Inflation Announcements and Social Dynamics*

Kinda Hachem † Jing Cynthia Wu ‡

Chicago Booth Chicago Booth

First Version: August 28, 2012

This Version: August 16, 2013

Abstract

We investigate the effectiveness of central bank communication when firms have heterogeneous inflation expectations that are updated through social dynamics. The bank's credibility evolves with these dynamics and determines how well its announcements anchor expectations. We show that trying to eliminate high inflation by introducing a low inflation target can lead to short-term overshooting if the introduction is insufficiently gradual. In contrast, combating deflation requires either aggressive announcements that are broadly consistent with price level targeting or QE-type announcements that allow the central bank to stem deflationary expectations without altering its inflation target.

Keywords: central bank communication, expectations heterogeneity, social dynamics, credibility, inflation targeting, quantitative easing

*We thank Veronica Guerrieri, Anil Kashyap, and Randy Kroszner for helpful suggestions. We have also benefited from discussions with Jon S. Cohen, Steve Davis, Jim Hamilton, Narayana Kocherlakota, Jim Pesando, Ricardo Reis, Michael Woodford, and seminar participants at the Federal Reserve Board, Chicago Booth, UBC, FRB New York, FRB Chicago, SED 2013, and the 2013 NBER Summer Institute EFSF Workshop. Both authors gratefully acknowledge financial support from the University of Chicago Booth School of Business. Cynthia Wu also gratefully acknowledges financial support from the IBM Faculty Research Fund at the University of Chicago Booth School of Business.

†The University of Chicago Booth School of Business, 5807 South Woodlawn Avenue, Chicago, IL 60637, USA, Kinda.Hachem@chicagobooth.edu

‡The University of Chicago Booth School of Business, 5807 South Woodlawn Avenue, Chicago, IL 60637, USA, Cynthia.Wu@chicagobooth.edu

1 Introduction

For many central banks, communication has become part of the policy toolkit. Inflation targeters provide a prime example as they rely on clear and transparent communication to anchor inflation expectations over various horizons (e.g. Carney (2012)). The rise of communication has also been visible since December 2008 when the Federal Reserve hit the zero lower bound and unveiled a series of unconventional programs to shore up the financial system and stem deflationary expectations. In this paper, we study the effectiveness of central bank communication when agents have heterogeneous inflation expectations that evolve through social dynamics.

While communication has gained attention among macroeconomists, many of the key insights depend on public uncertainty about the central bank's current or future actions (e.g. Melosi (2012) and Eggertsson and Pugsley (2006)). A similar dependence exists in game theoretic work, with asymmetric information between the public and the central bank used to explain the pre-Greenspan Fed's preference for ambiguity (e.g. Stein (1989) and Cukierman and Meltzer (1986)). Given the recent shift to transparency discussed in Woodford (2005) and Blinder et al. (2008), however, public information about policy goals may now be the relevant baseline. Interestingly though, increased transparency has not entirely eliminated heterogeneity in inflation expectations (e.g. Mankiw, Reis and Wolfers (2004)). Why does disagreement about future inflation persist despite clear announcements by the central bank? How can announcements be designed to achieve maximal anchoring of expectations? To answer such questions, one needs a model which speaks to the expectations formation process (e.g. Kroszner (2012), Boivin (2011), and King (2005)).

Our paper takes a step in this direction. We construct a simple model of inflation determination where monopolistically competitive firms must make decisions before the aggregate price level is known and thus rely on inflation forecasts. Empirical evidence points to two natural forecasting rules: one that is consistent with central bank announcements and one that is consistent with a random walk. We set up our model so that each rule is indeed an unbiased forecast of inflation when adopted by all firms. For example, if all firms use the central bank's announcements as a basis for forecasting (i.e., if the bank is highly credible), then firm decisions are such that the announcement is in fact realized. The opposite is true if the bank is not credible. The fraction of firms with announcement-consistent forecasts is thus a crucial variable in our model and we endogenize it using social dynamics. In particular, once inflation has been realized, firms can meet and potentially switch forecasting rules based on relative performance. A small and/or temporary divergence of realized inflation from the central bank's announcements may not have enough momentum to significantly affect credibility. However, prolonged divergence may convince some firms to abandon the central bank's cues in favor of more successful forecasting rules, limiting the extent to which future announcements will be realized. Combining our model of inflation determination with our model of social dynamics, we investigate how announcements can be tailored to limit divergence and build credibility.

Our analysis yields three main insights. First, we show that abruptly introducing a low inflation target to achieve a large disinflation can cause temporary overshooting of the target, even when the central bank is transparent and firms reset prices every period. In contrast, gradually introducing the target (i.e., via interim targets) directs the economy to the long-term goal more smoothly because the interim targets provide more scope for credibility-building when beliefs evolve through social dynamics. Our model thus suggests a novel explanation for overshooting differences among the inflation targeters studied in Mishkin and Schmidt-Hebbel (2007). Second, we show that gradual strategies are actually less effective in a deflation. Instead, the central bank can eliminate deflation more quickly by communicating an aggressive increase in its short-term inflation goals. Our results thus suggest that price-level targeting may have some communication-based benefits over inflation targeting during a deflation. Lastly, we show how two dimensions of quantitative easing - number of rounds and intensity of announcements - can be varied to guide the economy out of deflation without explicitly increasing short-term targets. While many papers have focused on the yield curve response to QE, our model sheds some light on the less pervasive but potentially key inflation response defined in Krishnamurthy and Vissing-Jorgensen (2011).[1]

Since our results are driven by the interaction between inflation determination and social dynamics, they are difficult to generate if expectations are homogeneous and rational as assumed in workhorse models of monetary policy.[2] The use of rule-based agents to bridge the gap between tractability and realism has recently gained attention in economic modeling, with Ellison and Fudenberg (1993) showing that even naive rules-of-thumb can achieve fairly efficient outcomes. Further work has also demonstrated how social dynamics between heterogeneous agents can change the predictions of more standard models (e.g. Arifovic, Bullard and Kostyshyna (2012)) and/or explain otherwise puzzling aggregate dynamics (e.g. Burnside, Eichenbaum and Rebelo (2013)).[3] Although there is a large literature on representative learning of central bank goals - see, for example, Orphanides and Williams (2005), Berardi and Duffy (2007), Eusepi and Preston (2010), and Branch and Evans (2011) - we are not aware of any papers that have introduced social dynamics into a model of inflation determination to endogenize the credibility of transparent communication. In this regard, we also differ from Arifovic et al. (2010) who allow the central bank to choose both inflation announcements and realized inflation in a cheap talk economy with social learning.

The rest of the paper proceeds as follows. Section 2 overviews our framework to highlight the key interactions, Section 3 explains the evolution of credibility through social dynamics, Section 4 builds a model of inflation determination for use in simulations, Sections 5 and 6 present the simulation results, and Section 7 concludes. All proofs appear in the Appendix.

[1] For more on the yield curve effect, see Gagnon et al. (2010), D'Amico and King (2010), Williams (2011), Hamilton and Wu (2012), and the references therein.

[2] See, for example, Clarida, Gali and Gertler (1999), Woodford (2003), Smets and Wouters (2003) and Christiano, Eichenbaum and Evans (2005).

[3] For more on agent-based models, see LeBaron (2001), Judd and Tesfatsion (2006), Colander et al. (2008), Ashraf and Howitt (2008), and Page (2012).

2 Basic Framework

We begin by describing how announcements, expectations, and inflation will come together in this paper. There are three important features. First, how is inflation determined given a set of inflation expectations? Second, how do central bank announcements affect expectations? Third, how does realized inflation affect expectations? We take each question in turn. The focus here will be on intuition, with more formal structure deferred until Sections 3 and 4.

Inflation Determination Consider a continuum of firms $i \in [0, 1]$. At the beginning of date t, firm i expects an inflation rate of $\widehat{\pi}_t^i$. The mean expectation is g and the standard deviation across all firms is σ. We are primarily interested in how σ affects the actual inflation rate realized at the end of date t. Our full model of inflation determination is developed in Section 4 so the goal at present is to distill only the key forces that will emerge. On one hand, we will have a Jensen's inequality effect which is fairly standard in the finance literature.[4] As σ increases, compounding of inflation expectations into price expectations skews the distribution of price expectations right and, all else constant, drives up market clearing prices. In finance, the market clearing price is typically a nominal bond price. In our model, it will be a nominal input price. On the other hand, as the nominal input price increases, low expectation firms in our environment have less incentive to operate (i.e., all else is not constant). The Jensen's effect, $J(\sigma)$, thus triggers a countervailing exit effect, $Q(\sigma)$, where $J(0) = Q(0) = 0$, $J'(\cdot) > 0$, and $Q'(\cdot) > 0$. With a natural upper bound on exit though, the Jensen's effect eventually outpaces the exit effect and yields the following:

Lemma 1 *If there exists a $\sigma_A > 0$ such that $Q(\sigma_A) = J(\sigma_A)$ and $Q'(\sigma_A) > J'(\sigma_A)$, then there exists a $\sigma_B > \sigma_A$ such that $Q(\sigma_B) = J(\sigma_B)$ and $Q(\sigma) > J(\sigma)$ for all $\sigma \in (\sigma_A, \sigma_B)$.*

Provided there is a positive relationship between input and output prices, the mapping of inflation expectations into realized inflation is $\pi_t^* \approx g + \lambda(\sigma)[J(\sigma) - Q(\sigma)]$, where $\lambda(\cdot) > 0$.

[4] See, for example, Veronesi and Yared (2000) and Xiong and Yan (2010).

Central Bank Announcements To introduce announcements into this environment, suppose the central bank releases an inflation forecast of $\overline{\pi}_t$ at the beginning of date t. Under Lemma 1, we can consider two forecasting rules for firms: Rule A which draws $\widehat{\pi}_t^i$ from $N(\overline{\pi}_t, \sigma_A)$ and Rule B which draws $\widehat{\pi}_t^i$ from $N(g_t, \sigma_B)$. We discuss these rules more formally in Section 5.1. For now, we only emphasize that Rule A forecasts according to the central bank's cues while Rule B represents an alternative forecasting rule.

Let $\xi_t \in [0,1]$ denote the fraction of firms that use Rule A. Lemma 1 implies $\pi_t^* \approx \overline{\pi}_t$ if $\xi_t = 1$ and $\pi_t^* \approx g_t$ if $\xi_t = 0$. Both rules thus provide unbiased forecasts of inflation if adopted by all firms but only Rule A achieves the central bank's announcement when $g_t \neq \overline{\pi}_t$. What happens when $g_t = \overline{\pi}_t$? In this case, the standard deviation of expectations across all firms is $\sigma(\xi_t) = \sqrt{\xi_t \sigma_A^2 + (1 - \xi_t)\sigma_B^2}$ so Lemma 1 implies $\pi_t^* < \overline{\pi}_t$ if $\xi_t \in (0,1)$. For any g_t then, $\xi_t = 1$ will allow the central bank to achieve inflation goals using only communication. As such, we can also interpret ξ_t as a measure of central bank credibility.

Conditional on the forecasting rules $N(\overline{\pi}_t, \sigma_A^2)$ and $N(g_t, \sigma_B^2)$, it is now apparent that the pass-through from an inflation announcement $\overline{\pi}_t$ to realized inflation π_t^* will hinge on ξ_t. Completing our framework thus requires modeling the evolution of ξ_t.

Evolution of Expectations We posit that agents whose forecasts are consistently outperformed by their peers will want to change how they forecast. The question then becomes how an agent discovers he is being outperformed and how much of this outperformance he attributes to one-time shocks rather than fundamentals. Social dynamics provide a simple yet powerful way to address these questions and endogenize the evolution of ξ_t.

The details are presented in Section 3. What we emphasize here is that ξ_t is endogenously determined through social dynamics. That ξ_t is endogenous rather than deterministic is fundamental. As we saw above, $\pi_t^* \approx \overline{\pi}_t$ if $\xi_t = 1$. However, to reach high ξ_t endogenously, firms must be convinced that the central bank's announcements provide a good basis for forecasting. The announcements will indeed provide a good basis if ξ has been high historically. Therefore, there is a two-way feedback between selection of forecasting rules and realized inflation which both disciplines how central bank announcements steer inflation and introduces non-trivial inflation dynamics into the model.

3 Social Dynamics

As described above, we adopt social dynamics to endogenize the fraction of firms who use central bank announcements as a basis for their forecasts. To this end, we initialize $\xi_1 = 0$ then let ξ_{t+1} evolve via tournament selection and mutation.

We use tournament selection to simulate the transmission of information in a complex world. Versions of this approach appear in Carroll (2003a), Arifovic et al. (2010), Arifovic, Bullard and Kostyshyna (2012), and Burnside, Eichenbaum and Rebelo (2013). Relative to other studies, our tournaments (1) favor agents who are endogenously more successful and (2) permit success to be judged over multiple observations. In particular, firms meet in pairs and compare forecast errors after the realization of π_t^*. Denote firm i's forecasting rule by $Rule_t^i \in \{A, B\}$ and consider a meeting between i and j. Firm i counts one strike against his rule if $Rule_t^i \neq Rule_t^j$ and $\left|\widehat{\pi}_t^i - \pi_t^*\right| > \left|\widehat{\pi}_t^j - \pi_t^*\right|$. Recall from Section 2 that π_t^* depends on ξ_t. For example, if ξ_t is too low, then π_t^* will be close to g_t so $|g_t - \overline{\pi}_t| \gg 0$ will allow Rule B to outperform Rule A in many meetings. Strikes will thus tend to be counted against Rule A, suggesting $\xi_{t+1} \leq \xi_t$. In contrast, if ξ_t is sufficiently high, then π_t^* will be close to $\overline{\pi}_t$ and strikes will tend to be counted against Rule B, suggesting $\xi_{t+1} \geq \xi_t$. This is the sense in which success is endogenous.

Whether strikes actually lead to $\xi_{t+1} \neq \xi_t$ depends on how stubborn firms are in their beliefs. Experimental evidence suggests that people are very reluctant to contradict their own information, even when Bayesian updating suggests they should (e.g., Weizsacker (2010) and Andreoni and Mylovanov (2012)). We thus allow firms to accumulate several strikes before deciding to switch forecasting rules. This is the sense in which success is judged over multiple observations. Going forward, we use S to denote the number of strikes needed for a switch (i.e., after S strikes, firm i switches rules and begins counting strikes against his new rule). We also refer to S as stubbornness, with higher S implying more stubborn beliefs. To gauge the importance of S, we will simulate our model for different values. As Section 5.3 will show, the extent of stubbornness is an important input into our social dynamics.

Strikes accumulate across meetings and periods so we must now specify how pairwise meetings come about. In our baseline specification, pairs are drawn randomly with replacement from the entire population. Drawing with replacement ensures that each firm can have zero to many meetings in a given period. Drawing from the full population ensures that even firms who do not operate are represented in tournaments. This is appealing since operation decisions are driven by expectations. In an alternative specification, we allow tournaments to occur within neighborhoods rather than at random. More specifically, imagine that firms lie along a circle and each firm meets its right and left neighbors every period. With interactions set up as such, firms always meet the same people. As we will see in Section 5, this will create clusters of firms that use the same forecasting rule. Firms at the center of a cluster are thus more likely to meet other firms using the same rule, increasing their *effective* stubbornness for any value of S. Notice that this suggests an alternative interpretation for our S: higher values of S are a stand-in for more localized interactions.

Lastly, to capture the fact that some changes may not be performance-driven, we incorporate mutations: at the beginning of date $t+1$, a very small fraction $\mu \in (0,1)$ of firms randomly switches rules regardless of strikes.

The timing of our social forces can be summarized as follows: (i) mutation turns the fraction of Rule A forecasters into $\widetilde{\xi}_t = (1-\mu)\,\xi_t + \mu\,(1-\xi_t)$ if $t \geq 2$; (ii) each firm i draws expectation $\widehat{\pi}_t^i$ from its forecasting rule; (iii) the set of expectations $\left\{\widehat{\pi}_t^i \mid i \in [0,1]\right\}$ determines π_t^* as per the model developed next in Section 4; (iv) tournament selection transforms ξ_t into ξ_{t+1} if $t = 1$ and $\widetilde{\xi}_t$ into ξ_{t+1} if $t \geq 2$.

4 Economic Model

While instructive, our discussion of inflation determination in Section 2 was highly stylized. We now develop a more formal model. Subsection 4.1 presents the environment. Subsections 4.2 and 4.3 then flesh out the details to derive the specific functions we use in our simulations.

4.1 Environment

Consider a continuum of firms, each producing a differentiated perishable good $i \in [0,1]$. Demand for good i is $D\left(\frac{\gamma_t P_t}{p_{it}}\right)$, where p_{it} is the price charged by firm i at date t, P_t is the aggregate price level, and $\gamma_t \in [1-\varepsilon, 1+\varepsilon]$ is an exogenous and independently distributed aggregate taste shock. Assume $D'(\cdot) > 0$ so that the demand for good i is decreasing in the relative price $\frac{p_{it}}{P_t}$ and increasing in the taste shock. Supply of good i is chosen by firm i which possesses a technology to transform inputs into output. The technology is summarized by $F(\ell_{it}, \cdot)$, where ℓ_{it} denotes the labor input used by firm i. The aggregate stock of labor is normalized to one and inelastically available at unit wage w_t.

