NBER WORKING PAPER SERIES

INFLATION ANNOUNCEMENTS AND SOCIAL DYNAMICS

Kinda Hachem
Jing Cynthia Wu

Working Paper 20161
http://www.nber.org/papers/w20161

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2014

We thank Veronica Guerrieri, Anil Kashyap, and Randy Kroszner for helpful suggestions. We have also benefited from discussions with Craig Burnside, Jon S. Cohen, Steve Davis, Jim Hamilton, Chris Hansen, Narayana Kocherlakota, Jim Pesando, Ricardo Reis, Michael Woodford, and seminar participants at the Federal Reserve Board, Chicago Booth, UBC, FRB New York, FRB Chicago, Duke, SED 2013, the 2013 NBER Summer Institute EFSF Workshop, and ASSA 2014. 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. Correspondence: kinda.hachem@chicagobooth.edu and cynthia.wu@chicagobooth.edu. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Inflation Announcements and Social Dynamics
Kinda Hachem and Jing Cynthia Wu
NBER Working Paper No. 20161
May 2014
JEL No. E17,E3,E58

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 find that trying to eliminate high inflation by abruptly introducing low inflation targets generates short-term overshooting. Gradual targets, in contrast, achieve a smoother disinflation. We present empirical evidence to support these predictions. Gradualism is not equally effective in other situations though: our model predicts aggressive announcements are more powerful when combating deflation.

Kinda Hachem
University of Chicago
Booth School of Business
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
kinda.hachem@chicagobooth.edu

Jing Cynthia Wu
University of Chicago
Booth School of Business
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
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. This paper proposes a new framework for understanding the effectiveness of central bank announcements.

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)). Such preferences have since subsided – see, for example, Woodford (2005) and Blinder et al. (2008) – so public information about policy goals is now a relevant baseline. Interestingly though, increased transparency has not 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 help answer such questions, models of the expectation formation process have been called for by prominent policy-makers (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. Existing 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 headline result is 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. We then present some new empirical evidence that corroborates the correlation between abruptness and overshooting predicted by our model.

Our next set of results concerns strategies for eliminating deflation. We start by showing that gradualism is actually less effective here: the central bank can eliminate deflation more quickly by communicating an aggressive increase in its short-term inflation goals. Price-level targeting may thus have some communication-based benefits during a deflation. We then show how two dimensions of quantitative easing – number of rounds and intensity of announcements – can be used to guide the economy out of deflation without explicit changes in short-term targets. While much of the QE literature has focused on yield curves,[1] our

[1]See, for example, Gagnon et al. (2010), D'Amico and King (2010), Williams (2011), Hamilton and Wu (2012), and the references therein.

model speaks to the inflation channel in Krishnamurthy and Vissing-Jorgensen (2011).

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 explains the evolution of credibility through social dynamics, Section 3 builds a model of inflation determination for use in simulations, Sections 4 and 5 present the simulation results, and Section 6 concludes. All proofs are collected in the Appendix.

2 Expectation Formation via Social Dynamics

Inflation expectations are an important equilibrium object in our model: they respond to realized inflation through social dynamics and affect realized inflation through price-setting.

[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).

This section explains how we use social dynamics to capture the evolution of inflation expectations. The next section then elaborates on the price-setting environment.

Consider a continuum of agents $i \in [0, 1]$. At the beginning of date t, agent i expects an inflation rate $\widehat{\pi}_t^i$. This expectation is drawn from a forecasting rule, with the choice of rule subject to social dynamics. Two mean rational rules – that is, two rules which yield unbiased forecasts when adopted by all agents – will suffice to expound these dynamics. In the short-run, selection between competing alternatives is the core of social dynamics so having at least two rules is important. In the long-run, selection algorithms typically converge to one alternative so focusing on mean rationality ensures convergence to an unbiased rule.

2.1 Forecasting Rules

Based on existing evidence, central bank announcements and random walk forecasts are natural candidates for our two forecasting rules: Atkeson and Ohanian (2001) find that random walk forecasts of inflation perform very well against more sophisticated statistical models, Faust and Wright (2012) find that the Fed's Greenbook forecasts are difficult to beat, and Gurkaynak et al. (2005) and Campbell et al. (2012) find that markets do indeed view FOMC statements as a source of new and reliable information about future economic conditions. While the viability of each rule will be determined endogenously, emergence of these approaches in the real world motivates their inclusion in the option set.

With at least some pass-through from professional forecasts to individual expectations (e.g., Carroll (2003b)), data from the Survey of Professional Forecasters (SPF) can be used to further discipline our forecasting rules. Of particular relevance is a special question on the 2012Q2 SPF which asked respondents whether their forecasts were consistent with the Fed's inflation target. Forecasters who self-identified as consistent formed a tight distribution around the Fed's target while the remaining forecasters formed a wider distribution around past inflation.[4] Let $\overline{\pi}_t$ denote the central bank's date t announcement and let π_{t-1}^* denote

[4]Raw data is available at www.phil.frb.org/research-and-data/real-time-center/survey-of-professional-forecasters/2012/survq212.cfm

realized inflation at the end of date $t-1$. Also define variances σ_A^2 and σ_B^2 such that $\sigma_A < \sigma_B$. The two forecasting rules can now be formally stated:

Definition 1 *Agent i is a Fed Follower (FF) at date t if* $\widehat{\pi}_t^i \sim N\left(\overline{\pi}_t, \sigma_A^2\right)$

Definition 2 *Agent i is a Random Walker (RW) at date t if* $\widehat{\pi}_t^i \sim N\left(\pi_{t-1}^*, \sigma_B^2\right)$

In our definitions, $N\left(\mu, \sigma^2\right)$ means a forecasting rule based primarily on μ but subject to individual interpretation as per σ. Writing $\widehat{\pi}_t^i \sim N\left(\mu, \sigma^2\right)$ then means a single draw from this rule. Fed Followers will thus form a tighter distribution around $\overline{\pi}_t$ than Random Walkers will form around π_{t-1}^*. Going forward, the fraction of Fed Followers is denoted by $\xi_t \in [0, 1]$. If $\xi_t = 1$, then everyone adopts the FF rule and mean rationality implies $\pi_t^* = \overline{\pi}_t$. In other words, the central bank can achieve inflation goals using only communication. We thus interpret ξ_t as a measure of central bank credibility.

2.2 Social Dynamics and the Evolution of Forecasts

To endogenize the evolution of ξ_t, we posit that agents whose forecasts are consistently outperformed by their peers will want to change how they forecast. Social dynamics then provides a useful approach to modeling how an agent discovers he is being outperformed and how much of this outperformance he attributes to one-time shocks rather than fundamentals. Initializing $\xi_0 = 0$, we let ξ_{t+1} evolve via tournament selection and mutation.

Tournament selection simulates information transmission in a complex world. Versions of the approach appear in Carroll (2003a), Acemoglu, Ozdaglar and ParandehGheibi (2010), Arifovic et al. (2010), Arifovic et al. (2012), and Burnside et al. (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, agents meet in pairs and compare forecast errors after the realization of π_t^*. Consider a meeting between agents i and j. Agent i counts one strike against his forecasting rule if j used a different rule and $\left|\widehat{\pi}_t^j - \pi_t^*\right| < \left|\widehat{\pi}_t^i - \pi_t^*\right|$. When ξ_t is sufficiently high, π_t^* will be close to $\overline{\pi}_t$ so $\left|\overline{\pi}_t - \pi_{t-1}^*\right| \gg 0$

will allow Fed Followers to outperform Random Walkers in many meetings. Strikes will thus tend to be counted against the RW rule, suggesting $\xi_{t+1} \geq \xi_t$. In contrast, when ξ_t is sufficiently low, π_t^* will be close to π_{t-1}^* so $|\overline{\pi}_t - \pi_{t-1}^*| \gg 0$ will instead allow Random Walkers to outperform Fed Followers in many meetings. Strikes will thus tend to be counted against the FF rule, suggesting $\xi_{t+1} \leq \xi_t$. This is the sense in which success is endogenous.

