_1 \hat{E}_t \pi_{t+1} + \nu_t \tag{20}$$

²⁵This is related to the finding of Cogley, Matthes, and Sbordone (2015) that an optimal disinflation policy with Taylor-type rules, and where private-sector agents learning about the policy rule, involves a more aggressive response coefficient.

where $\alpha_0 = \frac{\kappa^2 + \lambda(1-\beta)}{\lambda + \kappa^2} \bar{\pi}$, $\alpha_1 = \frac{\beta\lambda}{\lambda + \kappa^2}$, and ν_t is an appropriately defined white-noise shock. The reduced form (20) is identical to the Fisherian model of section 2 with α_1 determined by β, λ, κ . Proceeding in the same manner as in section 2, agents form their expectations from an AR(1) model of inflation. Figure 17 plots the resulting mean dynamics.

Figure 17 was created by solving the mean dynamics under the following parameterization: $\beta = 0.99, \kappa = .14, \lambda = .15, \sigma_u = .05$. As before, the learning coefficient for b was initialized above its rational expectations equilibrium value, and then the mean dynamics are solved to find the transition path back to the unique rational expectations equilibrium. The transition path leads temporarily to random-walk beliefs despite the fact that policy is now set optimally.²⁶ We remark, however, that small values of λ – for instance, where λ is equal to its micro-founded calibrated value – yield small values of α_1 , which will make random-walk beliefs less likely. In practice, policymakers are often modeled as using a higher value of λ than the micro-calibrated value. For example, Orphanides and Williams (2007) use as their benchmark equal weights on inflation and output volatility. Our finding that small λ , i.e. a strong weight on inflation stabilization, make random-walk beliefs less likely is consistent with Orphanides and Williams (2005b), who show that it is needed when agents are learning even if policymakers would otherwise place a higher weight on output stabilization.

4.4 Price-level Targeting

The results in this paper demonstrate that, in addition to the other consequences of having a higher long-run inflation target, the target will be destabilizing in an imperfect information environment. One popular alternative to inflation targeting is price-level targeting. Woodford (2003) and Vestin (2006) shows that a policy to target a price-level path can implement the optimal policy with commitment. Eggertsson and Woodford (2003) show that a price-level target can be an effective policy for pulling an economy away from the zero-lower bound.²⁷

²⁶Evans and Honkapohja (2003) show that the “fundamentals-based” nominal interest-rate rule, which aims to implement the optimal discretionary equilibrium using a rule that *assumes* rational expectations, will lead to indeterminacy and instability under learning. They also show that if policymakers use a suitable “expectations-based” rule consistent with optimal policy, then the optimal rational expectations equilibrium will be determinate and stable under learning. However, random-walk beliefs can still sometimes arise under their policy rule if agents use constant-gain learning and include lagged inflation as a regressor.

²⁷Woodford (2003) shows that price-level targeting rules are robustly optimal when policymakers can commit. An alternative form of the optimal commitment rules is $\pi_t = -(\lambda/\kappa)(x_t - x_{t-1})$. We considered this policy rule, as well, though to implement learning is cumbersome as the system is three-dimensional. However, following Evans and Honkapohja (2006), who formulate an “expectations-based” reaction function,

However, what if, as with the inflation target, the central bank is unable to perfectly communicate the precise value or timing for the price-level target? Can price-level targeting lead to temporarily unstable learning dynamics just as in the case of long-run inflation targets? To address this issue, this section considers a central bank that acts in accordance with the following price-targeting rule:

$$\kappa p_t + \lambda x_t = p^* \quad (21)$$

where p_t is the (log) price-level and p^* is the target value for the (log) price-level. This policy rule will implement the optimal policy under commitment for a zero long-run inflation target. The qualitative results below carry over to the case where the central bank targets a price-level path consistent with a non-zero long-run inflation target.

Using the identity $\pi_t = p_t - p_{t-1}$, plugging (21) into the benchmark NK aggregate supply equation (16) leads to the following equation for the price-level

$$p_t = \alpha_0 + \alpha_1 \hat{E}_t (p_{t+1} - p_t) + \alpha_2 p_{t-1} + \eta_t \quad (22)$$

where $\alpha_0 = (\kappa^2/(\lambda + \kappa^2))\bar{\pi}$, $\alpha_1 = \beta\lambda/(\lambda + \kappa^2)$, $\alpha_2 = \lambda/(\lambda + \kappa^2)$. Notice that p_t depends on $\hat{E}_t p_t$ under imperfect information because we assume that p_t is not contemporaneously observable. Again, suppose that private-sector agents forecast the price level according to the forecasting model

$$p_t = a + bp_{t-1} + \varepsilon_t \Rightarrow E_t p_t = a + bp_{t-1}, \quad E_t p_{t+1} = a(1 + b) + b^2 p_{t-1}.$$

The actual price-level process is found by plugging these expectations into (22), yielding

$$p_t = T(a, b)' X_{t-1} + \nu_t$$

where $T(a, b)' = (\alpha_0 + \alpha_1 ab, \alpha_1 b(b - 1) + \alpha_2)$.

Figure 18 plots the mean dynamics for the price-level targeting rule case under the same calibrated parameter values as the previous subsection (Figure 17) and $p^* = 10$. A key difference with the price-target rule, compared to the rules considered earlier, is that the rational expectations equilibrium exhibits non-zero serial correlation. For the chosen parameter values the REE value of b is approximately 0.7. Figure 18 initializes the learning coefficients at 0.77. The transition path first leads away from the rational expectations equilibrium, and then abruptly changes course heading towards a random-walk model for the price-level

it is straightforward to derive a bivariate system of expectational difference equations similar, in form, to the price-level targeting expression. As in the case of discretion, random-walk beliefs arise in this case for values of λ that are not too small.

before finally converging to the rational expectations equilibrium. Thus, the results of this section demonstrate that if there is imperfect information about the price-level target, then price-level targeting policy rules can also lead to temporarily unstable inflation dynamics.