Firms have to make pricing and production decisions before γ_t and P_t are realized - that is, before they know the actual demand for their products. At the beginning of date t, all firms forecast a taste shock of one. Each firm i also forecasts an aggregate price level of $\widehat{P}_t^i \equiv \exp(\widehat{\pi}_t^i)P_{t-1}^*$, where P_{t-1}^* is last period's realized price level and $\widehat{\pi}_t^i$ is the firm's inflation expectation for the current period. All firms make decisions for the current period based on their forecasts. We will return to the forecasting rules for $\widehat{\pi}_t^i$ shortly. For now, we only note that they constitute a deviation from rational expectations (RE) since $\widehat{\pi}_t^i$ is not updated whenever new information becomes available.

The timeline after $\widehat{P}_t^i$ has been set is as follows. First, the firm decides on a labor demand function $\ell\left(w_t; \widehat{P}_t^i\right)$ and an individual pricing function $p\left(w_t; \widehat{P}_t^i\right)$. Second, the wage is set by an auctioneer to clear the labor market. In particular, the auctioneer sets w_t^* to solve $\int \ell\left(w_t^*; \widehat{P}_t^i\right) di = 1$. Third, if $\ell\left(w_t^*; \widehat{P}_t^i\right) > 0$, then firm i posts price $p_{it}^* \equiv p\left(w_t^*; \widehat{P}_t^i\right)$ and uses inputs to produce its expected demand $y^*_{it} \equiv D\left(\frac{\widehat{P}^i_t}{p^*_{it}}\right)$. Fourth, the taste shock γ_t is realized and the aggregate price level is computed as a consumption-weighted average of individual prices. At aggregate price P_t, realized demand for good i is $D\left(\frac{\gamma_t P_t}{p^*_{it}}\right)$ which may differ from the available supply y^*_{it}. Consumption is thus the minimum of demand and supply so the auctioneer computes P^*_t to solve $P^*_t = \int \frac{c_{it}}{\int c_{jt} dj} p^*_{it} di$ and $c_{it} = \min\left\{y^*_{it}, D\left(\frac{\gamma_t P^*_t}{p^*_{it}}\right)\right\}$.[5]

Putting everything together, we now have the following mapping from a set of expectations $\left\{\widehat{P}^i_t \mid i \in [0,1]\right\}$ to the realized price level P^*_t:

$$\int_{O_t(w^*_t)} \ell\left(w^*_t; \widehat{P}^i_t\right) di = 1 \;\; where \;\; O_t\left(w^*_t\right) \equiv \left\{i \mid \ell\left(w^*_t; \widehat{P}^i_t\right) > 0\right\} \tag{1}$$

$$P^*_t \int_{O_t(w^*_t)} D\left(\frac{\min\left\{\widehat{P}^i_t, \gamma_t P^*_t\right\}}{p\left(w^*_t; \widehat{P}^i_t\right)}\right) di = \int_{O_t(w^*_t)} D\left(\frac{\min\left\{\widehat{P}^i_t, \gamma_t P^*_t\right\}}{p\left(w^*_t; \widehat{P}^i_t\right)}\right) p\left(w^*_t; \widehat{P}^i_t\right) di \tag{2}$$

Realized inflation is defined as $\pi^*_t \equiv \ln\left(P^*_t / P^*_{t-1}\right)$. The rest of this section puts structure on $D(\cdot)$, $p(\cdot)$, and $\ell(\cdot)$ to refine the mapping in equations (1) and (2).[6]

4.2 Additional Structure

We will assume $D\left(\frac{\gamma_t P_t}{p_{it}}\right) = \left(\frac{\gamma_t P_t}{p_{it}}\right)^{\frac{1}{1-\rho}}$ with $\rho \in (0,1)$ and derive $p(\cdot)$ and $\ell(\cdot)$ from a static profit maximization problem. The following proposition disciplines our approach:

Proposition 1 *If $p(\cdot)$ and $\ell(\cdot)$ maximize $p(\cdot) D\left(\frac{\widehat{P}^i_t}{p(\cdot)}\right) - w_t \ell(\cdot)$, then $F(\ell, \cdot) = \ell$ implies:*

1. *$p\left(w^*_t; \widehat{P}^i_t\right)$ is independent of $\widehat{P}^i_t$*

2. *If $\widehat{\pi}^i_t \sim N\left(g, \sigma^2\right)$ and $\ell\left(w^*_t; \widehat{P}^i_t\right) > 0$ for all i, then $\pi^*_t = g + \frac{\sigma^2}{2(1-\rho)}$*

[5] While our results are robust to different consumption aggregators, we use $\int c_{jt} dj$ to ensure that our consumption weights sum to one. Weights computed using a CES aggregator only sum to one if consumption is homogeneous across goods, a condition which does not hold in our framework.

[6] As equations (1) and (2) show, our model is one where actual inflation is determined entirely by expected inflation. In other words, the central bank can only change inflation by changing expectations. While our abstraction from conventional policy tools is done to isolate the effect of communication, current work by Campbell (2013) demonstrates that it may in fact be optimal for policymakers to rely on open mouth operations, even when open market operations are available.

The first part of Proposition 1 says that a linear one-to-one production technology delivers homogeneous price-setting behavior among firms, regardless of expectations. The second part says that it also delivers realized inflation above the mean expectation when expectations are heterogeneous and all firms produce. The latter is the Jensen's inequality result discussed in Section 2. It is exacerbated by more heterogeneity (i.e., higher σ) and more substitutability between goods (i.e., higher ρ). As σ increases, compounding of inflation expectations into price expectations skews the distribution of price expectations further right. The highest expectation firms thus drive input costs and prices up more than the lowest expectation firms drive them down. The effect is strongest when goods are more substitutable because high expectation firms foresee a huge increase in sales by undercutting the aggregate price level and thus participate more actively in the labor market.

As per Woodford (2013), "it is appealing to assume that people's beliefs should be rational, in the ordinary-language sense, though there is a large step from this to the RE hypothesis." Together with the results of Proposition 1, this prompts us to adopt $F(\ell, \cdot) \neq \ell$ and allow $\ell\left(w_t^*; \widehat{P}_t^i\right) = 0$ for some i. Why is $F(\ell, \cdot) \neq \ell$ desirable? Our taste shock is realized after the auctioneer computes w_t^* so, as long as $\frac{\partial p\left(w_t^*; \widehat{P}_t^i\right)}{\partial \widehat{P}_t^i} \neq 0$, the auctioneer cannot also compute P_t^* before firms post prices and undertake production. While our firms deviate from the RE hypothesis by not updating $\widehat{P}_t^i$ based on w_t^*, they would also be deviating from rationality in the ordinary-language sense if they did not update $\widehat{P}_t^i$ based on P_t^*. We thus adopt $F(\ell, \cdot) \neq \ell$ to generate $\frac{\partial p\left(w_t^*; \widehat{P}_t^i\right)}{\partial \widehat{P}_t^i} \neq 0$ and delay the revelation of P_t^*. Similar reasoning motivates our allowance of $\ell\left(w_t^*; \widehat{P}_t^i\right) = 0$. If all firms use the same general forecasting rule, namely $\pi_t^i = g + \varepsilon_{it}$ with $\varepsilon_{it} \sim N(0, \sigma^2)$, then $\pi_t^* = g$ (henceforth mean rationality) is desirable. In other words, the economy should not converge to a situation where all firms use a rule that is always wrong on average. Mathematically, this requires overcoming the σ^2 term generated by Jensen's inequality. Intuitively, it requires giving low expectation firms more pull to overcome the pull that compounding gives high expectation firms. Allowing low expectation firms to not produce is a natural step in this direction and we pursue it below.

4.3 The Full Model

Suppose good i is produced according to $F(\ell_{it}, z_{it}) = \ell_{it}^{\alpha} z_{it}^{1-\alpha}$, where ℓ_{it} is again labor, z_{it} is firm effort, and $\alpha \in (0,1)$. All firms also have a real outside option U, thus allowing for exit and (re)entry. Effort is exerted by the firm and imparts (real) disutility $\frac{z_{it}^{\theta}}{\theta}$, where $\theta > 1$. If θ is finite, then production exhibits diminishing returns to labor but constant returns overall. If θ is infinite, then effort is always unity and the production function is effectively decreasing returns to scale. Our numerical exercises will employ the limiting case of $\theta \to \infty$ to simplify the parameter space. However, we will derive the key properties of our model for any $\theta > 1$ to show that they do not hinge on decreasing returns to scale.

If firm i charges p_{it}, its anticipated demand is $\left(\frac{\widehat{P}_t^i}{p_{it}}\right)^{\frac{1}{1-\rho}}$ and it will need $\left(\frac{\widehat{P}_t^i}{p_{it}}\right)^{\frac{1}{\alpha(1-\rho)}} z_{it}^{1-\frac{1}{\alpha}}$ units of labor to produce this quantity. Therefore, in real terms, the firm solves:

$$\max\left\{\max_{p_{it}, z_{it}}\left[\frac{p_{it}}{\widehat{P}_t^i}\left(\frac{\widehat{P}_t^i}{p_{it}}\right)^{\frac{1}{1-\rho}} - \frac{w_t}{\widehat{P}_t^i}\left(\frac{\widehat{P}_t^i}{p_{it}}\right)^{\frac{1}{\alpha(1-\rho)}} z_{it}^{1-\frac{1}{\alpha}} - \frac{z_{it}^{\theta}}{\theta}\right], U\right\}$$

From the inner maximization problem, the price and effort choices of an operating firm are:

$$p\left(w_t; \widehat{P}_t^i\right) = \left[\left(\frac{w_t}{\alpha\rho}\right)^{\theta\alpha}\left(\frac{1}{\rho(1-\alpha)}\left(\widehat{P}_t^i\right)^{\frac{\theta-\rho}{1-\rho}}\right)^{1-\alpha}\right]^{\frac{(1-\rho)}{\theta(1-\alpha\rho)-\rho(1-\alpha)}} \tag{3}$$

$$z\left(w_t; \widehat{P}_t^i\right) = \left[\rho(1-\alpha)^{1-\alpha\rho}\left(\frac{\alpha}{w_t}\widehat{P}_t^i\right)^{\alpha\rho}\right]^{\frac{1}{\theta(1-\alpha\rho)-\rho(1-\alpha)}}$$

From the outer maximization problem, the set of operating firms is then:

$$O_t(w_t) = \left\{i \mid \widehat{P}_t^i \geq \overline{\psi}(\alpha, \rho, \theta, U)\, w_t\right\} \tag{4}$$

where

$$\overline{\psi}(\alpha, \rho, \theta, U) \equiv \frac{1}{\alpha\rho}\left[\frac{1}{\rho(1-\alpha)}\right]^{\frac{1-\alpha}{\theta\alpha}}\left[\frac{\theta U}{\theta(1-\alpha\rho)-\rho(1-\alpha)}\right]^{\frac{\theta(1-\alpha\rho)-\rho(1-\alpha)}{\theta\alpha\rho}}$$

If $i \notin O_t(w_t)$, then the firm's labor demand is $\ell\left(w_t; \widehat{P}_t^i\right) = 0$. Otherwise, the first order conditions from the inner problem yield:

$$\ell\left(w_t;\widehat{P}_t^i\right)=\left[\rho^{\theta}\left(1-\alpha\right)^{\rho(1-\alpha)}\left(\frac{\alpha}{w_t}\widehat{P}_t^i\right)^{\theta-\rho(1-\alpha)}\right]^{\frac{1}{\theta(1-\alpha\rho)-\rho(1-\alpha)}} \tag{5}$$

Notice that $U=0$ implies $O_t\left(w_t\right)=[0,1]$ so $\ell\left(w_t;\widehat{P}_t^i\right)>0$ for all firms. Having $U>0$ is thus what will permit $\ell\left(w_t;\widehat{P}_t^i\right)=0$ for some i. The market clearing wage is then determined according to equation (1), with $O_t\left(w_t^*\right)$ given more precisely by equation (4) and $\ell\left(w_t;\widehat{P}_t^i\right)$ as per equation (5). This yields:

$$w_t^*=\overline{\chi}\left(\alpha,\rho,\theta\right)\left[\int\limits_{O_t(w_t^*)}\left(\widehat{P}_t^i\right)^{\frac{\theta-\rho(1-\alpha)}{\theta(1-\alpha\rho)-\rho(1-\alpha)}}di\right]^{\frac{\theta(1-\alpha\rho)-\rho(1-\alpha)}{\theta-\rho(1-\alpha)}} \tag{6}$$

$$\textit{where}$$

$$\overline{\chi}\left(\alpha,\rho,\theta\right)\equiv\alpha\left[\rho^{\theta}\left(1-\alpha\right)^{\rho(1-\alpha)}\right]^{\frac{1}{\theta-\rho(1-\alpha)}}$$

Finally, the aggregate price level is determined according to equation (2), with $p\left(w_t;\widehat{P}_t^i\right)$ and $O_t\left(w_t^*\right)$ as per (3) and (4). The above equations show that operating firms with higher price expectations charge higher prices. They also hire more labor and exert more effort, resulting in more output. Furthermore, for any given wage, firms with higher price expectations are more likely to operate. The higher the wage though, the smaller the set of operating firms, the lower the output of each operating firm, and the higher the prices charged.

4.3.1 One Forecasting Rule

The following proposition establishes the key implications of our full model when all expectations are neatly captured by a single normal distribution:

Proposition 2 *If $\widehat{\pi}_t^i\sim N\left(g,\sigma^2\right)$ for all i, then:*

1. $\pi_t^*=g+f\left(\sigma\right)$

2. *If $\alpha=1$ and $U=0$, then $f\left(\sigma\right)=\frac{\sigma^2}{2(1-\rho)}$ as per Proposition 1.*

3. *If* $\alpha \leq 1$ *and* $U > 0$*, then the set of operating firms is shrinking in* σ*.*

4. *If* $\alpha = 1$ *and* $U \geq 1 - \rho$*, then* $f(\sigma_0) = 0$ *for a unique* $\sigma_0 > 0$*. Moreover,* $f'(\sigma_0) > 0$*.*

5. *There exist constants* $\underline{\alpha} \in (0,1)$ *and* $\overline{U} > \underline{U} > 0$ *such that* $\alpha \in (\underline{\alpha}, 1)$ *and* $U \in \left(\underline{U}, \overline{U}\right)$ *yield* $f(\sigma_A) = f(\sigma_B) = 0$ *for* $\sigma_B > \sigma_A > 0$*. Moreover,* $f'(\sigma_A) < 0$ *and* $f'(\sigma_B) > 0$*.*

The first part of Proposition 2 says that excess inflation $(\pi_t^* - g)$ depends only on the extent of expectations heterogeneity. The second part then says that linear production and no outside option return the results of Proposition 1. The remainder of Proposition 2 restricts attention to positive outside options. In particular, the third part establishes that σ decreases operation when the option is positive. As discussed in Subsection 4.2, higher σ amplifies the asymmetric effect that high expectation firms have on wage determination. Since higher wages cut into expected firm profits, the presence of a positive outside option means that more firms will choose not to operate. This puts downward pressure on wages and, as described in Section 2, helps offset the Jensen's inequality effect. Indeed, with linear production, the fourth part of Proposition 2 shows that a sufficiently lucrative outside option introduces a point $\sigma_0 > 0$ with no excess inflation (i.e., a point where $f(\cdot) = 0$ or, equivalently, a point of mean rationality). This is illustrated by the solid gray line in Figure 1. Existence of such a point is robust to non-linearities in production and, under some restrictions on α and U, the fifth part of Proposition 2 says that our full model actually produces two mean rational points. This is a more formal version of Lemma 1, with conditions to ensure the existence of both points.

Figure 2 provides a visual of the relevant parameter restrictions when $\rho = 0.9$ and $\theta \to \infty$. As both the proposition and the figure show, we will have two mean rational points if α is not too small and U falls within an intermediate range. Moreover, for any parameter combination represented by a blue dot in Figure 2 (i.e., for any combination that yields two mean rational points), the plot of $f(\cdot)$ will resemble the blue line in Figure 1. Notice from this line that $f(\cdot)$ is negative between the two mean rational points. In other words, realized inflation falls below the mean expectation. We refer to this situation as negative excess inflation. To see why it arises, recall the competing effects of higher σ on wages in equation (6). As σ increases, we know that the compounding of inflation expectations into price expectations skews the distribution of price expectations right and puts upward pressure on the wage through the labor demands of high expectation firms. As the wage increases though, we also know that low expectation firms find it more profitable to take their outside option and the resulting decline in operation puts downward pressure on the wage. For lower values of σ, the exit of low expectation firms dominates and dampens the wage but, when σ becomes sufficiently large, the labor demands of high expectation firms take over. The dependence of w_t^* on σ is thus U-shaped. What does this mean for prices? We know from equation (3) that individual prices respond positively to wages so, all else constant, the shape of w_t^* feeds into P_t^* and π_t^* (for a given P_{t-1}^*). Notice, however, that there is additional upward pressure on P_t^* at the price aggregation stage. Since only firms with sufficiently high expectations produce, the individual prices aggregated by equation (2) are $p\left(w_t^*; \widehat{P}_t^i\right)$ with $\widehat{P}_t^i$ high. Therefore, the pass-through from w_t^* to P_t^* varies across σ but, for $\alpha \in (\underline{\alpha}, 1)$ and $U \in \left(\underline{U}, \overline{U}\right)$, it is enough to generate two mean rational points with a U-shaped pattern in between. If the outside option is too high or the returns to labor are too low, then exit is too strong relative to labor demand and we get only one mean rational point with negative excess inflation to the left of that point and positive excess inflation to the right. If the outside option is too low, then exit is weak and we get positive excess inflation everywhere with no mean rational points.