How quickly do strikes lead to $\xi_{t+1} \neq \xi_t$? Experimental evidence suggests that agents 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 agents to accumulate several strikes before deciding to switch forecasting rules. This is the sense in which success is judged over multiple observations. We use S to denote the number of strikes needed for a switch: after S strikes, agent 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 simulate our model for different values and different accumulation rules. As Section 4.3 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 agent can have zero to many meetings in a given period. Drawing from the full population ensures that even agents who do not participate in economic activity are represented in tournaments. This is appealing since participation decisions are driven by expectations. In an alternative specification, we allow tournaments to occur locally rather than at random: agents lie along a circle and each agent meets his right and left neighbors every period. With interactions set up as such, agents always meet the same people. As we will see in Section 4, this creates clusters of agents that use the same forecasting rule. Agents at the center of a cluster are thus more likely to meet other agents using the same rule, increasing their *effective* stubbornness for any value of S. These results suggest 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 $\theta \in (0,1)$ of agents randomly switches rules regardless of strikes. The timing of our social forces can now be summarized as follows: (i) mutation turns the fraction of FFs into $\widetilde{\xi}_t = (1-\theta)\,\xi_t + \theta\,(1-\xi_t)$ if $t \geq 1$; (ii) each agent i draws expectation $\widehat{\pi}_t^i$ from his 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 3; (iv) tournament selection transforms ξ_t into ξ_{t+1} if $t=0$ and $\widetilde{\xi}_t$ into ξ_{t+1} if $t \geq 1$.

3 An Expectations-Based Model of Inflation

We now present a formal model to map inflation expectations into realized inflation. Interpret the continuum of agents in Section 2 as a continuum of firms, each producing a differentiated perishable good $i \in [0,1]$. The demand for good i in date t is $D_{it} = \left(\frac{\gamma_t P_t}{p_{it}}\right)^{\frac{1}{1-\rho}}$, where p_{it} is the price charged by firm i, 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 $\rho \in (0,1)$ so that the demand for each good is decreasing in its relative price $\frac{p_{it}}{P_t}$ and increasing in the taste shock γ_t. The supply of good i is then given by the production function $F(\ell_{it}) = \ell_{it}^{\alpha}$, where $\alpha \in (0,1)$ is a constant and ℓ_{it} is the labor input used by firm i. The aggregate stock of labor is normalized to one and inelastically available at unit wage w_t. As an alternative to production, each firm also has an exogenous outside option with real value $U > 0$. A firm is said to operate if and only if it does not take its outside option. We will elaborate on the choice of $\alpha < 1$ and the role of γ_t at the end of Subsection 3.1. The introduction of $U > 0$ is then discussed in Subsection 3.2.

3.1 Timing and Equilibrium

Firms have to make pricing and production decisions before P_t and γ_t are realized – that is, before they know the actual demand for their products. This assumption creates a role for expectations. In particular, at the beginning of date t, each firm i 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. Firm i will ultimately draw $\widehat{\pi}_t^i$ from one of the two forecasting rules described in Section 2 but, for now, we can imagine any draw that does not change within date t. To simplify the exposition and remain focused on inflation expectations, assume that all firms also forecast a taste shock of one. In other words, the forecasting rule for γ_t is simply the mean of the shock.

Individual Decisions Conditional on its forecast $\widehat{P}_t^i$ and the prevailing wage w_t, firm i chooses p_{it} and ℓ_{it} to solve a static profit maximization problem. Charging p_{it} for good i yields an expected demand of $\left(\frac{\widehat{P}_t^i}{p_{it}}\right)^{\frac{1}{1-\rho}}$ which in turn necessitates $\left(\frac{\widehat{P}_t^i}{p_{it}}\right)^{\frac{1}{\alpha(1-\rho)}}$ units of labor. Stated in real terms, firm i's problem is thus:

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

From the inner maximization problem, the pricing decision of an operating firm is:

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

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

$$O_t\left(w_t\right) = \left\{ i \mid \widehat{P}_t^i \geq \frac{1}{\alpha\rho} \left(\frac{U}{1-\alpha\rho}\right)^{\frac{1-\alpha\rho}{\alpha\rho}} w_t \right\} \tag{2}$$

If $i \notin O_t\left(w_t\right)$, 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:

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

Notice that operating firms with higher price expectations charge higher prices. They also hire more labor, resulting in more output. For any given wage, firms with higher expectations are also 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.

Wage Determination Given the individual decisions above, the wage is set to clear the labor market. More precisely, an auctioneer chooses w_t^* to solve $\int_{O_t(w_t^*)} \ell\left(w_t^*; \widehat{P}_t^i\right) di = 1$, with $O_t\left(w_t\right)$ and $\ell\left(w_t; \widehat{P}_t^i\right)$ as per equations (2) and (3) respectively. This yields:

$$w_t^* = \alpha\rho \left[\int_{O_t(w_t^*)} \left(\widehat{P}_t^i\right)^{\frac{1}{1-\alpha\rho}} di \right]^{1-\alpha\rho} \tag{4}$$

Following the determination of w_t^*, each firm $i \in O_t\left(w_t^*\right)$ posts price $p_{it}^* \equiv p\left(w_t^*; \widehat{P}_t^i\right)$ and hires labor to produce its expected demand $q_{it}^* \equiv \left(\frac{\widehat{P}_t^i}{p_{it}^*}\right)^{\frac{1}{1-\rho}}$. Price expectations are not updated based on w_t^* so, in this sense, the model deviates from rational expectations (RE). While some updating can certainly be accommodated, there must be residual heterogeneity for social dynamics to operate at the end of the period.

Realized Inflation After p_{it}^* and q_{it}^* have been set, the taste shock γ_t is realized and the aggregate price level is computed as a consumption-weighted average of individual prices. At price level P_t, the realized demand for good i is $\left(\frac{\gamma_t P_t}{p_{it}^*}\right)^{\frac{1}{1-\rho}}$ which may differ from the available supply q_{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 $c_{it} = \min\left\{q_{it}^*, \left(\frac{\gamma_t P_t^*}{p_{it}^*}\right)^{\frac{1}{1-\rho}}\right\}$.[5] The result is:

$$P_t^* \int_{O_t(w_t^*)} \left(\frac{\min\left\{\widehat{P}_t^i, \gamma_t P_t^*\right\}}{p\left(w_t^*; \widehat{P}_t^i\right)}\right)^{\frac{1}{1-\rho}} di = \int_{O_t(w_t^*)} \left(\frac{\min\left\{\widehat{P}_t^i, \gamma_t P_t^*\right\}}{p\left(w_t^*; \widehat{P}_t^i\right)}\right)^{\frac{1}{1-\rho}} p\left(w_t^*; \widehat{P}_t^i\right) di \tag{5}$$

with $p\left(\cdot\right)$, $O_t\left(\cdot\right)$, and w_t^* as per equations (1), (2), and (4) respectively. We now have a mapping from a set of price expectations $\left\{\widehat{P}_t^i \mid i \in [0,1]\right\}$ to the realized price level P_t^*. Recalling $\widehat{P}_t^i \equiv \exp(\widehat{\pi}_t^i) P_{t-1}^*$ and invoking $P_t^* \equiv \exp(\pi_t^*) P_{t-1}^*$, the mapping from a set of

[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.

inflation expectations $\left\{\widehat{\pi}_t^i \mid i \in [0,1]\right\}$ to realized inflation π_t^* is straightforward.[6]

Discussion We now elaborate on some of the modeling elements used above, namely the role of γ_t and the choice of $\alpha < 1$. 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." The principle of ordinary-language rationality motivates the timing of our taste shock. Without γ_t in equation (5), the auctioneer could compute P_t^* at the same time as w_t^*. 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^*. The choice of $\alpha < 1$ is similarly motivated. Notice from equation (1) that $\alpha = 1$ prompts all firms to set the same price – namely a constant mark-up over the wage – regardless of expectations. This has advantages and disadvantages. On one hand, it helps isolate how expectations affect realized inflation through just the labor market but, on the other, it eliminates the lag between w_t^* and P_t^* in the auctioneer's problem. With $\alpha < 1$, prices are not simple mark-ups so γ_t enters (5) and the lag is restored. Importantly, this restoration stems from production being non-linear in labor, not from production being decreasing returns to scale.[7] Our full model thus employs $\alpha < 1$, presenting $\alpha = 1$ as an intermediate step only when it helps unpack the results.