5 Conclusion

Long-run inflation targets, on the order of 4% per annum, have sometimes been recommended to guard against liquidity traps and binding zero constraints on nominal interest rates. These recommendations persist even though many welfare analyses caution against this approach since the distortions resulting from higher average inflation are often found to outweigh any gains from stabilizing inflation. Both the arguments for and against higher inflation targets have typically been made under the rational expectations assumption. This paper has revisited the issue of raising the inflation target, focusing on the question of whether higher targets do, in fact, lead to greater stability.

The primary results of this paper are as follows. First, although over time beliefs converge toward rational expectations, the combination of constant gain learning and a positive inflation target can lead agents in the economy to temporarily believe that the inflation process follows a random walk without drift. Such beliefs are temporarily (almost) self-confirming. When agents perceive the inflation process to be a random walk they will interpret recent innovations to inflation as permanent shifts in the mean inflation rate. These random walk beliefs arise for a very intuitive reason. The long-run inflation target, and imperfect information about that target, lead agents to estimate the mean inflation rate from real-time data. If data lead to a slight upward drift in the inflation rate, agents' econometric model will pick up that drift, leading to higher inflation expectations that feed back into higher inflation rates. This process is self-reinforcing and in some cases agents eventually come to believe that inflation follows a random walk. Crucially, we have shown that these beliefs are nearly self-fulfilling.

Implementing a higher target – say by moving the target from 2% to 4% – will introduce just the type of drift in inflation that can lead to random walk beliefs for standard policy coefficients in the Taylor rule. These random walk beliefs lead to a substantial overshooting of the inflation target. In addition, with a fixed target, occasional “unlikely” sequences of shocks can introduce drift to the inflation process that triggers random-walk beliefs and large deviations from the rational expectations equilibrium. Such departures from rational expectations can generate significant bursts of inflation, disinflation, and even deflation. In the New Keynesian model we find that such paths are more frequent with a higher level of

the inflation target.

In summary, in an imperfect information environment, increases in the inflation target, and also higher inflation targets, increase the chances of unstable inflation dynamics. We obtain these results for a wide range of Taylor rule parameters, and these findings also arise with some well-known forms of optimal monetary policy when there is imperfect knowledge of the inflation or price-level target. Unstable dynamics associated with increasing inflation targets can in principle be avoided by credibly and completely informing the private sector about the long-run inflation target *and* the timing of when that target will be implemented: we have seen that when agents know the mean inflation rate, and need only forecast its persistence, destabilizing inflation expectations do not arise. We are doubtful about success of this policy strategy for the reasons expressed by Ben Bernanke and by Barclays Capital in the quotations at the start of our paper. However, we have also seen that the risk of unstable dynamics is reduced if the inflation coefficient in the Taylor rule is increased substantially. Thus, an alternative way to view our results is that if a central bank does decide to increase its inflation target then it is essential to alter its interest rate rule, at least temporarily, to place a very high weight on achieving this target, e.g. by increasing the magnitude of the coefficient on inflation in the Taylor rule to values greatly in excess of those recommended by Taylor (1993) or found in most empirical studies.

Appendix

Proof to Propositions 1-2.

Propositions 1 and 2 provide asymptotic approximations to the learning algorithm

$$\begin{aligned}\theta_t &= \theta_{t-1} + \gamma S_{t-1} X_{t-1} (\pi_t - \theta'_{t-1} X_{t-1})' \\ S_t &= S_{t-1} + \gamma (X_t X_t' - S_{t-1})\end{aligned}$$

and where $\pi_t = T(\theta_{t-1})' X_{t-1} + \alpha^{-1} r_t$. It is possible to re-write the equations for real-time learning in the form

$$\phi_t^\gamma = \phi_{t-1}^\gamma + \gamma \mathcal{H}(\phi_{t-1}^\gamma, \bar{X}_t)$$

where $\bar{X}_t = (1, \pi_t, \pi_{t-1}, r_t)'$. Verifying many of the technical conditions required for convergence of the learning algorithm is simplified by the fact that the state dynamics, in a neighborhood of the equilibrium of interest, are conditionally linear and can be written as

$$\bar{X}_t \equiv \begin{bmatrix} X_t \\ X_{t-1} \\ r_t \end{bmatrix} = \begin{bmatrix} A(\phi_{t-1}) & 0 & 0 \\ I & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} \bar{X}_{t-1} + \begin{bmatrix} B & 0 \\ 0 & 0 \\ 0 & 1 \end{bmatrix} W_t$$

where $I, 0$ are conformable matrices, and

$$X_t = A(\phi_{t-1}) X_{t-1} + B W_t$$

Here $X_t = (1, \pi_t)'$ and $W_t = (1, r_t)'$ when dynamics are restricted to a neighborhood of the fundamentals REE or $W_t = (1, \nu_t)'$ when restricted to a set near the deflationary trap. The superscript γ highlights the dependence of the parameter estimates on γ . The stochastic approximation approach is to compare the solutions to the continuous time ODE and the discrete time algorithm, and then study the convergence of the continuous time approximating ODE. Thus, define the corresponding continuous time sequence for ϕ_t^γ as $\phi_t^\gamma = \phi_t^\gamma$ if $\tau_t^\gamma \leq \tau < \tau_{t+1}^\gamma$ where $\tau_t^\gamma = \gamma t$.