4.3.2 Two Forecasting Rules

We now investigate what happens if expectations are characterized by a mixture of normal distributions, namely a group of size ξ_t distributed according to $\widehat{\pi}_t^i \sim N\left(g_A, \sigma_A^2\right)$ and a group of size $1 - \xi_t$ distributed according to $\widehat{\pi}_t^i \sim N\left(g_B, \sigma_B^2\right)$. Notice that we now have two dimensions of heterogeneity: within groups (i.e., $\sigma_A > 0$ and $\sigma_B > 0$) and across groups (i.e.,

$g_A \neq g_B$ and/or $\sigma_A \neq \sigma_B$). It will be instructive to begin with $g_A = g_B = g$ for any g.[7] To simplify the exposition, define the following constants:

$$\kappa \equiv \frac{\theta-\rho(1-\alpha)}{\theta(1-\alpha\rho)-\rho(1-\alpha)}, \quad \delta \equiv \frac{(\theta-\rho)(1-\alpha)}{\theta(1-\alpha\rho)-\rho(1-\alpha)}, \quad \text{and} \quad \upsilon_j \equiv \frac{\sigma_j}{1-\rho}$$

Also define variables $y_t \equiv \frac{\pi_t^*-g}{1-\rho}$ and $x_t \equiv \frac{\ln\left(\overline{\psi}w_t^*/P_{t-1}^*\right)-g}{1-\rho}$. We use y_t (multiplied by $1-\rho$) to denote excess inflation when expectations come from two normal distributions and reserve $f(\cdot)$ for when they come from one normal distribution. Interpretation of x_t is with reference to equation (4). In particular, if the difference between a firm's inflation forecast and the mean forecast is greater than or equal to x_t, then the firm operates. The full model with a mixture of normal expectations yields x_t implicitly defined by:

$$x_t = \frac{1}{1-\rho}\ln\left(\overline{\psi}\overline{\chi}\right) + \frac{1}{(1-\rho)\kappa}\ln\left[\xi_t h\left((1-\rho)\kappa, \upsilon_A, x_t\right) + (1-\xi_t)\, h\left((1-\rho)\kappa, \upsilon_B, x_t\right)\right]$$

$$\textit{where}$$

$$h(\beta, \upsilon, x) \equiv \exp\left(\frac{(\beta\upsilon)^2}{2}\right)\Phi\left(\beta\upsilon - \frac{x}{\upsilon}\right)$$

If $x_t \geq \Upsilon(x_t, \xi_t) + \ln(\gamma_t)$, then $y_t = \Upsilon(x_t, \xi_t)$ where:

$$\Upsilon(x, \xi) \equiv \frac{1}{\rho}\ln\left(\frac{\theta(1-\alpha\rho)-\rho(1-\alpha)}{\theta U}\right) + \frac{\theta\alpha\kappa(1-\rho)x}{\theta-\rho(1-\alpha)} + \ln\left(\frac{\xi h(-\rho\delta,\upsilon_A,x)+(1-\xi)h(-\rho\delta,\upsilon_B,x)}{\xi h(-\delta,\upsilon_A,x)+(1-\xi)h(-\delta,\upsilon_B,x)}\right) \tag{7}$$

If $x_t < \Upsilon(x_t, \xi_t) + \ln(\gamma_t)$, then y_t solves:

$$y_t = \frac{1}{\rho}\ln\left(\frac{\theta(1-\alpha\rho)-\rho(1-\alpha)}{\theta U}\right) + \frac{\theta\alpha\kappa(1-\rho)x_t}{\theta-\rho(1-\alpha)} + \frac{1}{1-\rho}\ln\left(\frac{\xi_t N(x_t,y_t+\ln(\gamma_t),\upsilon_A)+(1-\xi_t)N(x_t,y_t+\ln(\gamma_t),\upsilon_B)}{\xi_t D(x_t,y_t+\ln(\gamma_t),\upsilon_A)+(1-\xi_t)D(x_t,y_t+\ln(\gamma_t),\upsilon_B)}\right) \tag{8}$$

$$\textit{where}$$

$$N(x, y, \upsilon) \equiv h(1-\rho\delta, \upsilon, x) - h(1-\rho\delta, \upsilon, y) + \exp(y)\, h(-\rho\delta, \upsilon, y)$$

$$D(x, y, \upsilon) \equiv h(1-\delta, \upsilon, x) - h(1-\delta, \upsilon, y) + \exp(y)\, h(-\delta, \upsilon, y)$$

[7] The derivations that follow parallel those in the proof of Proposition 2, Part 1 and are thus omitted. The only difference is the use of a mixture of normals rather than a single normal when evaluating any integrals.

The limiting cases of $\xi_t = 0$ and $\xi_t = 1$ return the full model with expectations characterized by a single normal distribution. For a mixture of distributions, we have:

Proposition 3 *Suppose $\widehat{\pi}_t^i \sim N(g, \sigma_A^2)$ for a group of measure ξ_t and $\widehat{\pi}_t^i \sim N(g, \sigma_B^2)$ for the rest, where $f(\sigma_A) = f(\sigma_B) = 0$. If $\alpha = 1$, then $y_t = 0$ for all $\xi_t \in (0, 1)$.*

Under $\alpha = 1$, Proposition 3 says that the entire population is mean rational when each subpopulation is individually mean rational. We know from Proposition 1 that all firms set the same price when $\alpha = 1$ so any heterogeneity in expectations only affects the economy through labor market clearing, namely equation (1). The latter aggregates linearly across subpopulations so, if the component distributions are parameterized such that $y_t = 0$, then their mixture will also deliver $y_t = 0$.

In contrast, Figure 3 shows what can happen when heterogeneity enters the more complicated aggregation defined by equations (1) and (2): the mixture distribution produces negative excess inflation even if each subpopulation possesses the mean rational property. Using $\rho = 0.9$ and $\theta \to \infty$ as before, panel (a) reveals that combinations of α and U which generate two distinct mean rational distributions also generate negative excess inflation for any mixture of these distributions. Panel (b) then provides a representative plot of y_t as a function of ξ_t. Notice that the shape of y_t over $\xi_t \in [0, 1]$ resembles the shape of $f(\cdot)$ over $\sigma \in [\sigma_A, \sigma_B]$. This is useful as it permits interpretation of our results vis-à-vis Figure 1: if one views $\xi_t \in (0, 1)$ and $g_A = g_B = g$ as generating a roughly normal aggregate distribution with $\sigma \in (\sigma_A, \sigma_B)$, then $\pi_t^* < g$ follows for the reasons discussed in Subsection 4.3.1.

5 Simulations: Introducing Inflation Targets

From the economic model of Section 4, we can calculate π_t^* conditional on ξ_t and the forecasting rules. Using social dynamics as per Section 3, we can then determine ξ_{t+1} conditional on π_t^*, ξ_t, and the forecasting rules. We now investigate how a central bank can use inflation announcements in this environment. We first parameterize the model then conduct simulations to study the effectiveness of announcing inflation targets to achieve a large disinflation.

5.1 Parameterization

We set $\rho = 0.9$ to capture an economy where goods are highly but not perfectly substitutable. We also take the limiting case of $\theta \to \infty$. To obtain two individually mean rational subpopulations (i.e., to obtain two forecasting rules that are each unbiased when adopted by all firms), we pick α and U from the blue region in Figure 2. We set $\alpha = 0.9$ and $U = 0.18$ but any choice from the aforementioned region will deliver qualitatively similar results. Lastly, we assume the taste shock is distributed according to $\gamma_t \sim \mathcal{U}[0.99, 1.01]$. These parameters deliver $\sigma_A = 0.0036$ and $\sigma_B = 0.0643$ as the solutions to $f(\cdot) = 0$.

To complete the characterization of the forecasting rules, we need g_A and g_B. In a special question, the 2012Q2 Philadelphia Fed Survey of Professional Forecasters (SPF) asked each respondent to indicate whether his/her forecasts were consistent with the Fed's inflation target. The results reveal that forecasters who self-identify as consistent with the Fed's target form a tight distribution around this target. In contrast, the remaining forecasters form a wider distribution around past inflation. As in Section 2, let $\overline{\pi}_t$ denote the central bank's date t announcement. Since individual expectations and professional forecasts are not unrelated (e.g. Carroll (2003b)), we can map the SPF results into our model as follows:

Definition 1 *Fed Followers (FFs):* $\widehat{\pi}_t^i \sim N\left(\overline{\pi}_t, \sigma_A^2\right)$

Definition 2 *Random Walkers (RWs):* $\widehat{\pi}_t^i \sim N\left(\pi_{t-1}^*, \sigma_B^2\right)$

The fraction of FFs is ξ_t. That our first group uses the central bank to guide its expectations can be interpreted vis-à-vis Faust and Wright (2012) who find that the Fed's Greenbook forecasts are difficult to beat. The existence of our first group is also consistent with evidence from Campbell et al. (2012) and Gurkaynak et al. (2005) that at least some market participants believe FOMC statements contain new and reliable information about future economic conditions. By endogenizing ξ_t, we are effectively endogenizing the fraction of participants with such beliefs. That our second group follows a random walk can then be interpreted vis-à-vis Atkeson and Ohanian (2001) who find that random walk forecasts of inflation perform very well against more sophisticated econometric models.

5.2 Results for Baseline Specification

In our baseline social dynamics, randomly matched firms compare forecasting performance and switch rules after being outperformed eight times (i.e., $S = 8$). We use 1000 firms and draw 1000 matches (with replacement) at the end of each period. Figure 4 presents results for the introduction of a 2% inflation target in an economy with 20% initial inflation. Blue lines average over 100 simulations while shaded areas are $[10\%, 90\%]$ confidence intervals.

Initially, $\xi_1 = 0$ so all firms are forecasting according to $\widehat{\pi}_1^i \sim N(20\%, \sigma_B^2)$ and the mean rational property yields $\pi_1^* = 20\%$. Consider first a central bank that introduces its target abruptly, announcing $\overline{\pi}_t = 2\%$ for all $t \geq 2$. The bank's announcement introduces a new forecasting rule which a small fraction of firms, $\mu = 0.02$, mutate towards. Figure 4(a) demonstrates that inflation converges to 2% but is followed by a temporary overshooting of the target. Recall from Section 4.3 that firms with low expectations (relative to their peers) are less likely to operate. FFs thus do not participate in the labor market early on, putting downward pressure on input prices and lowering inflation. To see why overshooting emerges, turn to the fraction of FFs just before the economy reaches 2%. With realized inflation near target and $\sigma_A < \sigma_B$, Fed Following is often a better forecasting rule than Random Walking. If beliefs were not stubborn (i.e., if S was low), RWs would switch very quickly and ξ_t would rise sharply. Virtually all firms would then forecast according to $\widehat{\pi}_t^i \sim N(2\%, \sigma_A^2)$ and we would thus observe $\pi_t^* = 2\%$. With stubbornness, however, the economy reaches 2% with a mix of FFs and RWs which, as per Subsection 4.3.2, generates $\pi_t^* < 2\%$. Over time though, RWs accumulate enough strikes to compel them to become FFs, returning inflation to target.

Figure 4(b) shows that overshooting can be avoided with gradual targets - that is, a path which interpolates between initial inflation and the long-run target. By achieving interim targets along this path, the central bank converts more firms into FFs on the way to 2%. In turn, the economy is very close to a situation where all firms forecast according to $\widehat{\pi}_t^i \sim N(2\%, \sigma_A^2)$ when 2% is actually reached. The prediction that gradual targets can avoid the overshooting associated with abrupt targets is consistent with empirical evidence. In particular, data from Mishkin and Schmidt-Hebbel (2007) reveals that countries such as Chile, Mexico, Columbia, and Peru introduced their targets more gradually than countries such as Canada, Sweden, the UK, and the Czech Republic. Incidentally, the first group seems to have experienced less overshooting than the second.

5.3 Results for Alternative Specifications

Lower Stubbornness The results so far have considered firms that are somewhat stubborn in their beliefs, refusing to switch forecasting rules at the first sign of a better rule. We now use $S = 1$ to see what happens if beliefs are less stubborn. Figure 5 shows that an abrupt introduction no longer leads to overshooting. As noted earlier, RWs who are not stubborn will switch rules very quickly once inflation approaches 2%, implying excess inflation of virtually zero. Notice, however, that $S = 1$ converges more slowly than $S = 8$ and with wider confidence bands. Slower convergence stems from fewer FFs persisting in early tournament selections. Without a large endowment of FFs, realized inflation remains relatively close to 20% for the first few periods so the huge gap between 20% and the mean FF forecast of 2% implies that FF forecasts are almost always outperformed by RW forecasts. Under low stubbornness, this will prompt FFs to switch rules very quickly and return to the labor market, thus mitigating the downward pressure through input prices. Wider confidence bands stem from ξ_t (and thus π_t^*) being more sensitive to the specific pattern of random meetings during tournament selection now that firms do not distinguish between one-time outperformance and sustained outperformance. Stubbornness thus has advantages and disadvantages for a central bank trying to achieve a drastic reduction in inflation. On one hand, higher stubbornness among FFs yields faster and more certain convergence to the bank's target but, on the other hand, higher stubbornness among RWs leads to a temporary overshooting of the target if the target is introduced abruptly.

Local Interactions We now return to $S = 8$ and relax the assumption that firms meet at random to compare forecasting rules. Suppose instead that tournaments occur locally, with each firm meeting its right and left neighbors every period. As discussed in Section 3, this set up will lead to clusters of FFs and clusters of RWs, increasing the effective stubbornness of firms at the center of each cluster. Figure 6 illustrates the clustering for a subset of 300 firms. A white (blue) dot at coordinate (i, t) means that firm i is a RW (FF) at date t. Figure 7 then shows that overshooting will be more pronounced under an abrupt target and the central bank will need to be more gradual to prevent it if interactions are local rather than random. This is consistent with local interactions generating more stubbornness.

5.4 Comparison to Benchmarks

Fixed Proportions A key insight from the above discussion is that the occurrence of overshooting hinges on the fraction of FFs when the economy reaches the long-run target. To better appreciate the role of social dynamics in determining this fraction, it will be instructive to compare our baseline results with a benchmark that fixes $\xi_t = \overline{\xi}$ for all t. The comparison is presented in Figure 8 for different values of $\overline{\xi}$. If $\overline{\xi} = 0$ (i.e., if the central bank is never credible and no one uses its announcements as a basis for forecasting), then inflation is stable at 20% and announcements are never effective. In contrast, if $\overline{\xi} = 1$ (i.e., if the central bank is always credible and everyone uses its announcements as a basis for forecasting), then the introduction of abrupt targets makes inflation fall to 2% immediately and with no overshooting. Consider now $\overline{\xi} \in (0, 1)$ so that the mix of FFs and RWs is constant but interior. The introduction of targets still succeeds in lowering inflation but we do not drop to 2% immediately. Moreover, if the mix of FFs and RWs is sufficiently interior, then inflation settles noticeably below 2%. With a constant mix, ξ_t is independent of how the central bank introduces its target and how well different rules perform so there is no mechanism to eliminate overshooting. Endogenizing credibility thus introduces an important channel through which central bank announcements affect inflation, providing a richer and more plausible set of dynamics.

Mutation Only Recall that our social dynamics have two elements: mutation and tournament selection. To see the impact of each, Figure 9 compares the full dynamics from Figure 4 against the results that would arise under only mutation. With just mutation, $\xi_{t+1} = (1-\mu)\,\xi_t + \mu\,(1-\xi_t)$ for all t so the fraction of FFs converges smoothly to 0.5. The contribution of tournaments over and above mutation is visible at several points. When targets are introduced abruptly, many FF forecasts are initially outperformed by RWs so tournaments slow the accumulation of FFs and extend the time needed to hit 2%. Around 2% though, the tables turn and many RW forecasts are outperformed by FFs so tournaments accelerate the accumulation of FFs and ensure convergence to 2%. Moreover, when targets are introduced gradually, the accelerated accumulation occurs before 2% is actually reached, allowing convergence to be achieved without even a temporary overshooting. Again then, letting credibility evolve within the model generates fundamentally different predictions.