3.2 Results with One Forecasting Rule

Before combining the model of Subsection 3.1 with the two forecasting rules in Section 2, it will be instructive to establish how the model works with one forecasting rule:

Proposition 1 *If $\widehat{\pi}_t^i \sim N(\mu, \sigma^2)$ for all i, then:*

[6] As equations (4) and (5) 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.

[7] In a previous version, we showed that the key properties of our model hold when production is instead given by $F(\ell_{it}, z_{it}) = \ell_{it}^{\alpha} z_{it}^{1-\alpha}$, where z_{it} is firm effort and the real disutility of such effort is $\frac{z_{it}^{\lambda}}{\lambda}$ with $\lambda \in (1, \infty]$. Notice that $F(\ell_{it}) = \ell_{it}^{\alpha}$ is just the limiting case of $\lambda \to \infty$.

1. $\pi_t^* = \mu + f(\sigma)$

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

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* $U \in (\underline{U}, \overline{U})$ *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 1 says that excess inflation, $\pi_t^* - \mu$, depends only on the extent of expectations heterogeneity. In the long-run, however, we should not observe excess inflation: if everyone draws inflation forecasts from $N(\mu, \sigma^2)$ but realized inflation ends up being $\pi_t^* \neq \mu$, then $N(\mu, \sigma^2)$ is a biased forecasting rule and its survival into the long-run would seem at odds with ordinary-language rationality. The rest of Proposition 1 establishes conditions under which heterogeneity ($\sigma > 0$) and unbiasedness ($\pi_t^* = \mu$) are consistent.

To fix ideas, consider the limiting case of $\alpha = 1$. As per Subsection 3.1, prices will be a constant mark-up over the wage and realized inflation will be determined in the labor market. We can thus focus on equation (4). Absent an outside option, all firms will operate so $O_t(w_t) = [0, 1]$ for any w_t. In other words, the operating set will be independent of the wage. The second part of Proposition 1 reveals that realized inflation will exceed the mean expectation in this case. The excess, $\frac{\sigma^2}{2(1-\rho)}$, is a Jensen's inequality term which arises whenever normal inflation expectations are compounded into log-normal price expectations and aggregated.[8] Mathematically, restoring unbiasedness without eliminating heterogeneity 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

[8]The Jensen's term is exacerbated by more heterogeneity in inflation expectations (i.e., higher σ) and more substitutability between goods (i.e., higher ρ). As σ increases, the compounding process 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.

expectation firms. Allowing low expectation firms to not produce by introducing a positive outside option is one way to achieve this, motivating our use of $U > 0$.[9]

The remainder of Proposition 1 restricts attention to positive outside options. The third part establishes that more heterogeneity in expectations decreases operation. 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 helps offset the Jensen's inequality effect. Indeed, with linear production, the fourth part of Proposition 1 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(a).[10] 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 1 says that our model actually produces two mean rational points.

As discussed in Section 2, the goal is to have two mean rational forecasting rules so we shall proceed with the restrictions at the end of Proposition 1: $\alpha \in (\underline{\alpha}, 1)$ and $U \in \left(\underline{U}, \overline{U}\right)$. By way of example, the blue dots in Figure 1(b) show the combinations of α and U that satisfy these restrictions when $\rho = 0.9$. For any such combination, Proposition 1 says the graph of $f(\cdot)$ will resemble the blue line in Figure 1(a). Notice from this line that $f(\cdot)$ is negative between the two mean rational points. In other words, excess inflation is negative. To understand why, recall the competing effects of higher σ on wages in equation (4). As σ increases, 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, low expectation firms find it more profitable to take their outside option and the resulting decline in operation

[9] In principle, one could instead assume normality (rather than log-normality) of price expectations. In practice though, prices cannot be negative so the appropriate assumption would be truncated normality which, like log-normality, is not symmetric.

[10] The ability of exit to counter Jensen's inequality does not hinge on inflation expectations being single draws from a normal distribution: if each firm were to treat its forecasting rule as a full prior, the proof of Part 4 shows that exit can be restored with two forecasting rules and appropriate bounds on U.

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. We know from equation (1) that individual prices respond positively to wages so, all else constant, the shape of w_t^* feeds into P_t^*. Notice, however, that $\alpha < 1$ implies 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 (5) 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. If the outside option is too low, then exit is weak and we get no mean rational points.

3.3 Results with Two Forecasting Rules

Consider now the determination of π_t^* when expectations are drawn from different forecasting rules. In particular, a group of size ξ_t forecasts according to $\widehat{\pi}_t^i \sim N\left(\mu, \sigma_A^2\right)$ and a group of size $1-\xi_t$ forecasts according to $\widehat{\pi}_t^i \sim N\left(\mu, \sigma_B^2\right)$. There are two dimensions of heterogeneity here: within group ($\sigma_A > 0$ and $\sigma_B > 0$) and across groups ($\sigma_A \neq \sigma_B$). We will also relax homogeneity of μ in our simulations but, for now, imagine $\overline{\pi}_t = \pi_{t-1}^* = \mu$ in Definitions 1 and 2. This approximates an economy which approaches its central bank's inflation target with a lot of expectations heterogeneity. To simplify notation, define:

$$y_t \equiv \frac{\pi_t^* - \mu}{1-\alpha\rho},\ x_t \equiv \frac{\ln\left(\frac{1}{\alpha\rho}\left(\frac{U}{1-\alpha\rho}\right)^{\frac{1-\alpha\rho}{\alpha\rho}} \frac{w_t^*}{P_{t-1}^*}\right) - \mu}{1-\alpha\rho},\ \text{and}\ \upsilon_j \equiv \frac{\sigma_j}{1-\alpha\rho}$$

The variable y_t is just excess inflation scaled up by a constant. The variable x_t provides a more compact way to express the operation condition in equation (2). 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. With a mixture of normal expectations, equations (2) and (4) yield x_t implicitly defined by:[11]

$$x_t = \frac{1}{\alpha\rho}\ln\left(\frac{U}{1-\alpha\rho}\right) + \ln\left[\xi_t \exp\left(\frac{\upsilon_A^2}{2}\right)\Phi\left(\upsilon_A - \frac{x_t}{\upsilon_A}\right) + (1-\xi_t)\exp\left(\frac{\upsilon_B^2}{2}\right)\Phi\left(\upsilon_B - \frac{x_t}{\upsilon_B}\right)\right]$$

The expression for excess inflation then comes from equation (5). If the realized taste shock is small enough to support $\gamma_t P_t^* \leq \widehat{P}_t^i$ for all operating firms, then $\gamma_t P_t^*$ will drop out of the consumption weights, leaving P_t^* and y_t^* explicitly defined. Otherwise, P_t^* and y_t^* will be implicitly defined. The threshold γ_t works out to:

$$\Upsilon(x_t, \xi_t) \equiv \frac{\left[\xi_t \exp\left(\frac{\upsilon_A^2}{2}\right)\Phi\left(\upsilon_A - \frac{x_t}{\upsilon_A}\right) + (1-\xi_t)\exp\left(\frac{\upsilon_B^2}{2}\right)\Phi\left(\upsilon_B - \frac{x_t}{\upsilon_B}\right)\right]^{1-\alpha}\left(\frac{U}{1-\alpha\rho}\right)^{\frac{1-\alpha\rho}{\alpha\rho}}}{\frac{\xi_t \exp\left(\frac{\left(\frac{\rho(1-\alpha)\upsilon_A}{1-\rho}\right)^2}{2}\right)\Phi\left(-\frac{\rho(1-\alpha)\upsilon_A}{1-\rho} - \frac{x_t}{\upsilon_A}\right) + (1-\xi_t)\exp\left(\frac{\left(\frac{\rho(1-\alpha)\upsilon_B}{1-\rho}\right)^2}{2}\right)\Phi\left(-\frac{\rho(1-\alpha)\upsilon_B}{(1-\rho)} - \frac{x_t}{\upsilon_B}\right)}{\xi_t \exp\left(\frac{\left(\frac{(1-\alpha)\upsilon_A}{1-\rho}\right)^2}{2}\right)\Phi\left(-\frac{(1-\alpha)\upsilon_A}{1-\rho} - \frac{x_t}{\upsilon_A}\right) + (1-\xi_t)\exp\left(\frac{\left(\frac{(1-\alpha)\upsilon_B}{1-\rho}\right)^2}{2}\right)\Phi\left(-\frac{(1-\alpha)\upsilon_B}{1-\rho} - \frac{x_t}{\upsilon_B}\right)}}$$