This Appendix sketches the proof to the propositions by making use of Propositions 7.8 and 7.9 of Evans and Honkapohja, and using arguments in Chapter 14 of Evans and Honkapohja and Branch and Evans (2011). The ‘‘mean dynamics’’ are the solution to the ODE

$$\frac{d\phi}{d\tau} = h(\phi)$$

where $h(\phi) = EH(\phi, \bar{X}_t)$. Notice, in particular, that this is the mean dynamics ODE given in the text:

$$\begin{aligned}\frac{d\theta}{d\tau} &= S^{-1} M(\phi) (T(\theta) - \theta) \\ \frac{dS}{d\tau} &= M(\phi) - S\end{aligned}$$

where $T(\theta) = (\alpha^{-1}(\alpha - 1)\bar{\pi} + \alpha^{-1}a(1 + b), \alpha^{-1}b^2)'$ in a subset of the fundamentals REE and $T(\theta) = (\tilde{\pi}, 0)'$ near the deflationary equilibrium.

Let $\tilde{\phi}(\tau, \phi_0)$ be the solution to the mean dynamics differential equation $\dot{\phi} = h(\phi)$ from an initial condition ϕ_0 . Define $U^\gamma(\tau) = \gamma^{-1/2} \left(\phi^\gamma(\tau) - \tilde{\phi}(\tau, \phi_0) \right)$. The two propositions in the text are based on U^γ converging to a Gaussian variable, in a sense made precise below. In particular, for small γ the probability distribution of $U^\gamma(\tau)$ converges to the probability distribution of the solution $U(t)$ to the differential equation

$$dU(\tau) = D_\phi h(\tilde{\phi}(\tau, \phi_0))U(\tau)d\tau + \mathcal{R}^{1/2}(\tilde{\phi}(\tau, \phi_0))dW(\tau)$$

The results below establish that $EU(\tau) = 0$ so that, as $\gamma \rightarrow 0$, $E\phi^\gamma(\tau) = \tilde{\phi}(\tau, \phi_0)$ and $\lim_{\tau \rightarrow \infty} \tilde{\phi}(\tau, \phi_0) = \phi^*$. Thus, key properties of the learning dynamics arise from a study of (i.) the asymptotic distribution for θ_t around the rational expectations equilibrium and (ii.) the mean dynamic path $\tilde{\phi}(\tau, \phi_0)$ where ϕ_0 are drawn from the asymptotic distribution.

The validity of the propositions in the text depend on verifying a set of technical conditions. The conditions required for Proposition 2 can be verified by using the arguments in Branch and Evans (2011), and so they are omitted here.

Proposition 2 uses the following result from Evans and Honkapohja (2001):

Proposition 3 (EH(2001)) *Consider the normalized random variables:*

$$U^\gamma(\tau) = \gamma^{-1/2} \left(\phi^\gamma(\tau) - \tilde{\phi}(\tau, \phi_0) \right).$$

As $\gamma \rightarrow 0$, the process $U^\gamma(\tau)$, $0 \leq \tau \leq T$, converges weakly to the solution $U(\tau)$ of the stochastic differential equation

$$dU(\tau) = D_\phi h(\tilde{\phi}(\tau, \phi_0))U(\tau)d\tau + \mathcal{R}^{1/2}(\tilde{\phi}(\tau, \phi_0))dW(\tau)$$

with initial condition $U(0) = 0$, where $W(\tau)$ is a standard vector Wiener process, and \mathcal{R} is a covariance matrix whose i, j th elements are

$$\mathcal{R}^{ij}(\phi) = \sum_{k=-\infty}^{\infty} Cov \left[\mathcal{H}^i(\phi, \bar{X}_k^\phi), \mathcal{H}^j(\phi, \bar{X}_0^\phi) \right]$$

Moreover, the solution to the stochastic differential equation has the following properties

$$\begin{aligned} EU(\tau) &= 0 \\ \frac{dVar(U(\tau))}{d\tau} &= D_\phi h(\tilde{\phi}(\tau, \phi_0))V_u(\tau) + V_u D_\phi h(\tilde{\phi}(\tau, \phi_0))' + \mathcal{R}(\tilde{\phi}(\tau, \phi_0)), \end{aligned}$$

where $V_u = \text{Var}(U(\tau))$. This result indicates that, for finite periods of time, the learning dynamics weakly converge to the solution of the ODE $\dot{\theta} = h(\theta)$, thus establishing Proposition 2.

Proposition 1 relies on the stochastic differential equation in the above result to have a stationary distribution asymptotically. Establishing this result requires stronger conditions. In particular,

A1 ϕ^* is a globally asymptotically stable rest point of the ODE $\dot{\phi} = h(\phi)$.

A2 $D_\phi h(\phi)$ is Lipschitz and all of the eigenvalues of $D_\phi h(\phi^*)$ have strictly negative real parts.