6 Simulations: Eliminating Deflation

Having seen how communication can be used to reduce inflation, we now investigate how it can be used to pull the economy out of deflation. We keep the parameterization as in Subsection 5.1 and start by considering announcements regarding the 2% target. We then move to announcements like the Fed's recent Quantitative Easing (QE) program which have the potential to skew the entire distribution of inflation expectations.

6.1 Using Targets

Suppose the economy starts at -2% inflation and the central bank announces that it will target 2% for all t. In our previous simulations, FFs were the low expectation firms and their initial impact was to decrease inflation via exit. Now, however, FFs are the high expectation firms so their initial impact is to increase inflation via price-setting. For our baseline specification, Figure 10 shows that the accumulation of FFs - at first through mutation and later through tournament selection - eventually brings the economy up to 2%. Figure 11(a) then shows that convergence to 2% now occurs more quickly under lower stubbornness. The difference between initial inflation of -2% and the mean FF forecast of 2% is such that FFs are not always outperformed by RWs in early tournaments. As a result, some Random Walkers incur strikes early on and, the faster they switch rules, the faster 2% is reached. As discussed in Subsection 5.3, local interactions increase effective stubbornness so, for a given value of S, local rather than random interactions would increase the time it takes to reach 2% when the economy begins at -2%. This is confirmed by Figure 11(b).

So far, we have assumed that the central bank keeps the target at 2% for all t. What happens if it instead decides to gradually lead the economy back to 2%? As Figure 12(a) reveals, a very gradual strategy involves more persistent deflation early on. Recall from Subsection 4.3 that firms with higher price expectations set higher prices. The central bank's gradual path initially implies $\overline{\pi}_1 = -2\% + \varepsilon$ where $\varepsilon > 0$ is small so the average FF only sets a slightly higher price than the average RW. The upward pressure from FF price-setting thus does not outpace the downward pressure from RW exit and the economy continues to experience deflation.

Lastly, Figure 12(b) shows that aggressive communication may be the best at eliminating deflation. The path we consider is one where the central bank announces short-term targets that are well above the long-run goal of 2%. Aggressive short-term targets induce any FFs to set very high prices, pushing realized inflation upwards. At the same time, however, the big gap between realized and targeted inflation does nothing to help the central bank accumulate more FFs and bring ξ_t towards 1. Therefore, when the target returns to 2% and the economy approaches it from above, we have the same overshooting problem we had in Figure 4(a). In order to eliminate this dip, the central bank would have to implement a gradual path on the way down to 2% and thus keep the economy above 2% for longer. To some extent, these results suggest that price-level targeting - which would indeed require the bank to balance out periods of deflation with periods of high inflation - has some advantages over inflation targeting in dealing with deflations when expectations exhibit some stubbornness.

6.2 Quantitative Easing

Our focus thus far has been on how central banks can use precise announcements about inflation to steer the course of actual inflation. To avoid any premature assumptions about the power of the bank, we restricted the direct effect of its announcements to just the mean of the FF distribution. If the bank has absolutely no followers, then the announcements are irrelevant as illustrated in Figure 8(a). If the bank has at least a few followers, then we showed how announcements can be tailored to endogenously increase this following by capitalizing on the social dynamics between firms.

We now consider what happens if central bank communication is more potent than previously assumed. In particular, instead of just changing the mean of the FF distribution, suppose announcements can directly affect the skewness of the economy-wide distribution. A practical example is what Krishnamurthy and Vissing-Jorgensen (2011) dub the inflation channel of QE - that is, the curtailment of deflation expectations due to publicity surrounding the Fed's recent large-scale asset purchases. We introduce this channel into our model via redraws. More precisely, some firms with deflationary expectations redraw their $\widehat{\pi}_t^i$'s after hearing that the central bank will take a proactive approach to stimulating the economy. Each redraw comes from the same distribution as the original draw so not all deflationary expectations will be eliminated. However, redraws do have the effect of skewing the RW and FF distributions so that more mass exists to the right of the mean. Since very few FFs actually expect deflation, the skew is stronger for RWs but, either way, the effect of QE communications is to increase expectations, reduce dispersion, enlarge the set of operating firms (by extension of Proposition 2, part 3), and put upward pressure on inflation.

We consider two dimensions of QE: rounds and intensity. In our context, rounds means the number of periods with media coverage about QE and, therefore, the number of periods that have redraws. Intensity means the fraction of deflationary firms that are exposed to this coverage and, therefore, the fraction that redraw in a given period. Our central bank again faces -2% inflation but, as an inflation targeter, would like to return the economy to 2% without changing its short-term targets (i.e., without pursuing the aggressive strategy in Figure 12(b)). Figures 13 and 14 illustrate how this can be achieved in our baseline specification ($S = 8$, random interactions) by varying rounds and intensity.

To isolate the effect of rounds, Figure 13 fixes intensity at 1. In other words, QE announcements are so pervasive and authoritative that all firms with deflationary expectations redraw. Panel (a) demonstrates that one round of QE announcements helps increase inflation but more time is needed to accumulate FFs and reach 2%. Panel (b) shows that two rounds of QE announcements actually push the economy above 2% for a short-time then below 2% for several periods. If the central bank wants to eliminate the dip back below target, it must increase the number of rounds. However, increasing rounds without decreasing intensity means that the bank has to tolerate more above-target inflation.

Figure 14(a) shows that two rounds with less than full intensity can bring the economy to 2% quickly and without any time above target. However, once the rounds run out, the economy dips back below target for several periods. Just as redraws skew the distributions and increase operation, the end of redraws unskews the distributions and decreases operation. Therefore, if the fraction of FFs is low when the redraws stop, exit among RWs returns the overshooting problem of Figure 4(a).

Finally, Figure 14(b) illustrates the outcome of many rounds and low intensity. Why do many rounds make it possible to find a monotonicity-inducing intensity? Avoiding the rise above 2% experienced in Figure 13(b) requires stopping QE right when inflation hits its target. At the same time, avoiding the dip below 2% experienced in Figure 14(a) requires a very high fraction of FFs when QE stops. Therefore, with T rounds of QE, the bank has T periods to accomplish two things: hit 2% and accumulate a lot of FFs. As we shorten T, accumulating a lot of FFs requires higher intensity. However, higher intensity also hastens the return to 2%. When firms are stubborn in their beliefs, a small increase in intensity will have a stronger effect on the speed of recovery than it will on the accumulation of FFs. The intensity increase needed to accumulate enough FFs thus exceeds the intensity increase needed to hit 2%. Stated otherwise, decreasing the number of rounds makes it harder to find an intensity that returns inflation to 2% monotonically.

7 Conclusion

This paper has investigated the effectiveness of central bank communication when price-setters with heterogeneous inflation expectations are subject to social dynamics. Prolonged periods of divergence between realized inflation and central bank announcements can lead to a loss of credibility through these dynamics and make future announcements much less effective. In this context, we identified how central bank communications can be tailored to endogenously build credibility. We demonstrated that the abrupt introduction of lower inflation targets can lead to a temporary overshooting of the target. In contrast, gradually introducing the target (i.e., via interim targets) directs the economy to the long-term goal more smoothly. Our next set of results concerned communications to guide the economy away from deflation. We found that avoiding a protracted deflation requires an aggressive rather than gradual approach, with price-level targeting conferring some communication-based benefits over inflation targeting. We then studied two dimensions of quantitative easing: number of rounds and intensity of announcements. Our results indicated that the inflation channel of QE is an effective way for an inflation targeting central bank to guide the economy out of deflation without announcing higher short-term targets. However, the mix of rounds and intensity needed to guide the economy out monotonically depends on the stubbornness of agents' beliefs.

References

Andreoni, James and Tymofiy Mylovanov, "Diverging Opinions," *American Economic Journal: Microeconomics*, 2012, *4* (1), pp. 209–232.

Arifovic, Jasmina, Herbert Dawid, Christophe Deissenberg, and Olena Kostyshyna, "Learning Benevolent Leadership in a Heterogenous Agents Economy," *Journal of Economic Dynamics and Control*, 2010, *34* (9), pp. 1768–1790.

_ **, James Bullard, and Olena Kostyshyna**, "Social Learning and Monetary Policy Rules," *The Economic Journal*, 2012.

Ashraf, Quamrul and Peter Howitt, "How Inflation Affects Macroeconomic Performance: An Agent-Based Computational Investigation," 2008. Working Paper, Brown University.

Atkeson, Andrew and Lee E. Ohanian, "Are Phillips Curves Useful for Forecasting Inflation?," *FRB Minneapolis Quarterly Review*, 2001, *25* (1), pp. 2–11.

Berardi, Michele and John Duffy, "The Value of Transparency when Agents are Learning," *European Journal of Political Economy*, 2007, *23* (1), pp. 9–29.

Blinder, Alan S., Michael Ehrmann, Marcel Fratzscher, Jakob De Haan, and David-Jan Jansen, "Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence," *Journal of Economic Literature*, 2008, *46* (4), pp. 910–945.

Boivin, Jean, "How People Think and How It Matters," 2011. Remarks to the Canadian Association for Business Economics, Kingston, ON.

Branch, William A. and George W. Evans, "Unstable Inflation Targets," 2011. Working Paper, University of California, Irvine.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo, "Understanding Booms and Busts in Housing Markets," 2013. Working Paper, Northwestern University.

Campbell, Jeffrey R., "Open Mouth Operations," 2013. Mimeo, FRB Chicago.

_ , Charles L. Evans, Jonas D.M. Fisher, and Alejandro Justiniano, "Macroeconomic Effects of FOMC Forward Guidance," *Brookings Papers on Economic Activity*, 2012, *Spring 2012*, pp. 1–54.

Carney, Mark, "A Monetary Policy Framework for All Seasons," 2012. Remarks to the U.S. Monetary Policy Forum, New York, NY.

Carroll, Christopher D., "The Epidemiology of Macroeconomic Expectations," in L. Blume and S. Durlauf, eds., *The Economy as an Evolving Complex System, III*, Oxford University Press, 2003.

_ , "Macroeconomic Expectations of Households and Professional Forecasters," *Quarterly Journal of Economics*, 2003, *118* (1), pp. 269–298.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 2005, *113* (1), pp. 1–45.

Clarida, Richard, Jordi Gali, and Mark Gertler, "The Science of Monetary Policy: A New Keynesian Perspective," *Journal of Economic Literature*, 1999, *37* (4), pp. 1661–1707.

Colander, David, Peter Howitt, Alan Kirman, Axel Leijonhufvud, and Perry Mehrling, "Beyond DSGE Models: Toward an Empirically Based Macroeconomics," *American Economic Review: Papers and Proceedings*, 2008, *98* (2), pp. 236–240.

Cukierman, Alex and Allan H. Meltzer, "A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information," *Econometrica*, 1986, *54* (5), pp. 1099–1128.

D'Amico, Stefania and Thomas B. King, "Flow and Stock Effects of Large-Scale Treasury Purchases," 2010. Finance and Economics Discussion Series 2010-52, Federal Reserve Board.

Eggertsson, Gauti B. and Benjamin Pugsley, "The Mistake of 1937: A General Equilibrium Analysis," *Monetary and Economic Studies*, 2006, *24* (S-1), pp. 1–58.

Ellison, Glenn and Drew Fudenberg, "Rules of Thumb for Social Learning," *Journal of Political Economy*, 1993, *101* (4), pp. 612–643.

Eusepi, Stefano and Bruce Preston, "Central Bank Communication and Expectations Stabilization," *American Economic Journal: Macroeconomics*, 2010, *2* (3), pp. 235–271.

Faust, Jon and Jonathan H. Wright, "Forecasting Inflation," 2012. Working Paper, John Hopkins University.

Gagnon, Joseph, Matthew Raskin, Julie Remache, and Brian Sack, "Large-Scale Asset Purchases by the Federal Reserve: Did They Work?," 2010. Federal Reserve Bank of New York Staff Reports.

Gurkaynak, Refet S, Brian Sack, and Eric T Swanson, "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements," *International Journal of Central Banking*, 2005, *1* (1), pp. 55–93.

Hamilton, James D. and Jing Cynthia Wu, "The Effectiveness of Alternative Monetary Policy Tools in a Zero Lower Bound Environment," *Journal of Money, Credit & Banking*, 2012, *44 (s1)*, 3–46.

Judd, Kenneth L. and Leigh Tesfatsion, *Handbook of Computational Economics: Agent-Based Computational Economics*, North Holland, 2006.

King, Mervyn, "Monetary Policy: Practice Ahead of Theory," 2005. Mais Lecture, Cass Business School, London.

Krishnamurthy, Arvind and Annette Vissing-Jorgensen, "The Effects of Quantitative Easing On Interest Rates: Channels and Implications for Policy," 2011. NBER Working Paper No. 17555.

Kroszner, Randall S., "Communications Strategy, Expectations Management, and Central Bank Credibility," 2012. Fall 2011 Brookings Papers on Economic Activity Conference, Washington, DC.

LeBaron, Blake, "Evolution and Time Horizons in an Agent Based Stock Market," *Macroeconomic Dynamics*, 2001, *5* (2), pp. 225–254.

Mankiw, N. Gregory, Ricardo Reis, and Justin Wolfers, "Disagreement about Inflation Expectations," *NBER Macroeconomics Annual 2003*, 2004, *18*, pp. 209–248.

Melosi, Leonardo, "Signaling Effects of Monetary Policy," 2012. Federal Reserve Bank of Chicago WP 2012-05.

Mishkin, Frederic S. and Klaus Schmidt-Hebbel, "Does Inflation Targeting Make a Difference?," 2007. NBER Working Paper No. 12876.

Orphanides, Athanasios and John C. Williams, "Imperfect Knowledge, Inflation Expectations, and Monetary Policy," in B. Bernanke and M. Woodford, eds., *The Inflation-Targeting Debate*, University of Chicago Press, 2005, pp. 201–234.

Page, Scott E., "Aggregation in Agent-Based Models of Economics," *The Knowledge Engineering Review*, 2012, *27* (2), pp. 151–162.

Pezzey, John C.V. and Jason J. Sharples, "Expectations of Linear Functions with respect to Truncated Multinormal Distributions," *Environmental Modelling and Software*, 2007, *22* (7), pp. 915–923.

Smets, Frank and Raf Wouters, "An Estimated Dynamic Stochastic General Equilibrium Model of The Euro Area," *Journal of the European Economic Association*, 2003, *1* (5), pp. 1123–1175.

Stein, Jeremy C., "Cheap Talk and the Fed: A Theory of Imprecise Policy Announcements," *American Economic Review*, 1989, *79* (1), pp. 32–42.

Veronesi, Pietro and Francis Yared, "Short and Long Horizon Term and Inflation Risk Premia in the US Term Structure," 2000. Working Paper, Chicago Booth.

Weizsacker, Georg, "Do We Follow Others when We Should? A Simple Test of Rational Expectations," *American Economic Review*, 2010, *100* (5), pp. 2340–2360.

Williams, John C., "Unconventional Monetary Policy: Lessons from the Past Three Years," 2011. FRBSF Economic Letter.

Woodford, Michael, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton, NJ: Princeton University Press, 2003.

_ , "Central Bank Communication and Policy Effectiveness," 2005. Proceedings of the Federal Reserve Bank of Kansas City Symposium at Jackson Hole, pp. 399-474.

_ , "Macroeconomic Analysis without the Rational Expectations Hypothesis," 2013. Forthcoming, Annual Review of Economics.

Xiong, Wei and Hongjun Yan, "Heterogeneous Expectations and Bond Markets," *Review of Financial Studies*, 2010, *23* (4), pp. 1433–1466.