If $\gamma_t \leq \Upsilon(x_t, \xi_t)$, then $y_t = x_t - \frac{\ln \Upsilon(x_t,\xi_t)}{1-\alpha\rho}$. Otherwise, y_t solves:

$$y_t = \frac{\alpha(1-\rho)}{1-\alpha\rho}\left[x_t - \frac{1}{\alpha\rho}\ln\left(\frac{U}{1-\alpha\rho}\right)\right] + \frac{1}{1-\alpha\rho}\ln\left(\frac{\xi_t h\left(x_t, y_t, \gamma_t, \upsilon_A, 1, \frac{\rho(1-\alpha)}{1-\rho}\right) + (1-\xi_t) h\left(x_t, y_t, \gamma_t, \upsilon_B, 1, \frac{\rho(1-\alpha)}{1-\rho}\right)}{\xi_t h\left(x_t, y_t, \gamma_t, \upsilon_A, \alpha, \frac{1-\alpha}{1-\rho}\right) + (1-\xi_t) h\left(x_t, y_t, \gamma_t, \upsilon_B, \alpha, \frac{1-\alpha}{1-\rho}\right)}\right)$$

where

$$\begin{aligned} h(x_t, y_t, \gamma_t, \upsilon, \beta, \delta) \equiv\; & \exp\left(\frac{(\beta\upsilon)^2}{2}\right)\left[\Phi\left(\beta\upsilon - \frac{x_t}{\upsilon}\right) - \Phi\left(\beta\upsilon - \frac{y_t + \frac{\ln(\gamma_t)}{1-\alpha\rho}}{\upsilon}\right)\right] \\ & + \exp\left(\frac{(\delta\upsilon)^2}{2} + \frac{(1-\alpha\rho)\left(y_t + \frac{\ln(\gamma_t)}{1-\alpha\rho}\right)}{1-\rho}\right)\Phi\left(-\delta\upsilon - \frac{y_t + \frac{\ln(\gamma_t)}{1-\alpha\rho}}{\upsilon}\right) \end{aligned}$$

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

Proposition 2 *Suppose $\widehat{\pi}_t^i \sim N(\mu, \sigma_A^2)$ for a group of size ξ_t and $\widehat{\pi}_t^i \sim N(\mu, \sigma_B^2)$ 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 2 says that the entire population is mean rational when each subpopulation is individually mean rational. We know from equation (1) that all firms set

[11] The derivations that follow parallel those in the proof of Proposition 1, 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 same price when $\alpha = 1$ so any heterogeneity in expectations only affects the economy through labor market clearing, namely equation (4). The latter aggregates linearly across subpopulations so, if the component distributions are each parameterized to deliver $\pi_t^* = \mu$, then their mixture will also deliver $\pi_t^* = \mu$.

In contrast, Figure 2 shows what can happen when $\alpha < 1$ also introduces heterogeneity into the price aggregation of equation (5): the mixture distribution produces negative excess inflation even if each subpopulation possesses the mean rational property. Using $\rho = 0.9$ 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(a): if one uses $N(\mu, \sigma_x^2)$ with some $\sigma_x \in (\sigma_A, \sigma_B)$ to approximate the aggregate distribution generated by $\xi_t \in (0, 1)$ and $\overline{\pi}_t = \pi_{t-1}^* = \mu$, then $\pi_t^* < \mu$ follows for the reasons in Subsection 3.2.

4 Introducing Inflation Targets

From the economic model of Section 3, we can calculate π_t^* conditional on ξ_t and the forecasting rules. Using social dynamics as per Section 2, 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 to achieve a large disinflation in this environment. The end of this section then presents some new empirical evidence that supports our core prediction.

4.1 Parameterization

We set $\rho = 0.9$ which captures high but imperfect substitutability between goods. To obtain two individually mean rational subpopulations, we pick α and U from the blue region in Figure 1(b). We set $\alpha = 0.9$ and $U = 0.18$ but any choice from the aforementioned region will deliver qualitatively similar results. Lastly, we assume that the taste shock is uniformly distributed according to $\gamma_t \sim \mathcal{U}[0.99, 1.01]$. These parameter choices deliver $\sigma_A = 0.0036$ and $\sigma_B = 0.0643$ as the solutions to $f(\cdot) = 0$. Substituting σ_A and σ_B into Definitions 1 and 2 completes the characterization of our forecasting rules. The mutation parameter is then set to $\theta = 0.02$ which is conservative compared to the social dynamics literature.

4.2 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 3 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, $\xi_0 = 0$ so all firms are Random Walkers who forecast according to $\widehat{\pi}_t^i \sim N(20\%, \sigma_B^2)$. The mean rational property thus yields $\pi_0^* = 20\%$.

Consider first a central bank that introduces its target abruptly, announcing $\overline{\pi}_t = 2\%$ for all $t \geq 1$. The bank's announcement introduces a new forecasting rule, Fed Following, which a small fraction of firms mutate towards. Figure 3(a) demonstrates that inflation converges to 2% but is followed by a temporary overshooting of the target. Recall from Section 3 that firms with low expectations (relative to their peers) are less likely to operate. Fed Followers 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 3.3, generates $\pi_t^* < 2\%$. Over time though, RWs accumulate enough strikes to compel them to become FFs, returning inflation to target.

Figure 3(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 Fed Followers on the way down to 2%. In turn, the economy is very close to a situation where all firms forecast according to $\widehat{\pi}_t^i \sim N\left(2\%, \sigma_A^2\right)$ when 2% is actually reached, resulting in $\pi_t^* = 2\%$.

4.3 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 investigate less stubborn beliefs. Figure 4(a) 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.

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 2, 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 5 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. Panels (b) and (c) in Figure 4 then show 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.

Negative Strikes Up to this point, we have only allowed for the accumulation of strikes against one's own forecasting rule. What if in between accumulating strikes an agent has several meetings where his rule outperforms the other? To investigate whether this changes our results we now allow for the de-accumulation of strikes. Firms still switch after S strikes and still add a strike when outperformed by a different forecasting rule but now they also subtract a strike when their rule is the outperformer. The results for abrupt introduction are shown in Figure 4(d). We have used $S = 8$ and random interactions so comparing against Figure 3(a) will isolate the effect of negative strikes. As realized inflation falls, the fraction of FFs starts rising sooner and the 2% target is reached more quickly in the version with negative strikes. However, once at 2%, the fraction of FFs does not accelerate as quickly as it did absent these strikes so overshooting is more pronounced. In essence, the introduction of negative strikes creates more effective stubbornness when inflation is near target because both forecasting rules have a chance to outperform and de-accumulate prior strikes.

Discussion The picture painted by our alternative specifications is two-fold. First, richer specifications such as local interactions and negative strikes matter insofar as they change effective stubbornness during the transition. Second, stubbornness has advantages and disadvantages for a central bank trying to achieve a large disinflation. On one hand, higher stubbornness among FFs yields faster and more certain convergence to the bank's target but, on the other, higher stubbornness among RWs leads to a temporary overshooting of the target if the target is introduced abruptly. This dichotomy prompts a policy tradeoff when stubbornness is high enough: abrupt targets will lower inflation quickly but create overshooting; gradual targets will eliminate overshooting but bring inflation down more slowly.

4.4 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 panels (a) and (b) of Figure 6 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, panels (c) and (d) of Figure 6 compare the full dynamics from Figure 3 against the results that would arise under only mutation. With just mutation, $\xi_{t+1} = (1 - \theta)\, \xi_t + \theta\, (1 - \xi_t)$ 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.

4.5 Empirical Evidence

As described above, an important prediction of our model is that abrupt targets lead to temporary overshooting while gradual targets do not. We now demonstrate that this prediction is indeed borne out in the data. Our point of departure is the set of 29 inflation targeters in Roger (2009) and Svensson (2010). Using these authors' dates for the formal adoption of inflation targeting in each country, we collect data on the path of inflation targets from individual central bank websites. Our data on actual inflation then come from the IMF's International Financial Statistics database, as per Mishkin and Schmidt-Hebbel (2007). The last observation in our sample is 2013Q3. To conduct the analysis, we exclude any countries that either began with an inflation target above actual inflation or eventually abandoned inflation targeting. This leaves us with the 19 countries listed in Table 1.