A3 There exist $q_1, q_2, q_3 \geq 0$ such that, for all $q > 0$ and all compact sets Q , there is a constant $\mu(q, Q)$ such that for all $x \in R^d, a \in Q$,

- i. $\sup_n E_{x,a}(1 + |\bar{X}_n|^q) \leq \mu(1 + |x|^q)$,
- ii. $\sup_n E_{x,a}(|H(\phi_n^\gamma, \bar{X}_{n+1})|^2) \leq \mu(1 + |x|^{q_1})$,
- iii. $\sup_n E_{x,a}(|\nu_{\phi_n^\gamma}(\bar{X}_{n+1})|^2) \leq \mu(1 + |x|^{q_2})$, where $\nu_\phi = \sum_{k \geq 0} (\Pi_\phi^k H_\phi - h(\phi))(y)$, and Π_ϕ is the stationary transition probability associated to the stationary Markov process \bar{X}_n ,
- iiii. $\sup_n E_{x,a}(|\phi_n^\gamma|^2) \leq \mu(1 + |x|^{q_3})$.

As noted in the text, there are three rest points to the ODE $\dot{\phi} = h(\phi)$, corresponding to the two REE with $b = 0$, i.e. $a = \bar{\pi}$ or $a = \tilde{\pi}$, and the other with $b = \alpha, a = \bar{\pi}$. The $b = \alpha$ REE is unstable under learning, and for some values of ϕ_t^γ the dynamics are explosive. For initial conditions sufficiently close to $b = 0$, and sufficiently small gain parameters γ , then the MSV REE is a stable rest point to the learning dynamics. However, to apply the approximation theorem below, the algorithm needs to rule out trajectories in the explosive region. Thus, the learning algorithm is supplemented with a “projection facility” that projects the iterates ϕ_t^γ into a confined set (see Evans and Honkapohja (2001) and Kushner and Yin (1997)). As a result of these assumptions the RE solution $(\bar{a}, \bar{b}) = (\bar{\pi}, 0)$ is a globally stable rest point of the ODE that satisfies (A1)-(A2).

It remains to verify (A3). Write $\bar{X}_n = \bar{A}(\phi_{n-1})\bar{X}_{n-1} + \bar{B}W_t$, where the expressions for \bar{A}, \bar{B} are given above. The eigenvalues of \bar{A} are zero and A , and the projection facility along with the conditional linearity ensures that \bar{X}_n remains in a compact subset of D , an open set around the REE $(\bar{\pi}, 0)$ or the REE $(\tilde{\pi}, 0)$, which in each case has a unique rest point to

$\dot{\phi} = h(\phi)$. Thus (A3.i) is immediate. Verifying conditions (A.ii)-(A.iv) is tedious, but given a projection facility that constrains ϕ_t to lie in a compact subset of D , it is straightforward to extend the arguments in Evans and Honkapohja (2001) (pg.335-336) for the Cobweb model to the present setting.

Proposition 1 arises from the following result in Evans and Honkapohja:

Proposition 4 (EH(2001)) *Consider the normalized random variables $U^{\gamma_k}(\tau) = \gamma_k^{-1/2} (\phi^{\gamma_k}(\tau) - \phi^*)$. For any sequences $\tau_k \rightarrow \infty, \gamma_k \rightarrow 0$, the sequence of random variables $(U^{\gamma_k}(\tau_k))_{k \geq 0}$ converges in distribution to a normal random variable with zero mean and covariance matrix*

$$C = \int_0^\infty e^{sB} \mathcal{R}(\theta^*) e^{sB'} ds,$$

where $B = D_\phi h(\phi^*)$.

It follows then that $\theta_t \sim N(\theta^*, \gamma C)$ for small γ and large t . Using arguments in Evans and Honkapohja (2001), Chapter 14.4, C is the solution to the matrix Riccati equation

$$D_\theta h(\phi^*)C + C(D_\theta h(\phi^*))' = -\mathcal{R}_\theta(\phi^*)$$

where $R = EH(\phi^*, \bar{X})H(\phi^*, \bar{X})'$. Straightforward calculations then lead to the expression for V in the text.

Overview of the New Keynesian Model with Trend Inflation.

The reduced-form equations (13)-(14) were derived by Ascari and Ropele (2007) from a standard New Keynesian framework and log-linearized around a non-zero steady-state inflation rate. This Appendix provides a brief overview of the model in Ascari and Ropele (2007).

There are a continuum of (identical) households whose flow utility is given by

$$U(C, N) = \frac{C_t^{1-\sigma}}{1-\sigma} - \chi N_t$$

Households maximize lifetime utility subject to the constraint,

$$P_t C_t + B_t \leq P_t w_t N_t + (1 + i_{t-1})B_{t-1} + \Pi_t + T_t$$

where P_t is the price of the final good, B_t are risk-free one period bonds with nominal net return i_{t-1} , Π_t are profits returned to households and T_t are lump-sum transfers. This formulation assumes the ‘‘cashless limit’’ that abstracts from money balances in the household’s problem. The household will select sequences of consumption, labor hours, and bond holdings to satisfy the first-order conditions

$$C_t^{-\sigma} = \beta \hat{E}_t \left(C_{t+1}^{-\sigma} (1 + i_t) \frac{P_t}{P_{t+1}} \right) \quad (23)$$

$$\chi C_t^\sigma = w_t \quad (24)$$

When $\hat{E} = E$, i.e. agents hold rational expectations, the conditions (23)-(24) have the usual interpretation.