Figure 1: Graphical Representation of Proposition 2

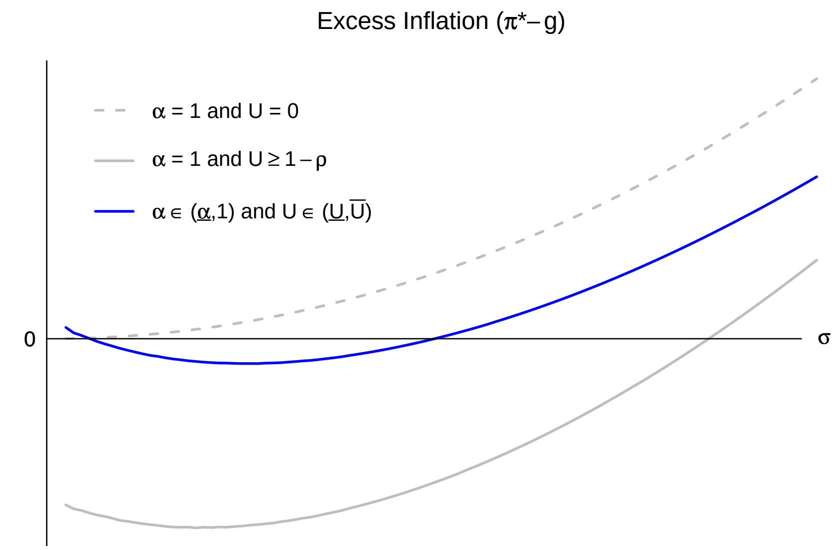

Figure 2: Parameter Space Example for Single Distribution

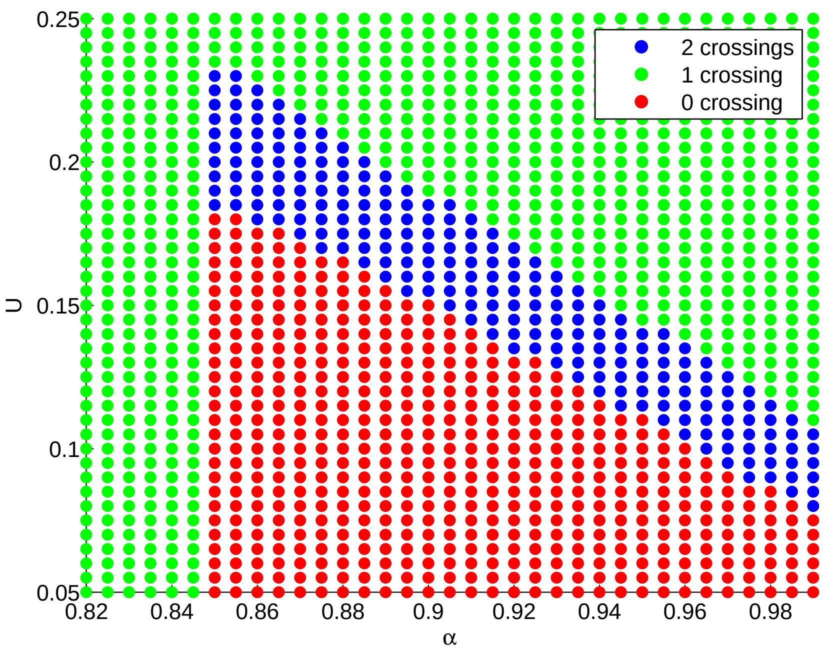

Figure 3: Parameter Space Example for Mixture Distribution

(a) Parameters with $y_t < 0$ when $\xi_t \in (0, 1)$

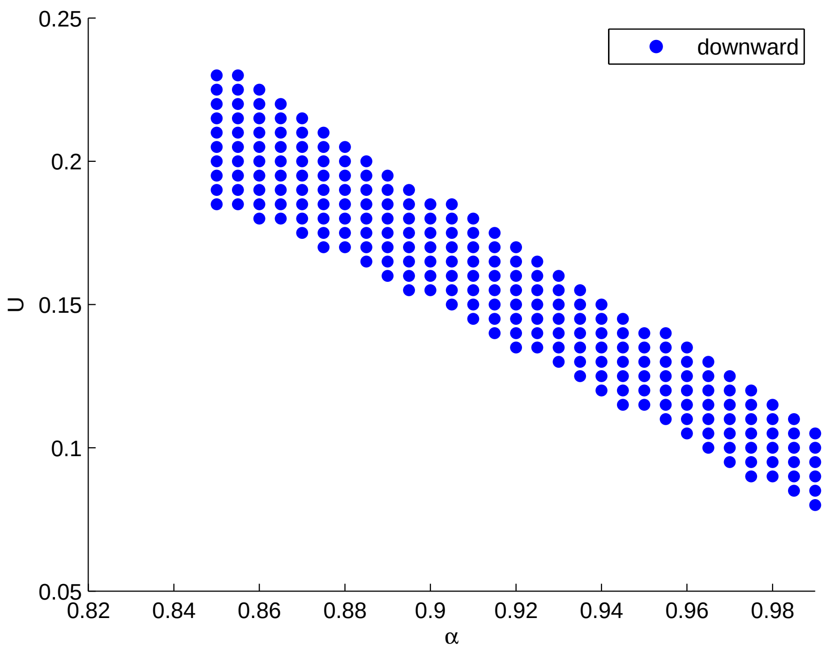

(b) Dependence of y_t on ξ_t

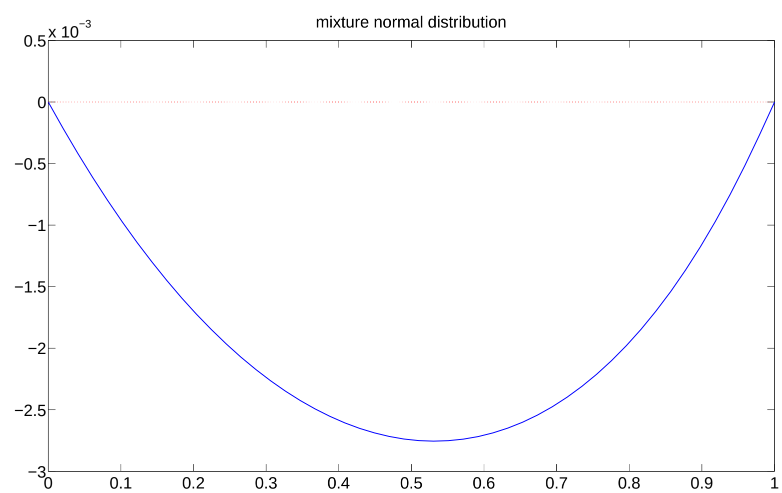

Figure 4: Introduction of IT: S = 8, Random Interactions

(a) Abrupt Strategy

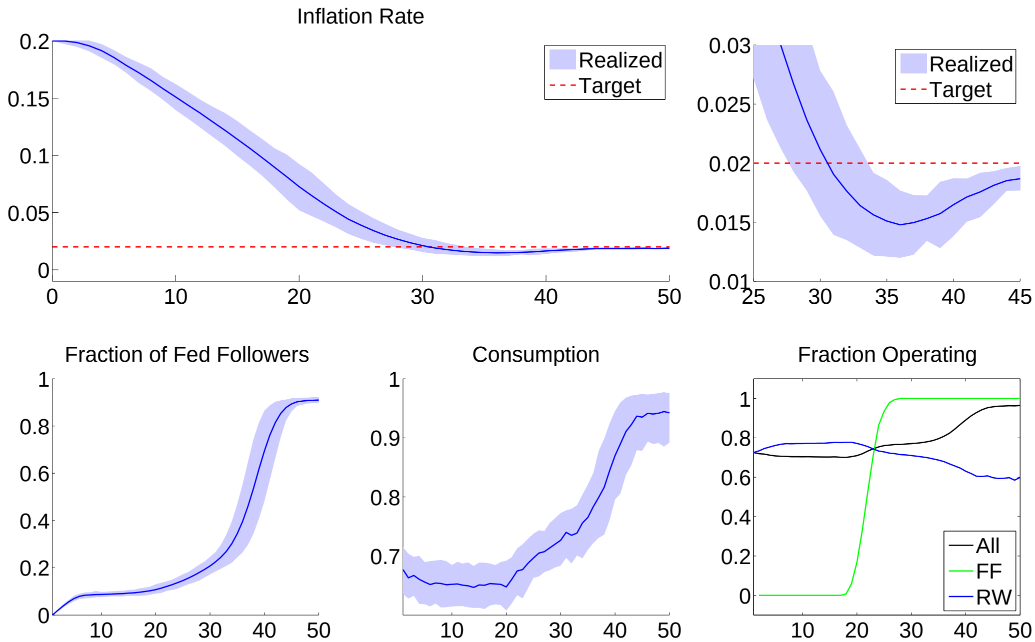

(b) Gradual Strategy

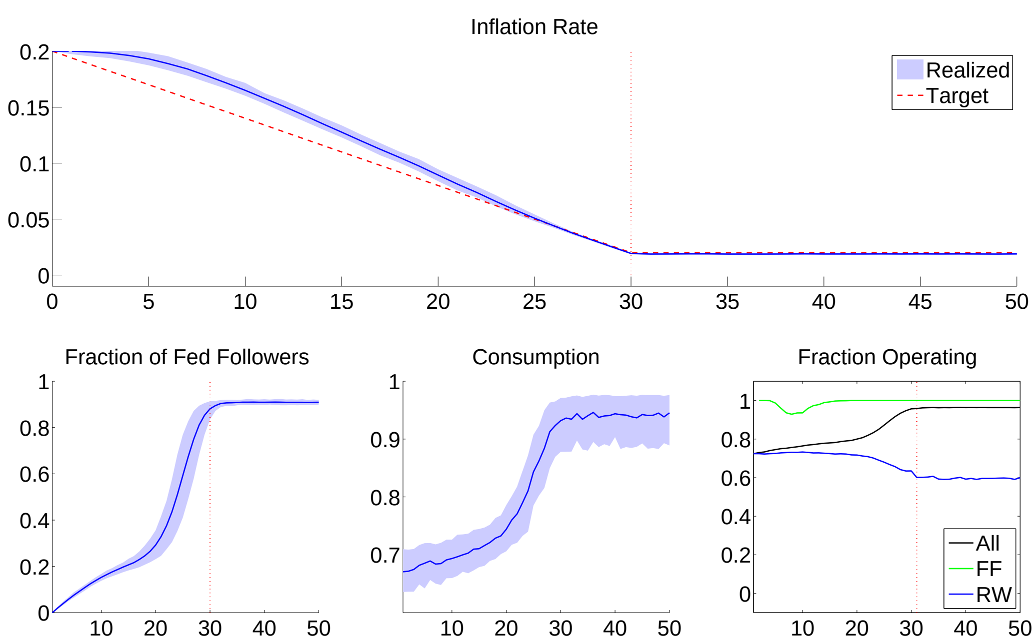

Figure 5: Introduction of IT: S = 1, Random Interactions

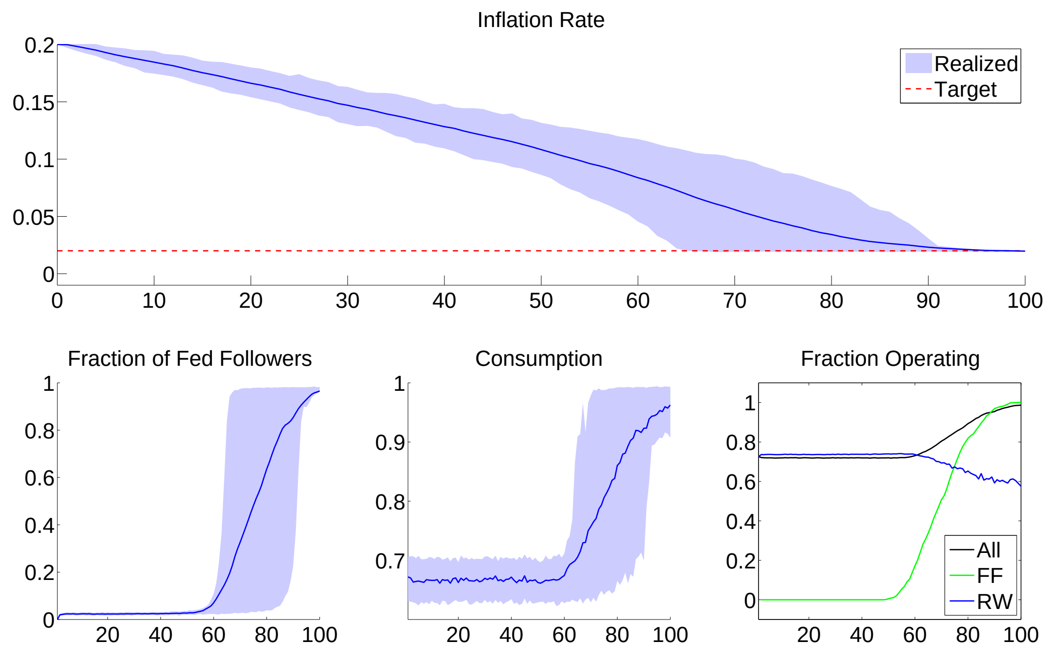

Figure 6: FF Accumulation for Different Interaction Types (S = 8)