We next construct some intuitive dummies for abruptness in a given country. The first, `abrupt1`, equals one if the targeted path is flat. This is the least subjective measure of abruptness so we adopt it as the baseline. The second dummy, `abrupt2`, equals one if the absolute value of the net change in inflation targets between the time of introduction and the end of our sample is less than 40% of the absolute value of the net change in realized inflation over the same period. The third dummy, `abrupt3`, equals one if the standard error of the inflation targets that make up the targeted path is smaller than the final target. The left panel in Table 1 shows the division of countries between abrupt and gradual according to each measure. By definition, any country with `abrupt1` = 1 will also have `abrupt2` = 1 and `abrupt3` = 1. To help visualize the data, Figure 7 compares the average inflation rates and targeted paths for abrupt versus gradual targeters. The top panel averages over countries with `abrupt1` = 1 while the bottom panel averages over countries with `abrupt1` =

0. Averages are for each point in time, with time 0 denoting the introduction of inflation targets. The plots line up quite well with the simulations in Figure 3: overshooting is more characteristic of abrupt inflation targeters.

For more formal evidence, we run some simple qualitative cross-country regressions. The benefit of this approach is that it is model-free and scale-free: the results are less susceptible to outliers and thus more robust for small samples. To construct a dummy variable for overshooting, let t_1 denote the first quarter in which actual inflation hits the targeted path from above and let $t_2 > t_1$ denote the first quarter in which actual inflation hits the final target from below.[12] Our main dependent variable, `overshoot1`, is a dummy that equals one if average inflation is less than the average target over the period t_1 to t_2. As an alternative, we also define `overshoot2` which equals one if average inflation is less than the average target for the period t_1 to 2013Q3. The right panel in Table 1 shows the division of countries between overshooting and no overshooting according to each measure.

The first two panels in Table 2 report the results of our qualitative regressions. The first three columns in the left panel regress `overshoot1` on `abrupt1`, `abrupt2`, and `abrupt3` respectively. The middle panel repeats the exercise using `overshoot2` as the dependent variable. The intercepts in our regressions are generally small and statistically insignificant whereas the abruptness coefficients are generally large and significant. On the whole, the range of intercepts suggests that the probability of overshooting is 0-29% for gradual targeters while the range of intercepts plus slopes suggests that this probability is 46-100% for abrupt targeters. The difference is economically and statistically significant, consistent with our model's prediction that abrupt targets are associated with overshooting. To check the robustness of our empirical results, we also run regressions that control for the distance between initial and desired inflation (`d2target`). Comparing the first and fourth columns in each panel of Table 2 reveals that the abruptness coefficient is largely unchanged, both quantitatively and qualitatively. In contrast, the coefficient on `d2target` is always small

[12] If actual inflation never hit the final target from below, we set t_2 equal to the last quarter in our sample (i.e., 2013Q3). If actual inflation also never hit the targeted path from above, we set t_1 equal to the first quarter in which the final target became effective.

and insignificant. Therefore, even controlling for initial conditions, there is evidence that overshooting is more characteristic of abrupt targeters.

To complement the qualitative results, the last panel in Table 2 presents some quantitative cross-country regressions. Our dependent variable, `overshoot_num`, is defined as the average difference between actual and targeted inflation over the period t_1 to t_2. Notice that `overshoot1` $= 1$ if and only if `overshoot_num` < 0 so overshooting will now be indicated by negative regression coefficients. The intercepts in the first three columns of this last panel measure the average overshooting among gradual targeters whereas the sum of the intercepts and abruptness coefficients measure the average overshooting among abrupt targeters. The intercepts are positive but generally insignificant whereas the abruptness coefficients are negative and generally significant. Classifying abruptness according to `abrupt1`, the first column says that actual inflation will average 0.97 percentage points below an abrupt target. The last column then says this result is robust to initial condition controls. The quantitative results thus align with the qualitative ones, underscoring our model's prediction.

5 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 4.1 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.

5.1 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, the top row of Figure 8 shows that the accumulation of FFs – first via mutation then via tournament selection – eventually brings the economy up to 2%.[13]

What happens if the central bank instead decides to gradually lead the economy back to 2%? As the middle row of Figure 8 reveals, a very gradual strategy involves more persistent deflation early on. Recall from Subsection 3.1 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, the bottom row of Figure 8 shows that aggressive communication may be the most effective at eliminating deflation. The path we consider is one where the central bank announces short-term targets well above its 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 3(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, this suggests that price-level targeting – which would indeed require periods of high inflation to balance out periods of deflation – has some advantages over inflation targeting when dealing with deflations and somewhat stubborn beliefs.

5.2 Quantitative Easing

Up to this point, the direct effect of central bank announcements has been limited to the mean of the FF distribution. We now consider more potent announcements which can di-

[13]Note that convergence to 2% would now occur 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% will be reached.

rectly affect the skewness of the economy-wide distribution. A practical example is what Krishnamurthy and Vissing-Jorgensen (2011) dub the inflation channel of QE: curbing deflationary expectations through the Fed's widely publicized 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, 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 to 2% without changing its short-term targets. Figure 9 illustrates how this can be achieved in our baseline specification ($S = 8$, random interactions) by varying rounds and intensity.

To isolate the effect of rounds, the top two rows of Figure 9 fix intensity at 1: all firms with deflationary expectations redraw. The first row demonstrates that one round of QE announcements helps increase inflation but more time is needed to accumulate FFs and reach 2%. The second row 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.

The third row of Figure 9 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 3(a).

Finally, the fourth row of Figure 9 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 the second row requires stopping QE right when inflation hits its target. At the same time, avoiding the dip below 2% experienced in the third row 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.

6 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 abruptly introducing a low inflation target leads to 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.

Empirical evidence was then presented to support the correlation between abruptness and overshooting predicted by our model. Our next set of results concerned communications to guide the economy away from deflation. We found that combating deflation requires either aggressive announcements that are broadly consistent with price-level targeting or QE-type announcements that stem deflationary expectations without changing inflation targets.

References

Acemoglu, Daron, Asuman Ozdaglar, and Ali ParandehGheibi, "Spread of (Mis)Information in Social Networks," *Games and Economic Behavior*, 2010, *70* (2), pp. 194–227.

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.

Roger, Scott, "Inflation Targeting at 20: Achievements and Challenges," 2009. IMF Working Paper WP/09/236.

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.

Svensson, Lars E.O., "Inflation Targeting," 2010. NBER Working Paper 16654.

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.

Table 1: Data

Country	abrupt1	abrupt2	abrupt3	overshoot1	overshoot2
Canada	0	1	1	1	1
Chile	0	0	0	0	0
Colombia	0	0	0	0	0
Czech Republic	0	1	1	1	1
Ghana	1	1	1	1	1
Guatemala	0	1	1	0	0
Hungary	0	0	1	0	0
Iceland	1	1	1	1	0
Indonesia	0	0	1	0	0
Israel	0	0	0	0	0
Mexico	0	0	0	0	0
New Zealand	0	1	1	0	0
Norway	1	1	1	1	1
Peru	0	0	0	1	0
Romania	0	0	1	0	0
South Africa	1	1	1	1	0
Sweden	1	1	1	1	1
Turkey	0	0	0	0	0
United Kingdom	0	0	1	1	1

Notes: The left panel classifies countries as abrupt versus gradual targeters based on the abruptness measures defined in Subsection 4.5. The right panel then codes whether or not the targeted path was overshot based on the overshooting measures in the same subsection.

Table 2: Regressions

	overshoot1				overshoot2				overshoot_num			
intercept	0.29**	0.20	0.17	0.33**	0.21*	0.10	0.00	0.32*	0.39	0.66**	0.40	0.15
	(0.02)	(0.16)	(0.41)	(0.05)	(0.10)	(0.47)	(1.00)	(0.07)	(0.15)	(0.04)	(0.40)	(0.67)
abrupt1	0.71***			0.68**	0.39			0.28	-1.36**			-1.14*
	(0.00)			(0.01)	(0.12)			(0.29)	(0.02)			(0.05)
abrupt2		0.58***				0.46**				-1.33***		
		(0.01)				(0.03)				(0.01)		
abrupt3			0.45*				0.46**				-0.53	
			(0.08)				(0.05)				(0.35)	
d2target				0.00				-0.01				0.02
				(0.70)				(0.33)				(0.33)

Notes: Regression coefficients with p-values in brackets. * denotes rejection at the 10% level, ** at the 5% level, and *** at the 1% level. All variables are as defined in Subsection 4.5.