When $\hat{E} \neq E$, the equation (23) can be justified in several ways. Our preference is to posit (23) as a behavioral relation that dictates that boundedly rational households choose their consumption to equate their expected marginal rate of substitution with the marginal rate of transformation. This is called Euler-equation learning and is another benchmark approach in the learning literature. See, for example Evans and Honkapohja (2013). In the current setting, the Euler-equation learning interpretation of (23) is completed by imagining a representative agent with a long history of data who observed a strong positive correlation between their own consumption and aggregate consumption. Euler-equation learning is a bounded-optimality approach closely related to the more general “shadow price” learning approach of Evans and McGough (2014).²⁸

The final good Y_t is produced by perfectly competitive firms using intermediate goods $Y_t(i)$ produced using a CES production function $Y_t = \left(\int_0^1 Y_t(i)^{(\zeta-1)/\zeta} di \right)^{\zeta/(\zeta-1)}$, $\zeta > 1$. The final goods firms choose their inputs to maximize profits, taking prices as given, resulting in the demand for input i $Y_t(i) = (P_t(i)/P_t)^{-\zeta} Y_t$. Intermediate goods are produced by a continuum of firms with technology $Y_t(i) = N_t(i)$. Intermediate goods producers take the demand for their good as given when setting prices optimally. However, they also face the Calvo risk where with probability α the firm’s price will remain unchanged each period. This leads to an expression for price setting that is identical to that of Woodford, except that the optimal re-set price also depends on the cumulative gross inflation rates over the period that a price might remain fixed.

Ascari and Ropele (2007) show that the steady-state properties depend on the trend inflation rate and, in particular, under most plausible parameterizations positive trend inflation leads to a lower steady-state output. Ascari and Ropele then demonstrate that a log-linearization, around a steady-state with gross inflation Π , of the equilibrium conditions

²⁸An alternative approach has been advanced by Preston (2006) in which boundedly rational agents solve their perceived dynamic programming problem, assuming that their beliefs will not change over time. This “anticipated utility” approach is typically implemented, e.g., by obtaining an IS equation that depends on expectations of interest rates and inflation over all future horizons. It would also be possible to develop Bayesian approaches that do not rely on a large exogenous component to consumption. However, as is evident from Cogley and Sargent (2008), this approach would be technically difficult to implement. We hypothesize that our results will be robust to these alternative non-RE approaches.

lead to the following reduced-form equations:

$$\begin{aligned}\hat{x}_t &= E_t \hat{x}_{t+1} - \sigma^{-1} (i_t - E_t \hat{\pi}_{t+1} - \hat{r}_t) \\ \hat{\pi}_t &= \kappa \hat{x}_t + \beta \Pi E_t \hat{\pi}_{t+1} + (\Pi - 1) \beta (1 - \alpha \Pi^{\zeta-1}) E_t \left((\zeta - 1) \hat{\pi}_{t+1} + \hat{\phi}_{t+1} \right), \\ \hat{\phi}_t &= (1 - \alpha \beta \Pi^{\zeta-1}) (1 - \sigma) \hat{x}_t + \alpha \beta \Pi^{\zeta-1} E_t \left((\zeta - 1) \hat{\pi}_{t+1} + \hat{\phi}_{t+1} \right)\end{aligned}$$

where $\hat{x}, \hat{\pi}, \hat{i}$ are log deviations from a steady-state with gross inflation factor Π . Iterating forward on the ϕ equation leads to the equations in the text. Ascari and Ropele show that $\kappa = (\Pi - 1)(\sigma - 1)\beta(1 - \alpha\Pi^{\zeta-1}) + \sigma\lambda(\Pi)$, $\lambda(\Pi) = (1 - \alpha\Pi^{\zeta-1})(1 - \alpha\beta\Pi^{\zeta})/\alpha\Pi^{\zeta-1}$.

By setting $\Pi = 1$, i.e. linearizing around a zero inflation steady-state, these equations reduce to the benchmark New Keynesian model

$$\begin{aligned}\hat{x}_t &= E_t \hat{x}_{t+1} - \sigma^{-1} (\hat{i}_t - E_t \hat{\pi}_{t+1} - r_t) \\ \hat{\pi}_t &= \beta E_t \hat{\pi}_{t+1} + \kappa \hat{x}_t\end{aligned}$$

References

- ASCARI, G., AND T. ROPELE (2007): “Optimal Monetary Policy Under Low Trend Inflation,” *Journal of Monetary Economics*, 54, 2568–2583.
- ASCARI, G., AND A. SBORDONE (2014): “The Macroeconomics of Trend Inflation,” *Journal of Economic Literature*, 52(3), 679–739.
- BALL, L. (2013): “The Case for Four Percent Inflation,” working paper.
- BRANCH, W., AND G. W. EVANS (2006): “A Simple Recursive Forecasting Model,” *Economic Letters*, 91, 158–166.
- BRANCH, W. A., AND G. W. EVANS (2011): “Learning about Risk and Return: A Simple Model of Bubbles and Crashes,” *American Economic Journal: Macroeconomics*, 3(3), 159–191.
- BULLARD, J., AND I.-K. CHO (2005): “Escapist Policy Rules,” *Journal of Economic Dynamics and Control*, 29, 1841–1866.
- CHO, I.-K., AND K. KASA (2008): “Learning Dynamics and Endogenous Currency Crises,” *Macroeconomic Dynamics*, 12, 257–285.
- CHO, I.-K., N. WILLIAMS, AND T. J. SARGENT (2002): “Escaping Nash Inflation,” *Review of Economic Studies*, 69, 1–40.