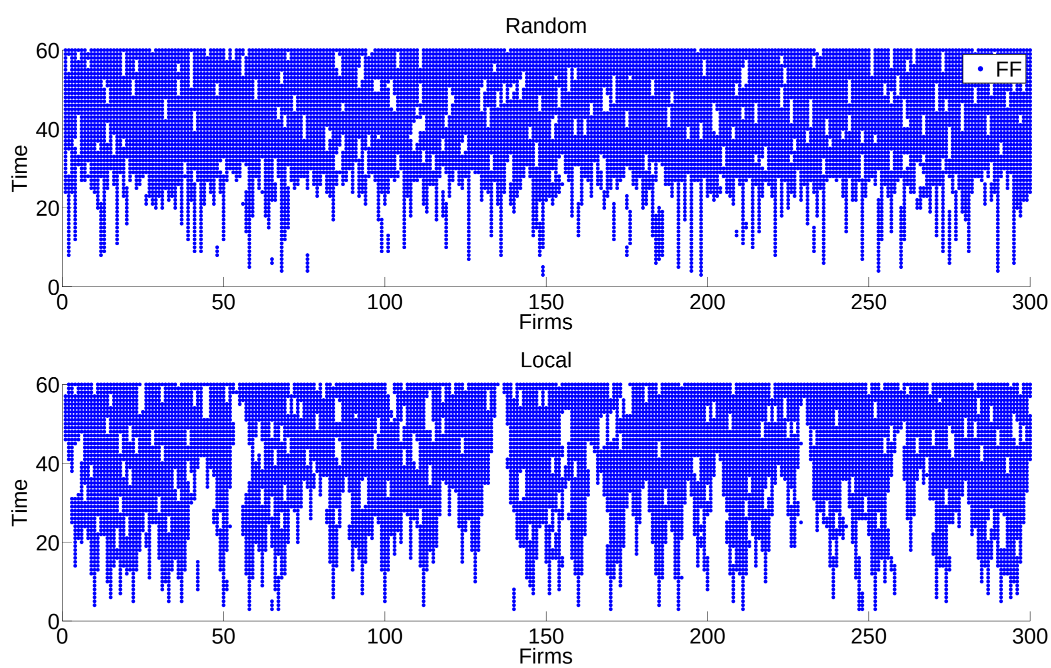

Figure 7: Introduction of IT: S = 8, Local Interactions

(a) Abrupt Strategy

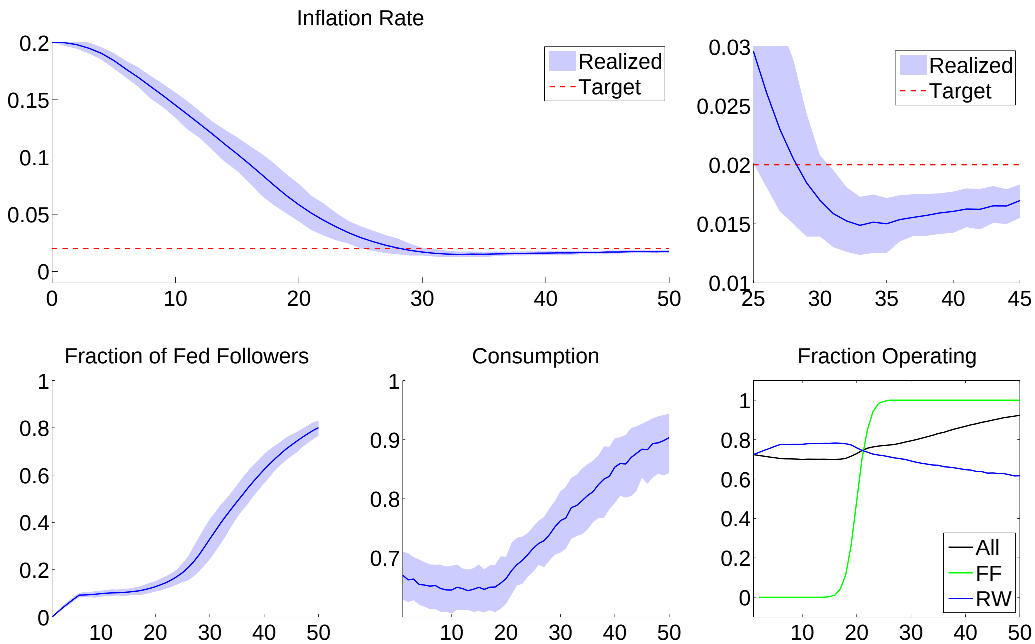

(b) Gradual Strategy

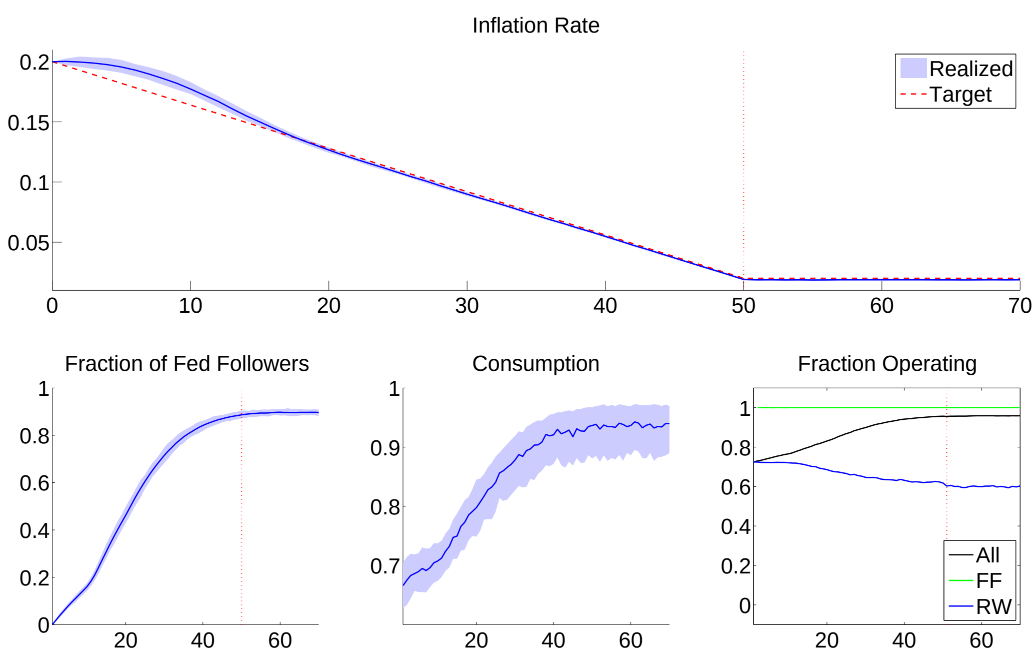

Figure 8: Comparison to Fixed Proportions

(a) Abrupt Strategy

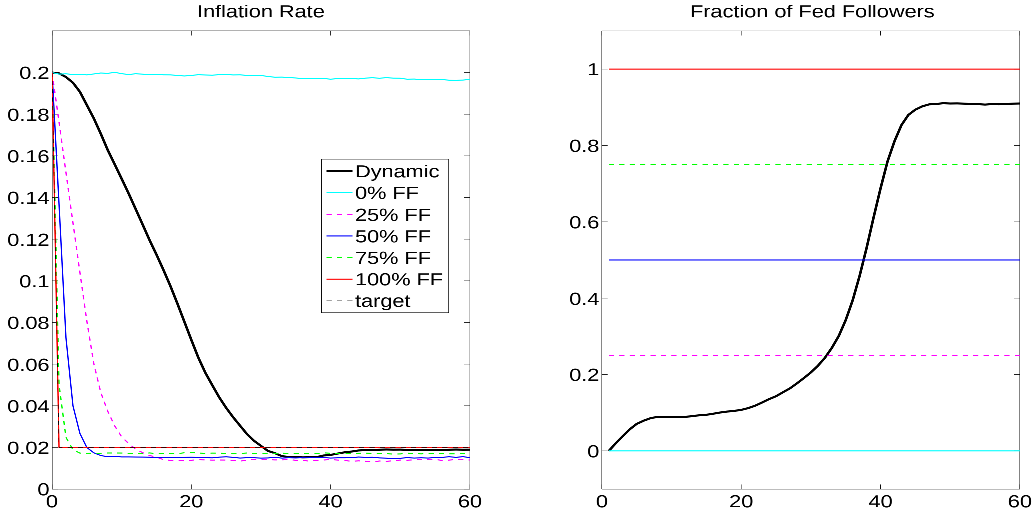

(b) Gradual Strategy

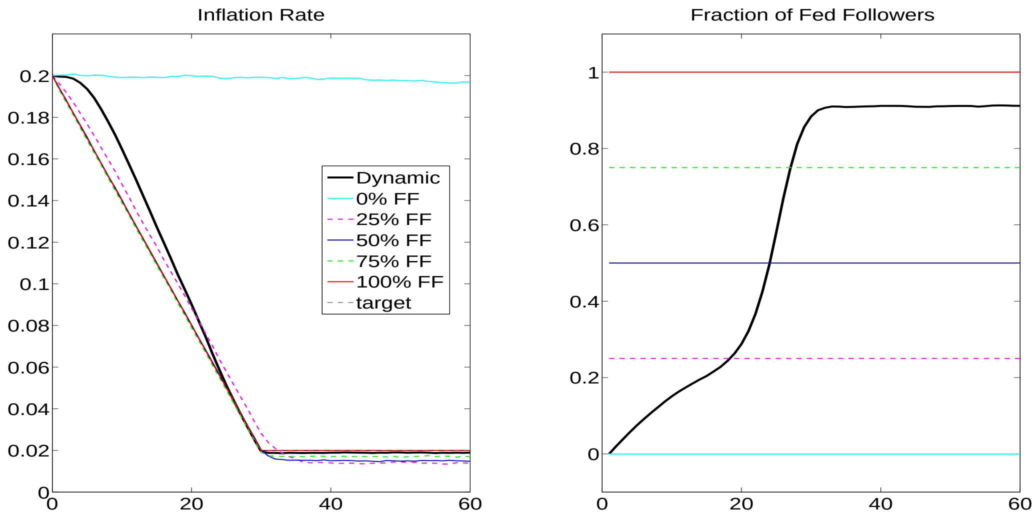

Figure 9: Comparison to Mutation-Only

(a) Abrupt Strategy

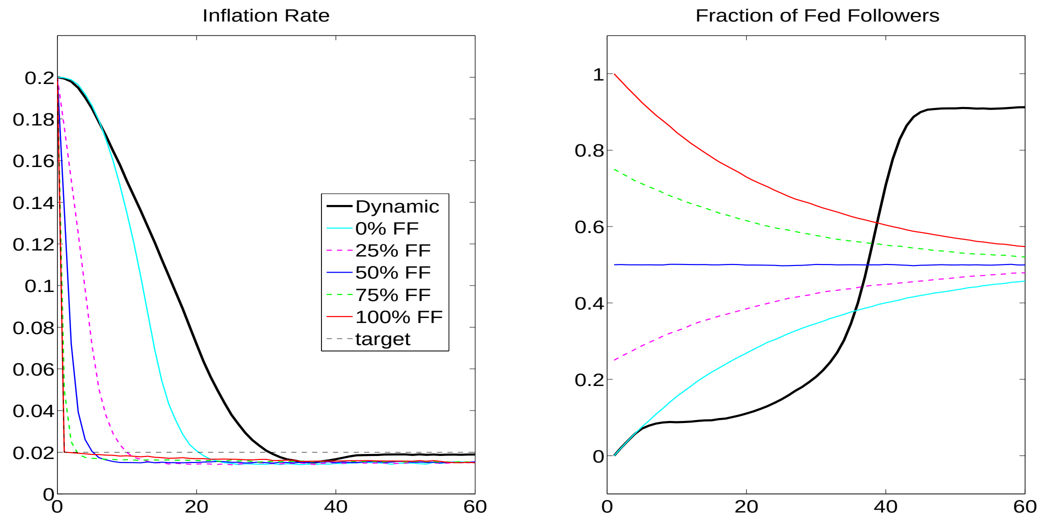

(b) Gradual Strategy

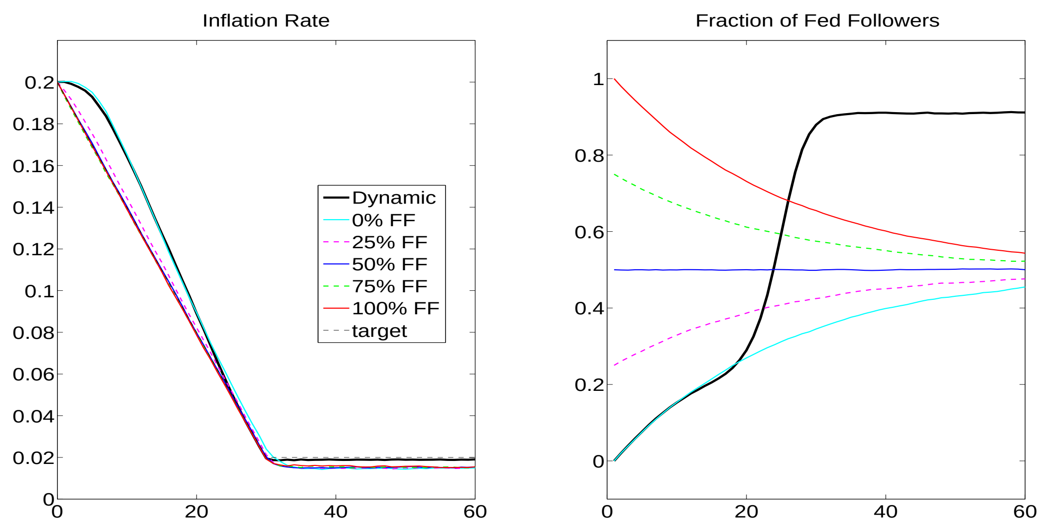

Figure 10: Eliminating Deflation: S = 8, Random Interactions

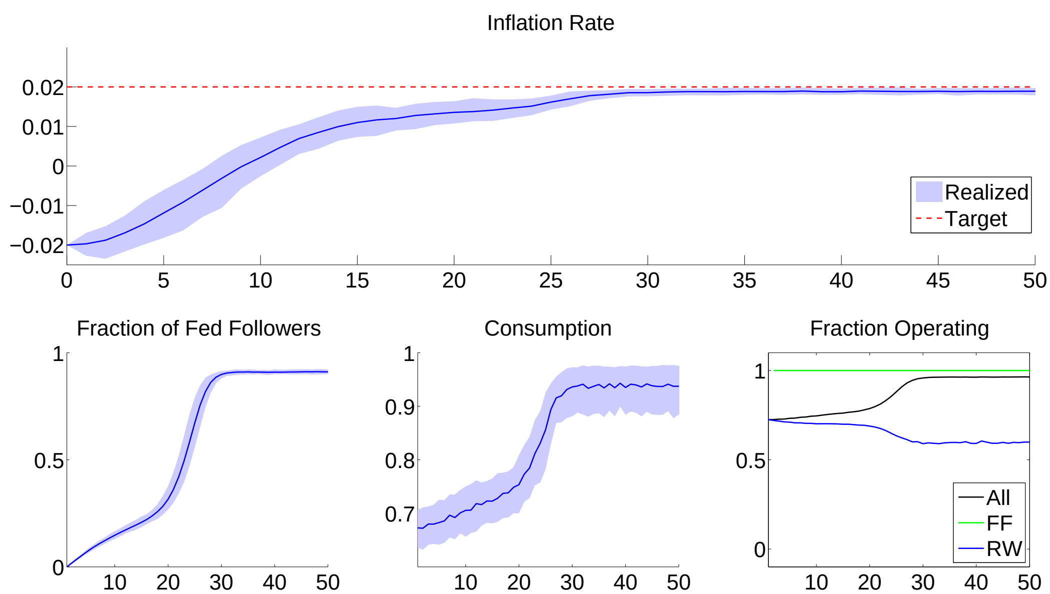

Figure 11: Eliminating Deflation: Alternative Specifications

(a) S=1, Random Interactions

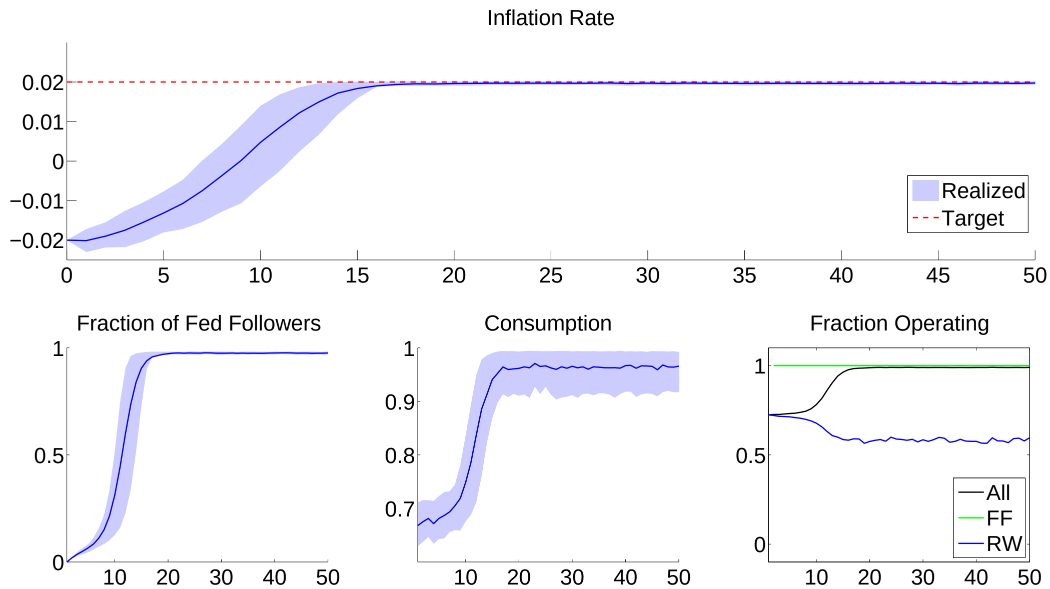

(b) S=8, Local Interactions

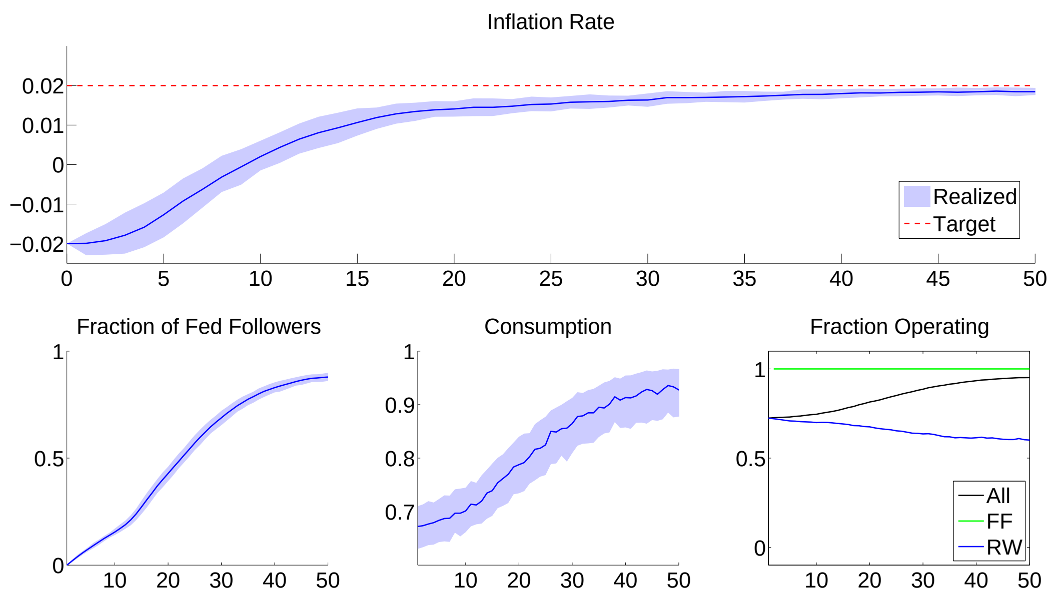

Figure 12: Eliminating Deflation: S = 8, Random Interactions

(a) Gradual Strategy

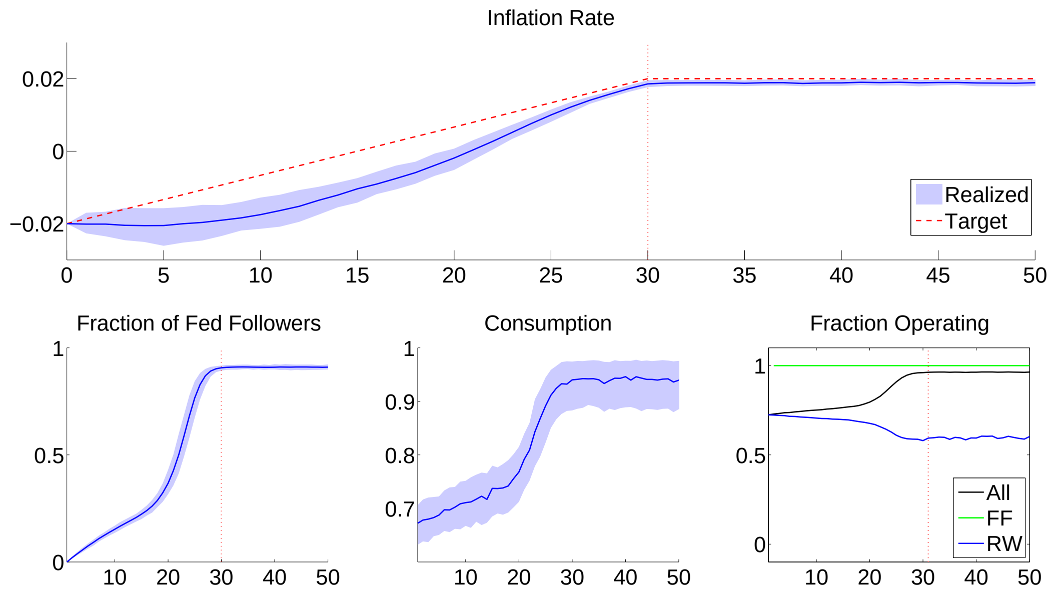

(b) Aggressive Strategy

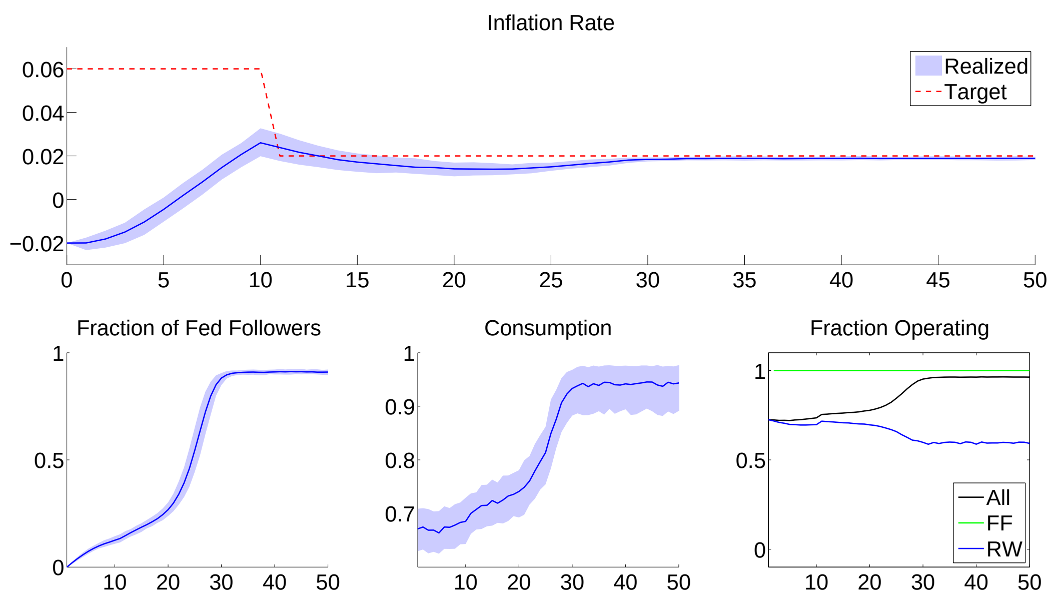

Figure 13: QE Announcements

(a) Rounds = 1; Intensity = 1

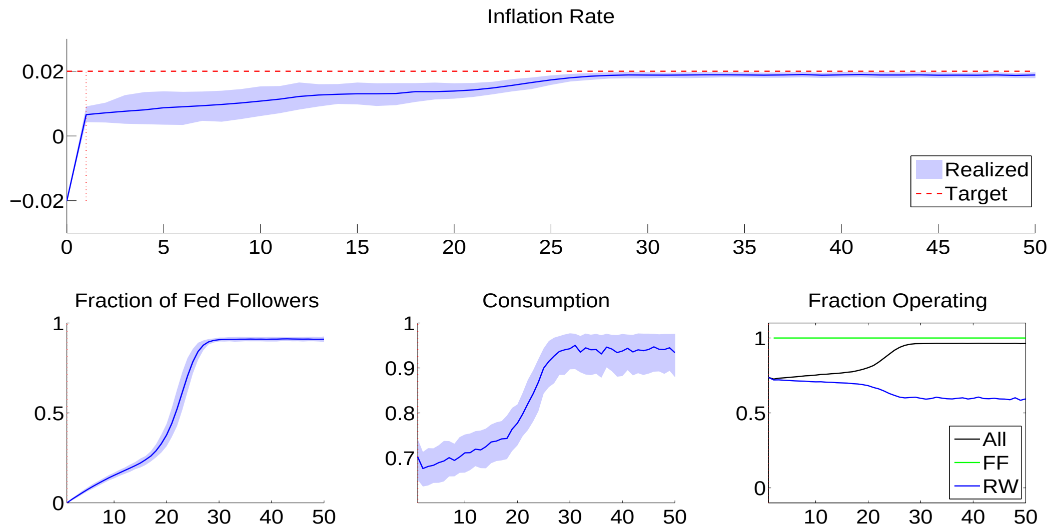

(b) Rounds = 2; Intensity = 1

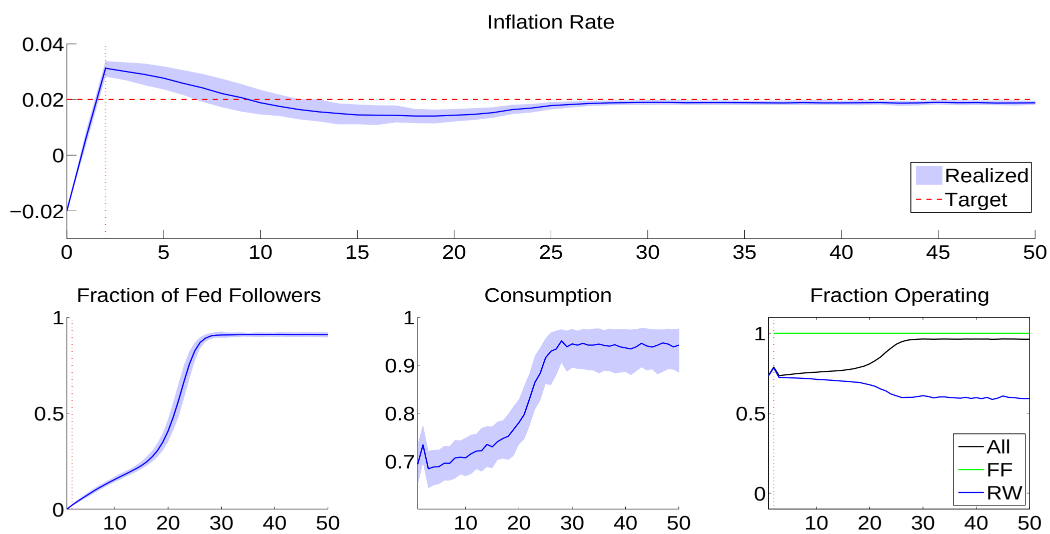

Figure 14: QE Announcements

(a) Rounds = 2; Intensity = 0.75

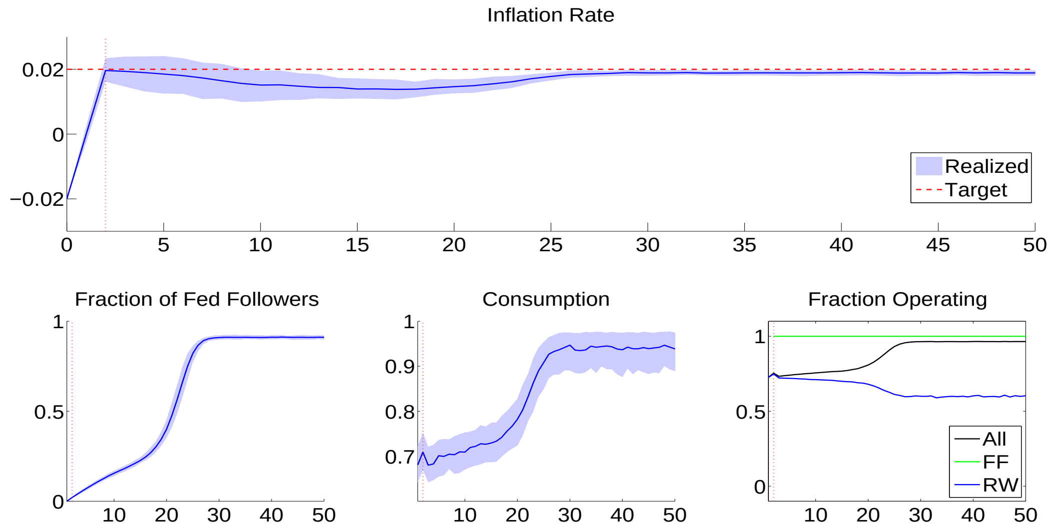

(b) Rounds = 25; Intensity = 0.1

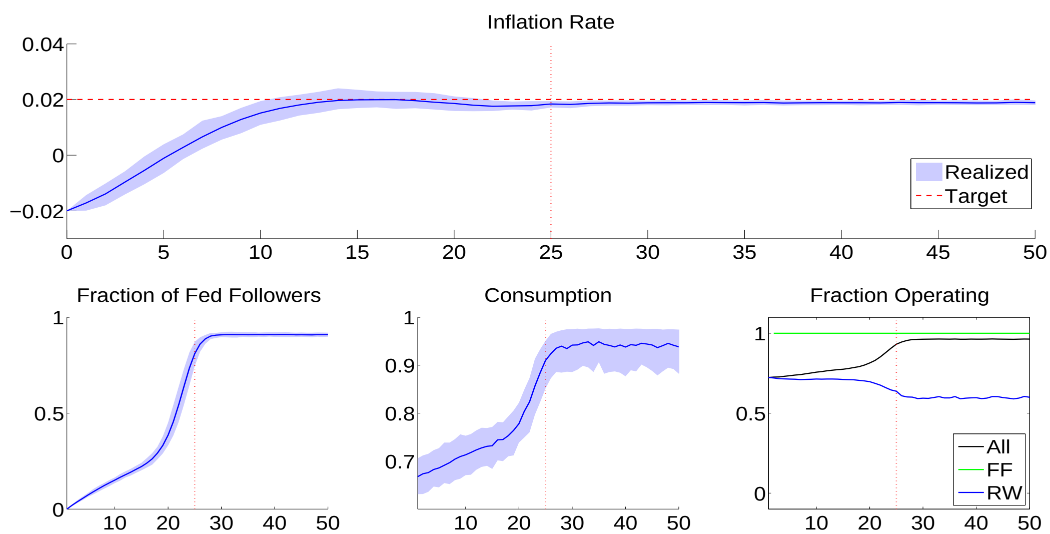

Appendix - Proofs

Proof of Proposition 1

Recall that firm i aims to produce its expected demand $\left(\frac{\widehat{P}_t^i}{p_{it}}\right)^{\frac{1}{1-\rho}}$. Under $F(\ell,\cdot)=\ell$, this will require $\left(\frac{\widehat{P}_t^i}{p_{it}}\right)^{\frac{1}{1-\rho}}$ units of input so i's profit maximization problem is $\max_{p_{it}} (p_{it}-w_t)\left(\frac{\widehat{P}_t^i}{p_{it}}\right)^{\frac{1}{1-\rho}}$. Optimization yields $p\left(w_t;\widehat{P}_t^i\right)=\frac{w_t}{\rho}$ and thus $\ell\left(w_t;\widehat{P}_t^i\right)=\left(\frac{\rho\widehat{P}_t^i}{w_t}\right)^{\frac{1}{1-\rho}}$. Notice that $p\left(w_t;\widehat{P}_t^i\right)$ does not depend on $\widehat{P}_t^i$. Substituting $p\left(w_t^*;\widehat{P}_t^i\right)$ into equation (2) gives $P_t^*=\frac{w_t^*}{\rho}$ and substituting $\ell\left(w_t^*;\widehat{P}_t^i\right)$ into equation (1) gives $\frac{w_t^*}{\rho}=\left[\int\left(\widehat{P}_t^i\right)^{\frac{1}{1-\rho}}di\right]^{1-\rho}$. Combining these two expressions and using the definitions of $\widehat{\pi}_t^i$ and π_t^* then yields $\pi_t^*=(1-\rho)\ln\left(\int\exp\left(\frac{\widehat{\pi}_t^i}{1-\rho}\right)di\right)$. With $\widehat{\pi}_t^i\sim N\left(g,\sigma^2\right)$, we can use the moment generating function of the normal distribution to simplify the preceding integral. Given $\ell\left(w_t^*;\widehat{P}_t^i\right)>0$ for all i, the integral is taken over the entire set so the moment generating function produces $\pi_t^*=g+\frac{\sigma^2}{2(1-\rho)}$. ■

Proof of Proposition 2