Figure 1: One Forecasting Rule

(a) Illustration of Proposition 1

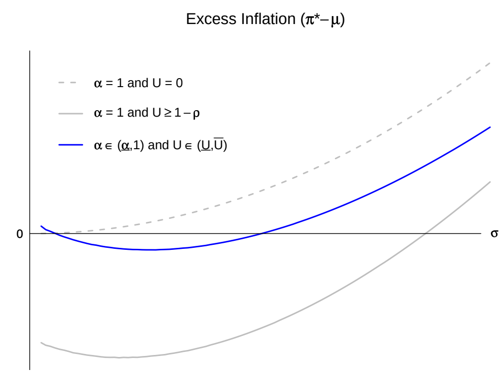

(b) Parameter Space Example

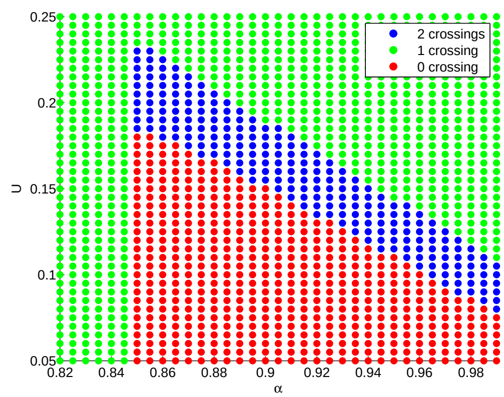

Figure 2: Two Forecasting Rules

(a) Parameter Space Example

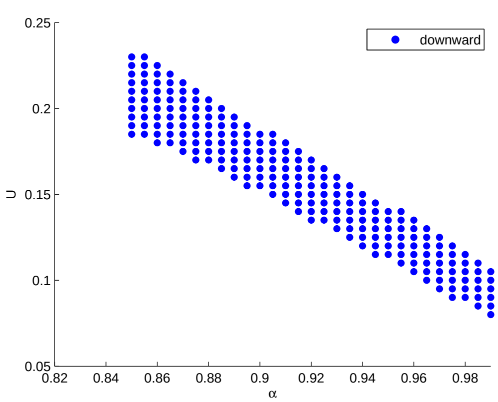

(b) Dependence of y_t on ξ_t

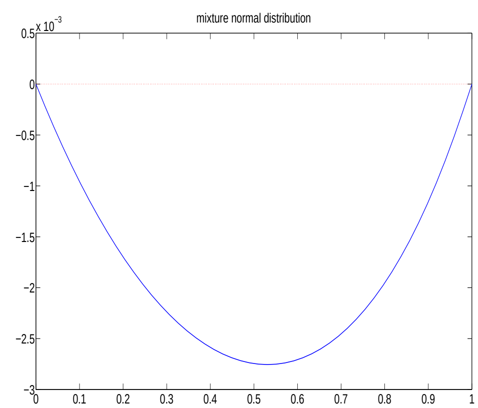

Figure 3: Introduction of IT, Baseline Specification

(a) Abrupt Strategy

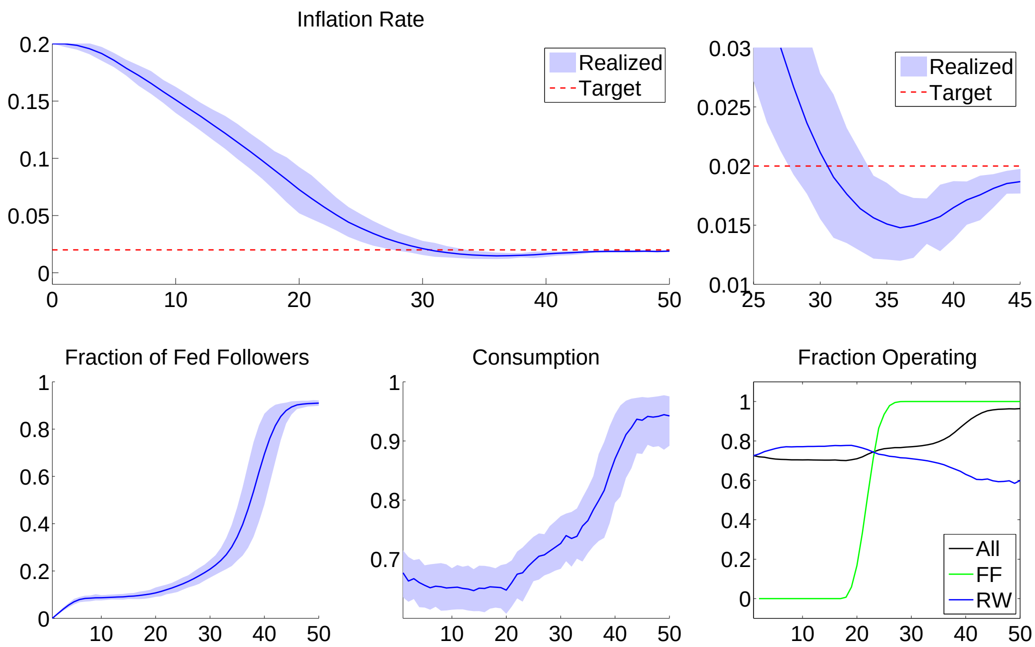

(b) Gradual Strategy

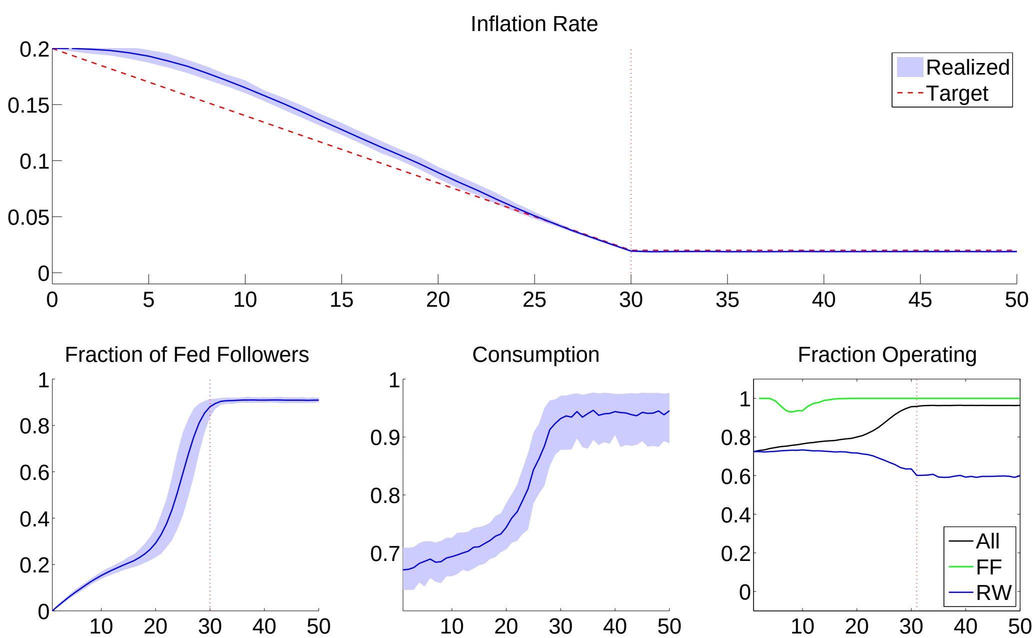

Figure 4: Alternative Specifications

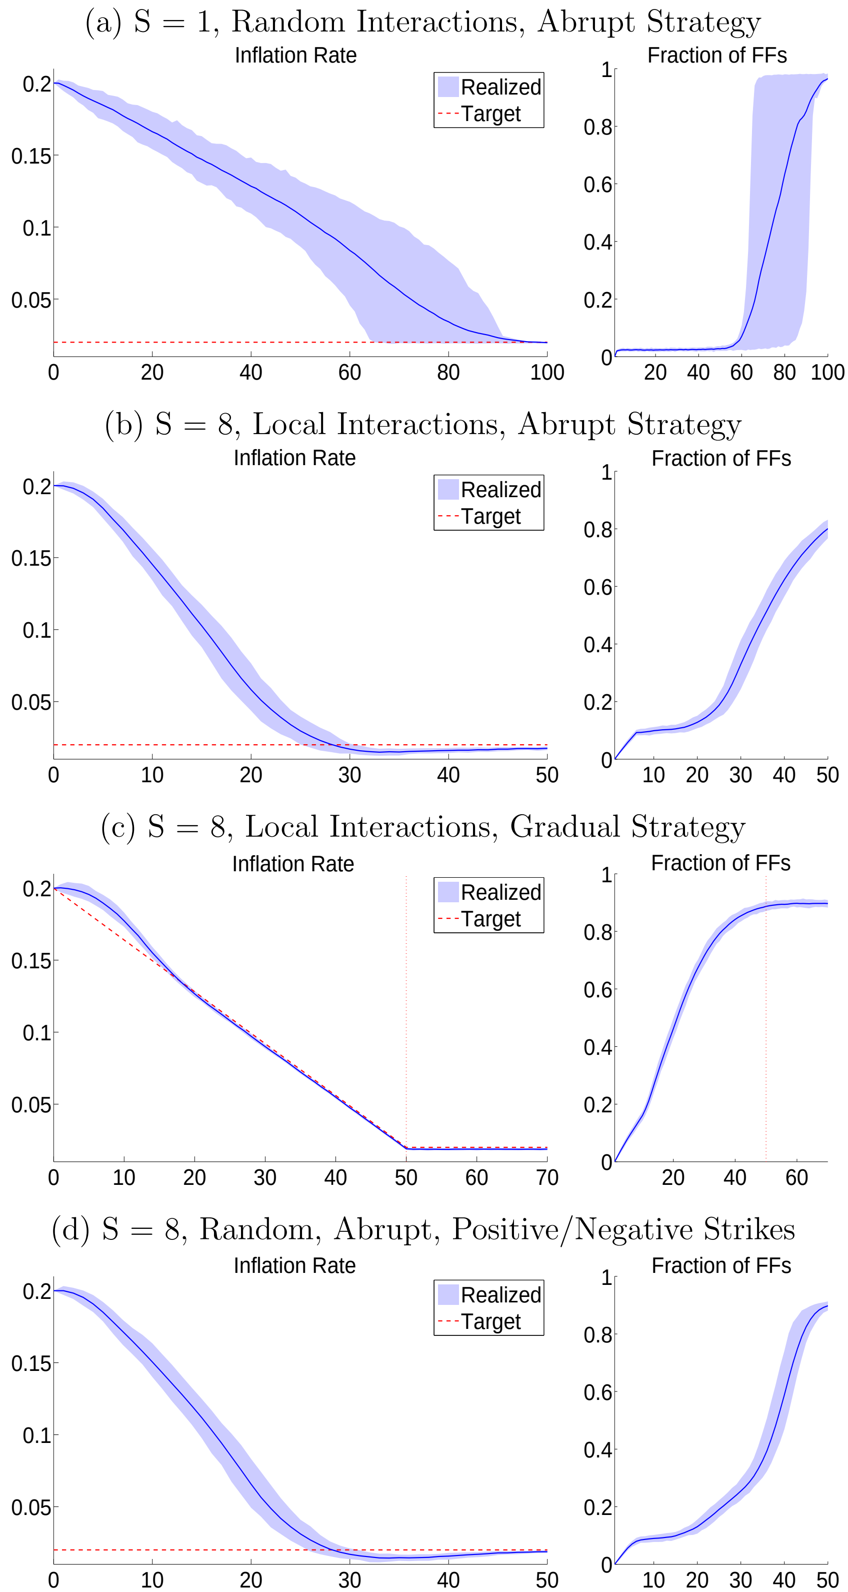

Figure 5: FF Accumulation for Different Interaction Types (S = 8)