- COGLEY, T., C. MATTHES, AND A. SBORDONE (2015): “Optimized Taylor Rules for Disinflation When Agents are Learning,” *Journal of Monetary Economics*, 72(131-47).
- COGLEY, T., AND T. J. SARGENT (2008): “Anticipated Utility and Rational Expectations as Approximations of Bayesian Decision Making,” *International Economic Review*, 49.
- COGLEY, T., AND A. SBORDONE (2008): “Trend Inflation, Indexation, and Inflation Persistence in the New Keynesian Phillips Curve,” *American Economic Review*, 98.
- COIBION, O., Y. GORODNICHENKO, AND J. WIELAND (2012): “The Optimal Inflation Rate in New Keynesian Models,” *Review of Economic Studies*, 79, 1371–1406.
- EGGERTSSON, G. B. (2008): “Great Expectations and the End of the Depression,” *American Economic Review*, 98(4), 1476–1516.
- EGGERTSSON, G. B., AND M. WOODFORD (2003): “The Zero Interest-Rate Bound and Optimal Monetary Policy,” *Brookings Panel on Economic Activity*.
- EUSEPI, S., AND B. PRESTON (2010): “Central Bank Communication and Expectations Stabilization,” *American Economic Journal: Macroeconomics*, 2(3), 235–271.
- (2011): “Expectations, Learning and Business Cycle Fluctuations,” *American Economic Review*, 101, 2844–2872.
- EVANS, G. W. (2013): “The Stagnation Regime of the New Keynesian Model and Current US Policy,” in *Macroeconomics at the Service of Public Policy*, ed. by T. Sargent, and J. Vilmunen, chap. Chapter 3. Oxford University Press.
- EVANS, G. W., E. GUSE, AND S. HONKAPOHJA (2008): “Liquidity Traps, Learning and Stagnation,” *European Economic Review*, 52, 1438–1463.
- EVANS, G. W., AND S. HONKAPOHJA (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, New Jersey.
- (2003): “Expectations and the Stability Problem for Optimal Monetary Policies,” *Review of Economic Studies*, 70, 807–824.
- (2006): “Monetary Policy, Expectations and Commitment,” *Scandinavian Journal of Economics*, 108, 15–38.

- (2013): “Learning as a Rational Foundation for Macroeconomics and Finance,” in *Rethinking Expectations: The Way Forward for Macroeconomics*, ed. by R. Frydman, and E. S. Phelps. Princeton University Press.
- EVANS, G. W., S. HONKAPOHJA, AND N. WILLIAMS (2010): “Generalized Stochastic Gradient Learning,” *International Economic Review*, 51, 237–262.
- EVANS, G. W., AND B. MCGOUGH (2014): “Learning to Optimize,” Working paper.
- GERAATS, P. (2002): “Central Bank Transparency,” *Economic Journal*, 112, F532–F565.
- GURAYNAK, R., A. LEVIN, AND E. SWANSON (2010): “Does Inflation Targeting Anchor Long-Run Inflation Expectations? Evidence from the US, UK, and Sweden,” *Journal of the European Economic Association*, 8(6), 1208–1242.
- KOBAYASHI, T., AND I. MUTO (2013): “A Note on Expectational Stability under Non-zero Trend Inflation,” *Macroeconomic Dynamics*, 17, 681–693.
- KRUGMAN, P. R. (1998): “It’s Baaack: Japan’s Slump and the Return of the Liquidity Trap,” *Brookings Papers on Economic Activity*, (2), 137–205.
- KUSHNER, H. J., AND G. G. YIN (1997): *Stochastic Approximation Algorithms and Applications*. Springer-Verlag, Berlin.
- LANSING, K. (2009): “Time Varying U.S. Inflation Dynamics and the New Keynesian Phillips Curve,” *Review of Economic Dynamics*, 12(2), 304–326.
- MARCET, A., AND T. J. SARGENT (1989): “Convergence of Least-Squares Learning in Environments with Hidden State Variables and Private Information,” *Journal of Political Economy*, 97, 1306–1322.
- MCGOUGH, B. (2006): “Shocking Escapes,” *Economic Journal*, 116, 507–528.
- NUNES, R. (2009): “Learning the Inflation Target,” *Macroeconomic Dynamics*, 13, 167–188.
- ORPHANIDES, A., AND J. C. WILLIAMS (2005a): “The Decline of Activist Stabilization Policy: Natural Rate Misperceptions, Learning and Expectations,” *Journal of Economic Dynamics and Control*, 29, 1927–1950.
- (2005b): “Imperfect Knowledge, Inflation Expectations, and Monetary Policy,” in *The Inflation-Targeting Debate*, ed. by B. Bernanke, and M. Woodford, chap. 5, pp. 201–234. University of Chicago Press.

- (2007): “Robust Monetary Policy with Imperfect Knowledge,” *Journal of Monetary Economics*, 54, 1406–1435.
- PRESTON, B. (2006): “Adaptive Learning, Forecast-based Instrument Rules and Monetary Policy,” *Journal of Monetary Economics*, 53, 507–535.
- RUDEBUSCH, G. (2001): “Is the Fed Too Timid? Monetary Policy in an Uncertain World,” *Review of Economics and Statistics*, 83(2), 203–217.
- SARGENT, T. J. (1999): *The Conquest of American Inflation*. Princeton University Press, Princeton NJ.
- SARGENT, T. J., AND N. WILLIAMS (2005): “Impacts of Priors on Convergence and Escapes from Nash Inflation,” *Review of Economic Dynamics*, 8, 360–391.
- SCHMITT-GROHE, S., AND M. URIBE (2011): “The Optimal Rate of Inflation,” in *Handbook of Monetary Economics*, ed. by B. M. Friedman, and M. Woodford, vol. 3B, pp. 653–722. Elsevier.
- SMETS, F., AND R. WOUTERS (2007): “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach,” *American Economic Review*, 97(3), 586–606.
- SUMMERS, L. H. (1991): “Panel Discussion: Price Stability, how Should Long-Term Monetary Policy be Determined?,” *Journal of Money, Credit and Banking*, 23, 625–631.
- TAYLOR, J. (1993): “Discretion versus Policy Rules in Practice,” *Carnegie-Rochester Conference Series in Public Policy*, 39, 195–214.
- VESTIN, D. (2006): “Price-level versus Inflation Targeting,” *Journal of Monetary Economics*, 53, 1361–1376.
- WILLIAMS, J. C. (2006): “Monetary Policy in a Low Inflation Economy with Learning,” in *Monetary Policy in an Environment of Low Inflation: Proceedings of the Bank of Korea International Conference 2006*, pp. 199–228. Bank of Korea.
- WILLIAMS, J. C., AND D. REIFSCHNEIDER (2000): “Three Lessons for Monetary Policy in a Low Inflation Era,” *Journal of Money, Credit and Banking*, pp. 936–966.
- WILLIAMS, N. (2004): “Escape Dynamics in Learning Models,” Discussion paper, working paper, Princeton University.