Part 1 Let $\phi(\cdot)$ and $\Phi(\cdot)$ denote the standard normal PDF and CDF respectively. From Pezzey and Sharples (2007), the moment generating function of a truncated normal random variable with mean 0 and variance σ^2 is:

$$\int_{x\geq c}\exp(rx)\,\phi\left(x,\sigma^2\right)dx=\exp\left(\frac{r^2\sigma^2}{2}\right)\Phi\left(r\sigma-\frac{c}{\sigma}\right)\tag{9}$$

Using $\widehat{P}_t^i=\exp\left(\widehat{\pi}_t^i\right)P_{t-1}$ and $\widehat{\pi}_t^i=g+\varepsilon_{it}$ with $\varepsilon_{it}\sim N\left(0,\sigma^2\right)$ in equation (4), we can rewrite the operation constraint as:

$$\varepsilon_{it}\geq\ln\left(\frac{\overline{\psi}w_t}{P_{t-1}}\right)-g\equiv X\tag{10}$$

To ease notation, define the following constants:

$$\kappa_1\equiv\frac{\theta-\rho(1-\alpha)}{\theta(1-\alpha\rho)-\rho(1-\alpha)},\ \kappa_2\equiv\frac{\theta\alpha}{\theta(1-\alpha\rho)-\rho(1-\alpha)},\ and\ \kappa_3\equiv\frac{(\theta-\rho)(1-\alpha)}{(1-\rho)[\theta(1-\alpha\rho)-\rho(1-\alpha)]}$$

Combining equations (9) and (10) with the wage equation in (6) then yields an implicit definition of X which is independent of g:

$$X = \frac{1}{\rho\kappa_2}\ln\left(\frac{\kappa_2 U}{\alpha}\right) + \frac{1-\alpha}{[\theta-\rho(1-\alpha)]\kappa_2}\ln\left(\frac{1}{\rho(1-\alpha)}\right) + \frac{\kappa_1\sigma^2}{2} + \frac{1}{\kappa_1}\ln\Phi\left(\kappa_1\sigma - \frac{X}{\sigma}\right) \tag{11}$$

Turn now to inflation. Substitute the firm pricing equation (3) into the price aggregator (2) and simplify to get:

$$\frac{\int_{\widehat{\pi}_t^i \geq \ln\left(\frac{\overline{\psi} w_t^*}{P_{t-1}^*}\right)} \exp\left(\min\left\{\frac{\widehat{\pi}_t^i}{1-\rho}, \frac{\pi_t^* + \ln(\gamma_t)}{1-\rho}\right\} - \kappa_3\widehat{\pi}_t^i\right) di}{\int_{\widehat{\pi}_t^i \geq \ln\left(\frac{\overline{\psi} w_t^*}{P_{t-1}^*}\right)} \exp\left(\min\left\{\frac{\widehat{\pi}_t^i}{1-\rho}, \frac{\pi_t^* + \ln(\gamma_t)}{1-\rho}\right\} - \rho\kappa_3\widehat{\pi}_t^i\right) di} = \frac{\left[\left(\frac{w_t^*}{\alpha\rho}\right)^{\kappa_2}\left(P_{t-1}^*\right)^{\kappa_3}\right]^{(1-\rho)}}{P_t^*\left[\rho\left(1-\alpha\right)\right]^{\frac{(1-\alpha)(1-\rho)}{\theta(1-\alpha\rho)-\rho(1-\alpha)}}} \tag{12}$$

Combining equations (11) and (12) then taking logs yields:

$$\begin{aligned} \pi_t^* \quad = \quad & \ln\left(\frac{\int_{\widehat{\pi}_t^i \geq \ln\left(\frac{\overline{\psi} w_t^*}{P_{t-1}^*}\right)} \exp\left(\min\left\{\frac{\pi_t^* + \ln(\gamma_t)}{1-\rho} - \rho\kappa_3\widehat{\pi}_t^i, \kappa_1\widehat{\pi}_t^i\right\}\right) di}{\int_{\widehat{\pi}_t^i \geq \ln\left(\frac{\overline{\psi} w_t^*}{P_{t-1}^*}\right)} \exp\left(\min\left\{\frac{\pi_t^* + \ln(\gamma_t)}{1-\rho} - \kappa_3\widehat{\pi}_t^i, \kappa_2\widehat{\pi}_t^i\right\}\right) di}\right) \\ & + (1-\rho)\,\kappa_2\left[g + \frac{\kappa_1\sigma^2}{2} + \frac{(1-\rho)\kappa_3}{(\theta-\rho)\kappa_1\kappa_2}\ln\left(\frac{1}{\rho(1-\alpha)}\right) + \frac{1}{\kappa_1}\ln\Phi\left(\kappa_1\sigma - \frac{\ln\left(\frac{\overline{\psi} w_t^*}{P_{t-1}^*}\right) - g}{\sigma}\right)\right] \end{aligned} \tag{13}$$