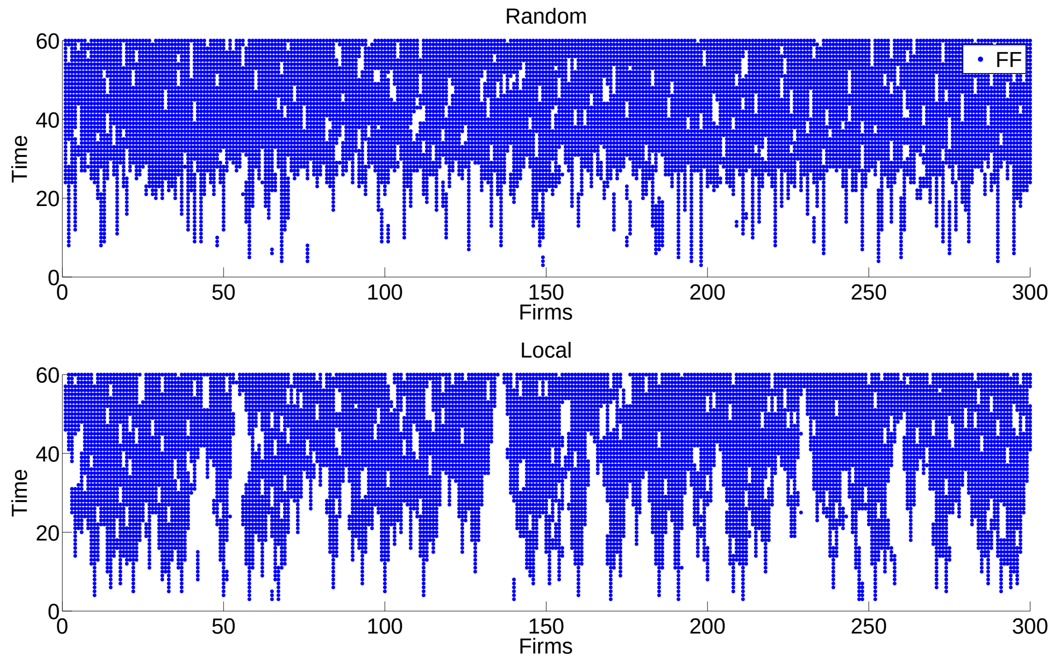

Figure 6: Benchmarks

(a) Baseline vs Fixed Proportions, Abrupt Strategy

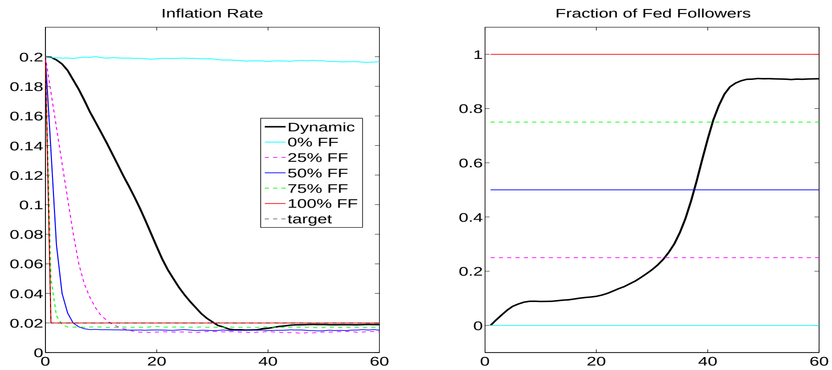

(b) Baseline vs Fixed Proportions, Gradual Strategy

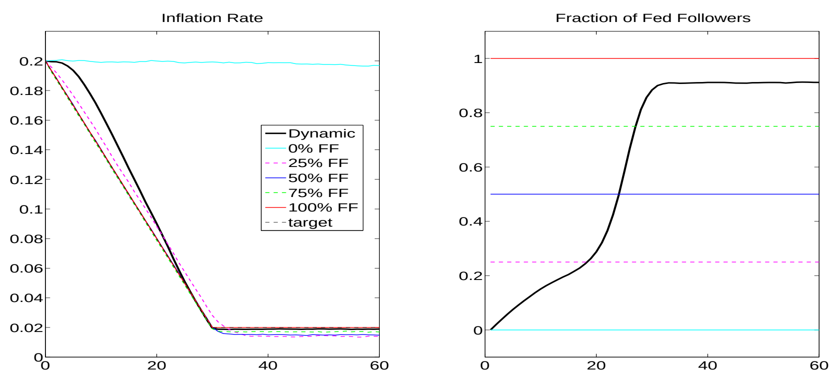

(c) Baseline vs Mutation-Only, Abrupt Strategy

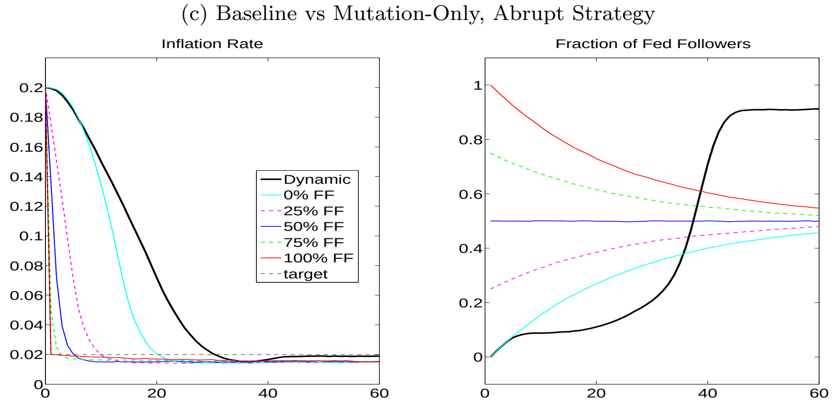

(d) Baseline vs Mutation-Only, Gradual Strategy

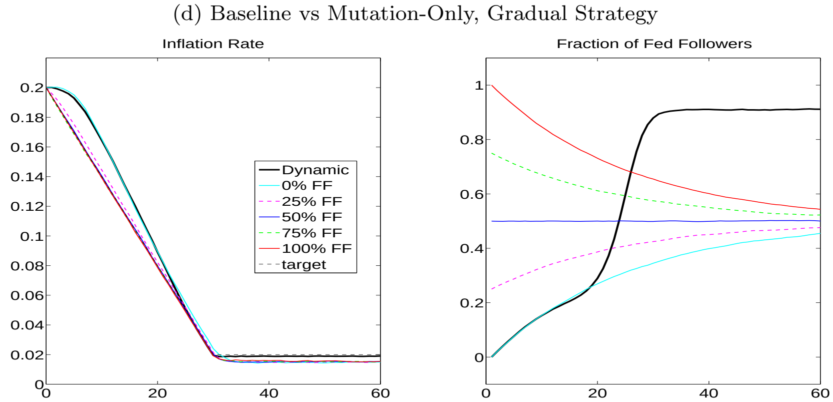

Figure 7: Abrupt vs Gradual Targeters in the Data

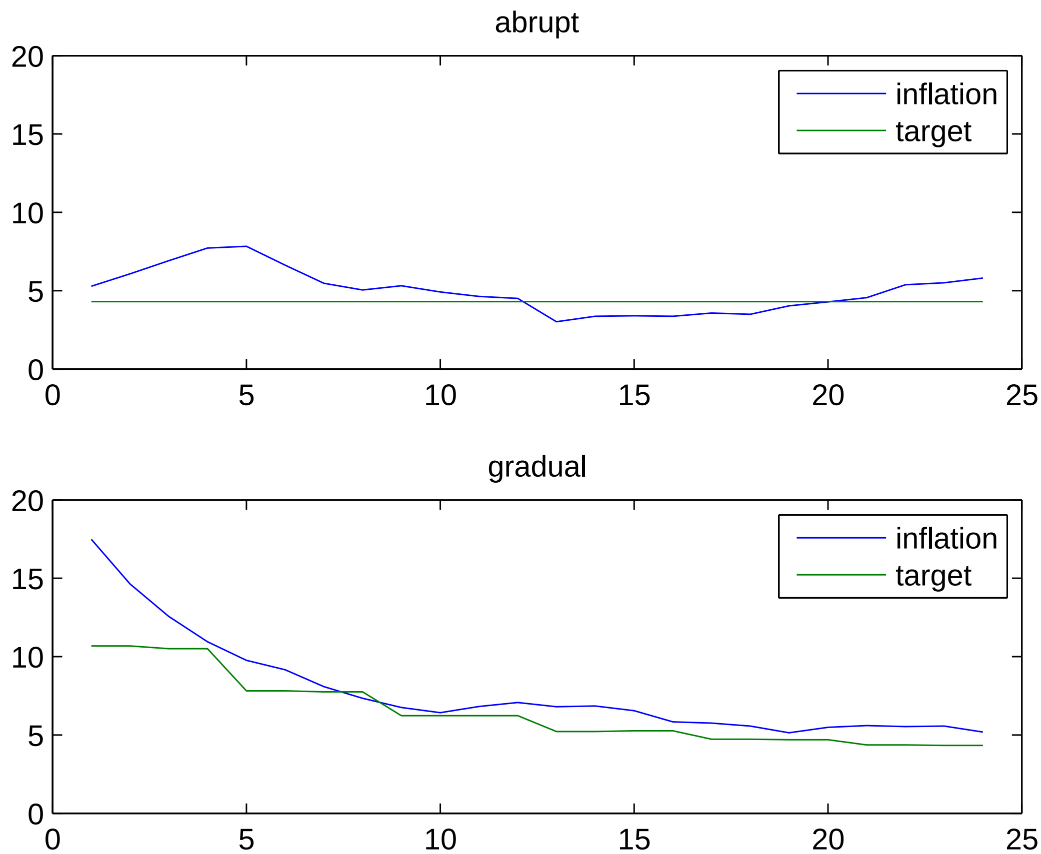