WOODFORD, M. (2003): *Interest and Prices: Foundations of a Theory of Monetary Policy*.
Princeton University Press, Princeton, NJ.

Table 1: Calibration. Note: For definitions of parameters, see Appendix.

β	0.995
ζ	10
α_π	1.5
α_x	0.125
α	0.67
σ	1.00
χ	1.00
σ_r^2	0.1
σ_u^2	0.003

Figure 1: T-map dynamics in the Fisherian model with $\alpha = 1.1$, $\bar{\pi} = 4\%$.




Figure 2: Confidence Ellipses around REE for constant gain learning. Each ellipse corresponds to a different inflation target. The targets are .5%, 1%, 2%, 3%, 4%, 5%, expressed in annualized rates.




Figure 3: Mean Dynamics in the Fisherian Model. Initial conditions are drawn from the confidence ellipse.




Figure 4: Change in inflation target from 2% to 4%. Economy begins in the REE. $\alpha = 1.1, \sigma_r^2 = .003, \gamma = .02$.




Figure 5: Fisherian inflation dynamics with a 4% target. Left panel plots an inflationary episode, right panel plots a disinflationary episode.

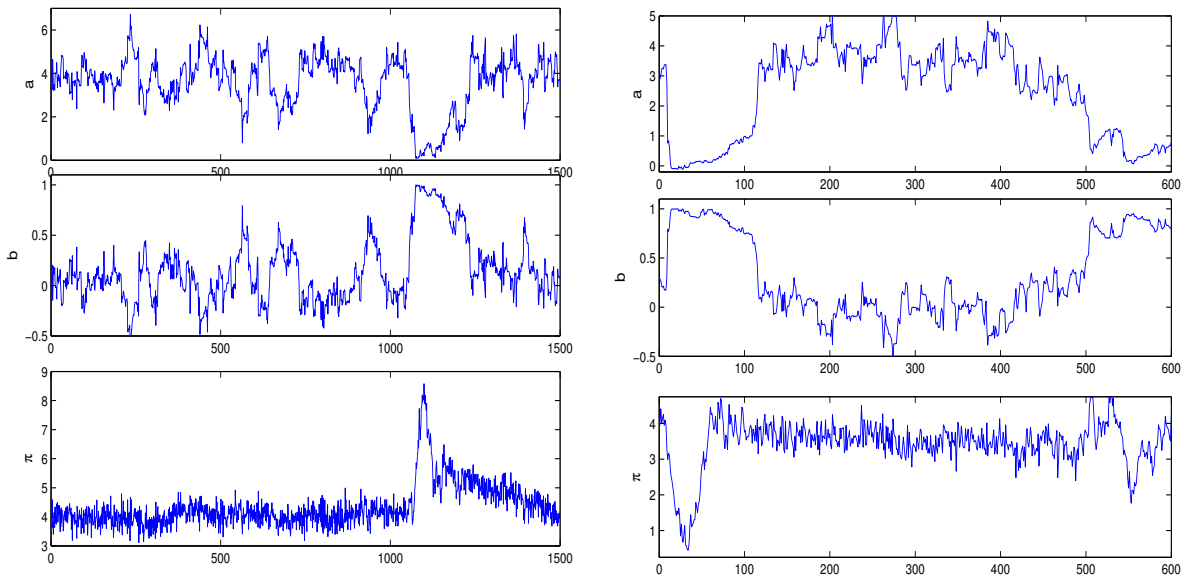


Figure 6: Alternative implementations of new target. Left panel implements target incrementally over horizons of 1,8,16,32,64,128,256 quarters. Right panel initially sets a higher inflation target then lowers the target each quarter until reaching the desired 4% target.