Now use $\widehat{\pi}_t^i = g + \varepsilon_{it}$ and $\varepsilon_{it} \sim N\left(0, \sigma^2\right)$ with equation (9) to simplify (13). It will be useful to define the following:

$$\begin{aligned} \Upsilon\left(X,\sigma\right) \quad \equiv \quad & \left[\kappa_1\kappa_2\left(1-\rho\right)^2 - \frac{(1+\rho)(\theta-\rho)^2(1-\alpha)^2}{[\theta(1-\alpha\rho)-\rho(1-\alpha)]^2}\right]\frac{\sigma^2}{2(1-\rho)} \\ & + \frac{(1-\alpha)(1-\rho)}{\theta-\rho(1-\alpha)}\left[\ln\left(\frac{1}{\rho(1-\alpha)}\right) + \frac{\theta\alpha}{1-\alpha}\ln\Phi\left(\kappa_1\sigma - \frac{X}{\sigma}\right)\right] + \ln\left(\frac{\Phi\left(-\rho\kappa_3\sigma - \frac{X}{\sigma}\right)}{\Phi\left(-\kappa_3\sigma - \frac{X}{\sigma}\right)}\right) \end{aligned} \tag{14}$$

If $X \geq \Upsilon\left(X,\sigma\right) + \ln\left(\gamma_t\right)$, then $\pi_t^* - g = \Upsilon\left(X,\sigma\right)$. Otherwise:

$$
\begin{aligned}
\pi_t^* - g \quad = \quad & \left[\kappa_1^2 - \kappa_2^2 + (1-\rho)\,\kappa_1\kappa_2\right] \tfrac{\sigma^2}{2} + \tfrac{(1-\alpha)(1-\rho)}{\theta-\rho(1-\alpha)} \left[\ln\left(\tfrac{1}{\rho(1-\alpha)}\right) + \tfrac{\theta\alpha}{1-\alpha} \ln\Phi\left(\kappa_1\sigma - \tfrac{X}{\sigma}\right)\right] \\
& + \ln\left(\frac{\Phi\left(\kappa_1\sigma - \frac{X}{\sigma}\right) - \Phi\left(\kappa_1\sigma - \frac{\ln(\gamma_t)+\pi_t^*-g}{\sigma}\right) + \exp\left(\left(\rho^2\kappa_3^2-\kappa_1^2\right)\frac{\sigma^2}{2} + \frac{\ln(\gamma_t)+\pi_t^*-g}{1-\rho}\right)\Phi\left(-\rho\kappa_3\sigma - \frac{\ln(\gamma_t)+\pi_t^*-g}{\sigma}\right)}{\Phi\left(\kappa_2\sigma - \frac{X}{\sigma}\right) - \Phi\left(\kappa_2\sigma - \frac{\ln(\gamma_t)+\pi_t^*-g}{\sigma}\right) + \exp\left(\left(\kappa_3^2-\kappa_2^2\right)\frac{\sigma^2}{2} + \frac{\ln(\gamma_t)+\pi_t^*-g}{1-\rho}\right)\Phi\left(-\kappa_3\sigma - \frac{\ln(\gamma_t)+\pi_t^*-g}{\sigma}\right)}\right)
\end{aligned}
\tag{15}
$$

Either way, we have a definition of $\pi_t^* - g$ which is independent of g. $\square$

Part 2 Follows directly from Proposition 1. $\square$

Part 3 The fraction of firms not operating is $\Delta \equiv \Phi\left(\frac{X}{\sigma}\right)$. Taking derivatives yields $\frac{d\Delta}{d\sigma} \propto \frac{dX}{d\sigma} - \frac{X}{\sigma}$ so what we want to show is $\frac{dX}{d\sigma} > \frac{X}{\sigma}$. Using equation (11) from the proof of Part 1 above:

$$
\frac{dX}{d\sigma} = \kappa_1\sigma + \frac{\frac{X}{\sigma}}{1 + \frac{\Phi\left(\kappa_1\sigma - \frac{X}{\sigma}\right)}{\phi\left(\kappa_1\sigma - \frac{X}{\sigma}\right)}\kappa_1\sigma}
\tag{16}
$$

The desired inequality is thus $\left(\kappa_1\sigma - \frac{X}{\sigma}\right)\frac{\Phi\left(\kappa_1\sigma - \frac{X}{\sigma}\right)}{\phi\left(\kappa_1\sigma - \frac{X}{\sigma}\right)} > -1$. Using $x\Phi(x) > -\phi(x)$ as shown next completes the proof: $x\Phi(x) = x\int_{-\infty}^{x} \phi(t)\,dt > \int_{-\infty}^{x} t\phi(t)\,dt = -\int_{-\infty}^{x} \phi'(t)\,dt = -\phi(x)$. $\square$

Part 4 If $\alpha = 1$, then the equations in Part 1 reduce to:

$$
X = \tfrac{1-\rho}{\rho}\ln\left(\tfrac{U}{1-\rho}\right) + \tfrac{\sigma^2}{2(1-\rho)} + (1-\rho)\ln\Phi\left(\tfrac{\sigma}{1-\rho} - \tfrac{X}{\sigma}\right)
\tag{17}
$$

$$
\pi_t^* - g = \tfrac{\sigma^2}{2(1-\rho)} + (1-\rho)\ln\Phi\left(\tfrac{\sigma}{1-\rho} - \tfrac{X}{\sigma}\right)
\tag{18}
$$

We can thus write $f(\sigma) = X - \frac{1-\rho}{\rho}\ln\left(\frac{U}{1-\rho}\right)$ with X dependent on σ as per (17). To make this dependency explicit, we further write $X(\sigma)$ in place of just X. Consider any $\sigma_0 > 0$ satisfying $f(\sigma_0) = 0$. That is, consider any $\sigma_0 > 0$ satisfying $X(\sigma_0) = \frac{1-\rho}{\rho}\ln\left(\frac{U}{1-\rho}\right)$. If $U \geq 1-\rho$, then $X(\sigma_0) \geq 0$ which, given $\frac{dX}{d\sigma} > \frac{X}{\sigma}$ from the proof of Part 3, implies $X'(\sigma_0) > 0$. Notice

$f'(\cdot) = X'(\cdot)$. This means that, if $U \geq 1-\rho$, then any $\sigma_0 > 0$ satisfying $f(\sigma_0) = 0$ must also satisfy $f'(\sigma_0) > 0$. There is therefore at most one $\sigma_0 > 0$ such that $f(\sigma_0) = 0$. To show that there is exactly one such $\sigma_0 > 0$, it will be sufficient to show $\lim_{\sigma\to 0^+} f(\sigma) < 0$ and $\lim_{\sigma\to\infty} f(\sigma) > 0$. Equation (17) yields $X(0) \equiv \lim_{\sigma\to 0^+} X(\sigma) = (1-\rho)\left[\frac{1}{\rho}\ln\left(\frac{U}{1-\rho}\right) + \ln\Phi\left(\lim_{\sigma\to 0^+}\frac{-X(\sigma)}{\sigma}\right)\right]$. Notice that $X(0) > 0$ is impossible while $X(0) < 0$ is only possible if $U < 1-\rho$. Therefore, $U \geq 1-\rho$ implies $X(0) = 0$ and thus $\lim_{\sigma\to 0^+} f(\sigma) = -\frac{1-\rho}{\rho}\ln\left(\frac{U}{1-\rho}\right) < 0$. Equation (17) also yields $\lim_{\sigma\to\infty} X(\sigma) = \infty$ and thus $\lim_{\sigma\to\infty} f(\sigma) = \infty$. Putting everything together, we can now conclude that there is exactly one $\sigma_0 > 0$ such that $f(\sigma_0) = 0$. Moreover, $f'(\sigma_0) > 0$. $\square$

Part 5 Define $\tau(U) \equiv \frac{1}{\rho}\ln\left(\frac{\theta U}{\theta(1-\alpha\rho)-\rho(1-\alpha)}\right) + \frac{1-\alpha}{\theta-\rho(1-\alpha)}\ln\left(\frac{1}{\rho(1-\alpha)}\right)$ and $\widetilde{U} \equiv \tau^{-1}(0)$. Also define $\overline{U} \equiv 1-\alpha\rho-\frac{\rho(1-\alpha)}{\theta}$ and focus on $U \leq \overline{U}$. If θ is finite, then $\widetilde{U} < \overline{U}$ and it will suffice to establish the result for some subset of $\left(\widetilde{U}, \overline{U}\right)$. We proceed with this case before turning to θ infinite. At $\sigma = 0$, equation (11) yields $X(0) \equiv \lim_{\sigma\to 0^+} X(\sigma) = \frac{1}{\kappa_2}\left[\tau(U) + \frac{\theta\alpha}{\theta-\rho(1-\alpha)}\ln\Phi\left(\lim_{\sigma\to 0^+}\frac{-X(\sigma)}{\sigma}\right)\right]$. Notice that $X(0) > 0$ is impossible while $X(0) < 0$ is only possible if $\tau(U) < 0$ or, equivalently, $U < \widetilde{U}$. Therefore, $U > \widetilde{U}$ implies $X(0) = 0$ and thus $m \equiv \lim_{\sigma\to 0^+}\frac{-X(\sigma)}{\sigma} = \Phi^{-1}\left(\exp\left(-\frac{[\theta-\rho(1-\alpha)]\tau(U)}{\theta\alpha}\right)\right)$. Using equation (16) yields $\lim_{\sigma\to 0^+}\frac{1}{\sigma}\left[X'(\sigma) - \frac{X(\sigma)}{\sigma}\right] = \left(1+\frac{\Phi(m)m}{\phi(m)}\right)\kappa_1$. Using equations (14) and (15) then yields $f(0) \equiv \lim_{\sigma\to 0^+} f(\sigma) = \frac{1-\rho}{\rho}\ln\left(\frac{\overline{U}}{U}\right)$ and $f'(0) \equiv \lim_{\sigma\to 0^+} f'(\sigma) = \frac{(\theta-\rho)(1-\alpha)}{\theta\overline{U}}\left[\frac{\phi(m)}{\Phi(m)} - \frac{\theta\alpha(1-\rho)}{(\theta-\rho)(1-\alpha)}m\right]$. At this point, it will be instructive to consider $U = \overline{U}$ and hence $f(0) = 0$. Define $\eta \equiv [\rho(1-\alpha)]^{\frac{1-\alpha}{\theta\alpha}}$. If $\frac{\theta\alpha(1-\rho)}{(\theta-\rho)(1-\alpha)} > \frac{\phi\left(\Phi^{-1}(\eta)\right)}{\eta\Phi^{-1}(\eta)}$, then $f'(0) < 0$. By properties of the standard normal, $\frac{\phi\left(\Phi^{-1}(\eta)\right)}{\eta\Phi^{-1}(\eta)}$ is decreasing in η for $\eta > 0.5$. Moreover, $\frac{d\eta}{d\alpha} \propto \frac{1}{\alpha}\ln\left(\frac{1}{\rho(1-\alpha)}\right) - 1$ which is positive for α sufficiently large. Since $\frac{\theta\alpha(1-\rho)}{(\theta-\rho)(1-\alpha)}$ is increasing in α, it then follows that $f'(0) \ll 0$ if α is above some threshold $\underline{\alpha} \in (0,1)$. Next, notice that $f(0)$, m, and $f'(0)$ are continuous in U, with $\frac{df(0)}{dU} < 0$, $\frac{dm}{dU} > 0$, and $\frac{df'(0)}{dU} < 0$. Therefore, there exists an $\epsilon > 0$ such that $f(\sigma_A) = 0$ and $f'(\sigma_A) < 0$ for some $\sigma_A > 0$ if $\alpha \in (\underline{\alpha}, 1)$ and $U \in \left(\underline{U}, \overline{U}\right)$, where $\underline{U} \equiv \max\left\{\widetilde{U}, \overline{U}-\epsilon\right\}$. That there is a $\sigma_B > \sigma_A$ such that $f(\sigma_B) = 0$ and $f'(\sigma_B) > 0$ follows from $\lim_{\sigma\to\infty} f(\sigma) = \infty$. Turn now to $\theta \to \infty$. In this case, $\widetilde{U} = \overline{U} = 1-\alpha\rho$ so we will establish the result for some subset of $\left(0,\overline{U}\right)$. At $U=\overline{U}$, we have $X(0)=0$, $\lim_{\sigma\to 0^+}\frac{-X(\sigma)}{\sigma}=\infty$, $f(0)=0$, and $f'(0)=\frac{\alpha(1-\rho)}{1-\alpha\rho}X'(0)$. Since $X'(0)\equiv\lim_{\sigma\to 0^+}X'(\sigma)\cong\frac{X(h)-X(0)}{h-0}=\frac{X(h)}{h}\xrightarrow{h\to 0^+}-\infty$, it follows that $f'(0)<0$. Taken together, $f(0)=0$ and $f'(0)<0$ imply existence of a $\sigma>0$ such that $f(\sigma)<0$. Combined with $\lim_{\sigma\to\infty}f(\sigma)=\infty$, this then implies existence of a $\sigma_B>0$ satisfying $f(\sigma_B)=0$ and $f'(\sigma_B)>0$. For $f(\cdot)$ continuous in U, we can thus find an $\widehat{\epsilon}>0$ such that there also exists a $\sigma_B>0$ satisfying $f(\sigma_B)=0$ and $f'(\sigma_B)>0$ when $U\in\left(\overline{U}-\widehat{\epsilon},\overline{U}\right)$. Now, for any $U\in\left(0,\overline{U}\right)$, we have $X(0)\in(-\infty,0)$, $f(0)=0$, and $f'(0)=\frac{\alpha(1-\rho)}{1-\alpha\rho}X'(0)$, where $X'(0)=\frac{1-\alpha\rho}{X(0)}\lim_{\sigma\to 0^+}\left(\frac{X(\sigma)}{\sigma}\right)^2\phi\left(-\frac{X(\sigma)}{\sigma}\right)=0$. Also notice $\Upsilon(X(0),0)=0$ so, for small taste shocks, the properties of $f''(\cdot)$ around zero are dictated by equation (15). After some algebra (available upon request), we obtain $f''(0)\equiv\lim_{\sigma\to 0^+}f''(\sigma)=\frac{(1-\alpha)(\alpha-\rho)+\alpha(1-\rho)^2}{(1-\alpha\rho)^2(1-\rho)}$ which is positive for $\alpha\in\left(1-\frac{\rho(1-\rho)}{2}-\sqrt{\left(1-\frac{\rho(1-\rho)}{2}\right)^2-\rho},1\right)\equiv(\underline{\widehat{\alpha}},1)$. Therefore, there must exist a $\sigma_A\in(0,\sigma_B)$ satisfying $f(\sigma_A)=0$ and $f'(\sigma_A)<0$ when $U\in\left(\overline{U}-\widehat{\epsilon},\overline{U}\right)$ and $\alpha\in(\underline{\widehat{\alpha}},1)$. ■

Proof of Proposition 3

To simplify notation, define $\lambda\equiv\frac{1}{\rho}\ln\left(\frac{U}{1-\rho}\right)$. If $\alpha=1$, then the mixture equations reduce to:

$$\exp(y_t)=\xi_t\exp\left(\frac{\upsilon_A^2}{2}\right)\Phi\left(\upsilon_A-\frac{y_t+\lambda}{\upsilon_A}\right)+(1-\xi_t)\exp\left(\frac{\upsilon_B^2}{2}\right)\Phi\left(\upsilon_B-\frac{y_t+\lambda}{\upsilon_B}\right)\tag{19}$$

Under $\xi_t=0$, equation (19) yields $\exp(y_t)=\exp\left(\frac{\upsilon_B^2}{2}\right)\Phi\left(\upsilon_B-\frac{y_t+\lambda}{\upsilon_B}\right)$. Under $\xi_t=1$, it yields $\exp(y_t)=\exp\left(\frac{\upsilon_A^2}{2}\right)\Phi\left(\upsilon_A-\frac{y_t+\lambda}{\upsilon_A}\right)$. Since $y_t=f(\sigma_B)=0$ at $\xi_t=0$ and $y_t=f(\sigma_A)=0$ at $\xi_t=1$, it follows that $\exp\left(\frac{\upsilon_A^2}{2}\right)\Phi\left(\upsilon_A-\frac{\lambda}{\upsilon_A}\right)=1$ and $\exp\left(\frac{\upsilon_B^2}{2}\right)\Phi\left(\upsilon_B-\frac{\lambda}{\upsilon_B}\right)=1$. Consider now $\xi_t\in(0,1)$. If $y_t<0$, then $\exp\left(\frac{\upsilon_i^2}{2}\right)\Phi\left(\upsilon_i-\frac{y_t+\lambda}{\upsilon_i}\right)>\exp\left(\frac{\upsilon_i^2}{2}\right)\Phi\left(\upsilon_i-\frac{\lambda}{\upsilon_i}\right)=1$ for $i\in\{A,B\}$ so equation (19) implies $y_t>0$ which is a contradiction. If $y_t>0$, then $\exp\left(\frac{\upsilon_i^2}{2}\right)\Phi\left(\upsilon_i-\frac{y_t+\lambda}{\upsilon_i}\right)<\exp\left(\frac{\upsilon_i^2}{2}\right)\Phi\left(\upsilon_i-\frac{\lambda}{\upsilon_i}\right)=1$ for $i\in\{A,B\}$ so equation (19) implies $y_t<0$ which is a contradiction. Therefore, $y_t=0$ for $\xi_t\in(0,1)$. ■