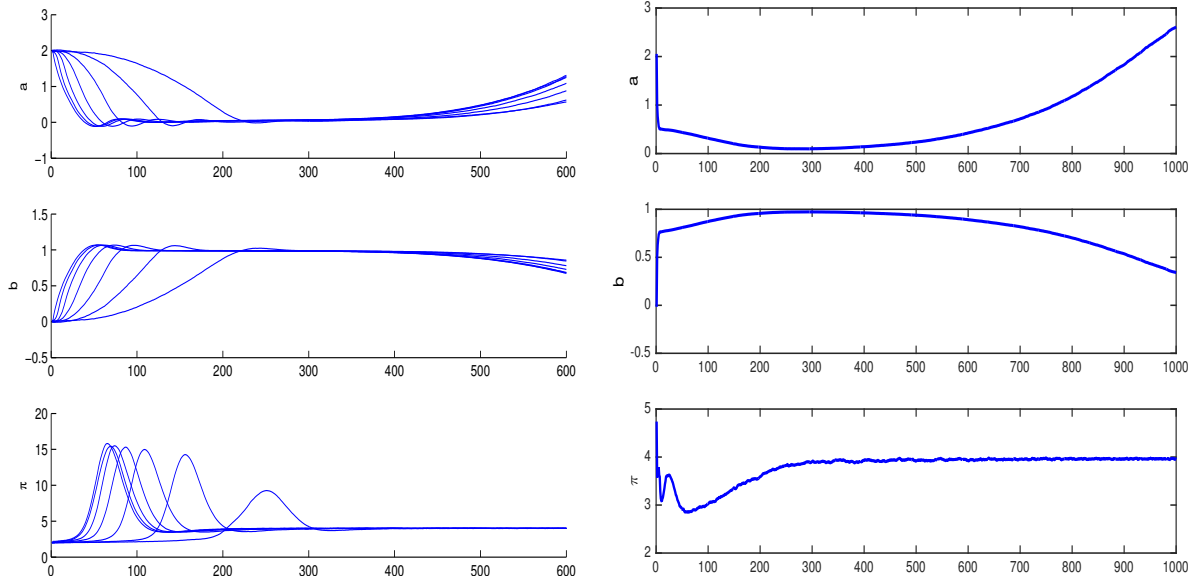


Figure 7: Confidence Ellipse for 2% and 4% targets. Computed via monte-carlo simulations assuming an asymptotic normal distribution for the coefficients. The coefficients are centered around $(0,0)$ by de-meaning the simulated series.




Figure 8: Mean dynamics in NK Model with a 4% target.




Figure 9: Mean dynamics in NK Model with alternative inflation targets.




Figure 10: Random-walk beliefs for various inflation targets. Figure plots the largest root of agents' VAR matrix B in the New Keynesian model with the baseline calibration.




Figure 11: Increasing the Inflation target from 2% to 4%.




Figure 12: Inflation dynamics in NK Model with a 4% target.

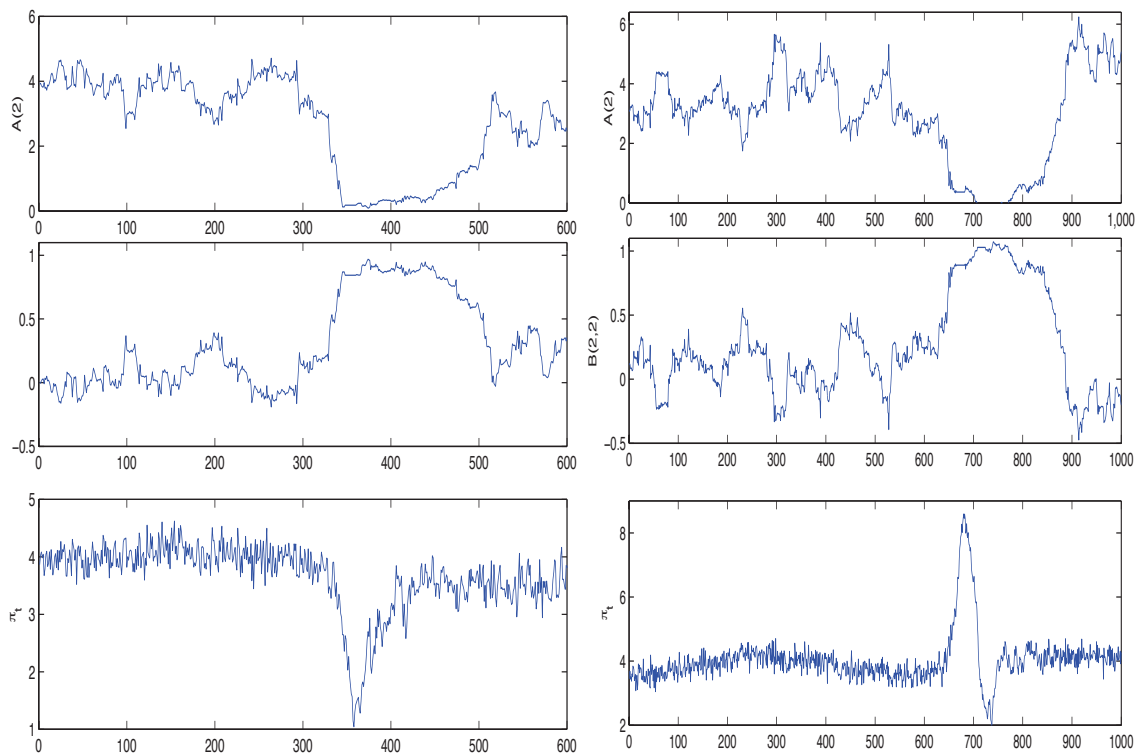


Figure 13: A Policy Frontier for Random-walk beliefs. All pairs of policy coefficients (α_x, α_π) below the line can lead to random-walk beliefs in the mean dynamics for the NK model with the baseline calibration, and policy coefficients above the line do not exhibit random-walk beliefs. Circle corresponds to the Taylor rule.




Figure 14: Inflation dynamics in NK Model with a 4% target and a zero lower bound.




Figure 15: Policy target communication.




Figure 16: Increasing the target from 2% to 4% and increasing α_π from 1.5 to 3.0.




Figure 17: Mean dynamics in NK Model with a 4% target and optimal policy without commitment.




Figure 18: Mean dynamics in NK Model with a price-level targeting rule.