Notes: Average inflation targets and average realized inflation. The horizontal axis is time in quarters since IT adoption. The top panel averages over abrupt targeters (`abrupt1` = 1) while the bottom panel averages over gradual targeters (`abrupt1` = 0).

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

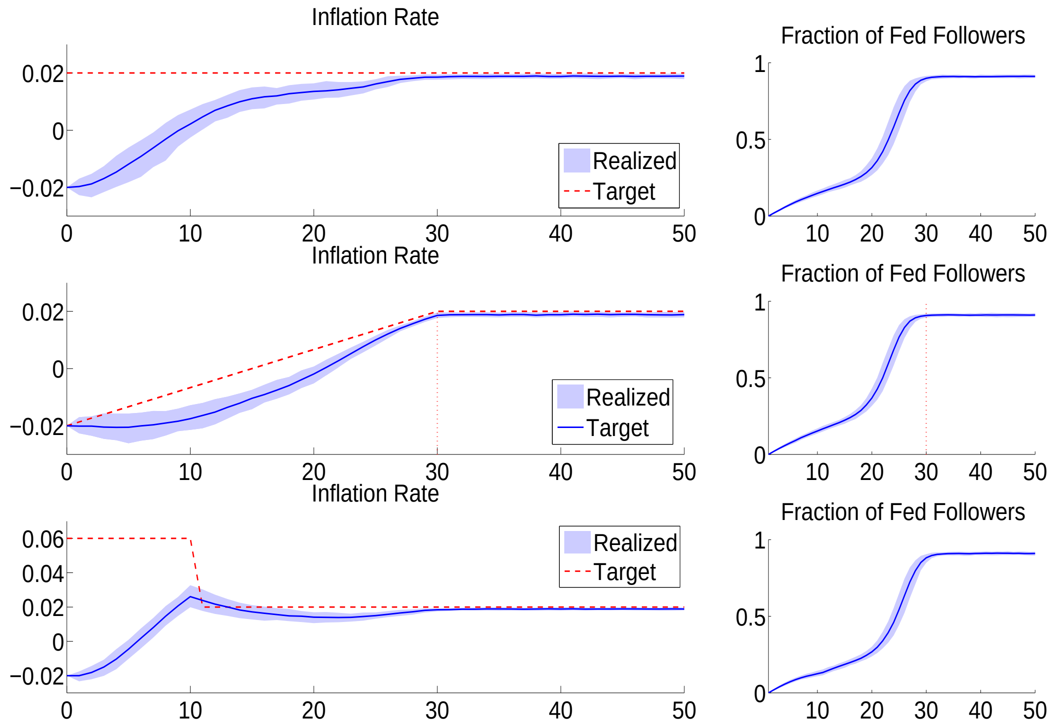

Notes: The top row depicts a flat target, the middle row depicts a gradual strategy, and the bottom row depicts an aggressive strategy.

Figure 9: QE Announcements

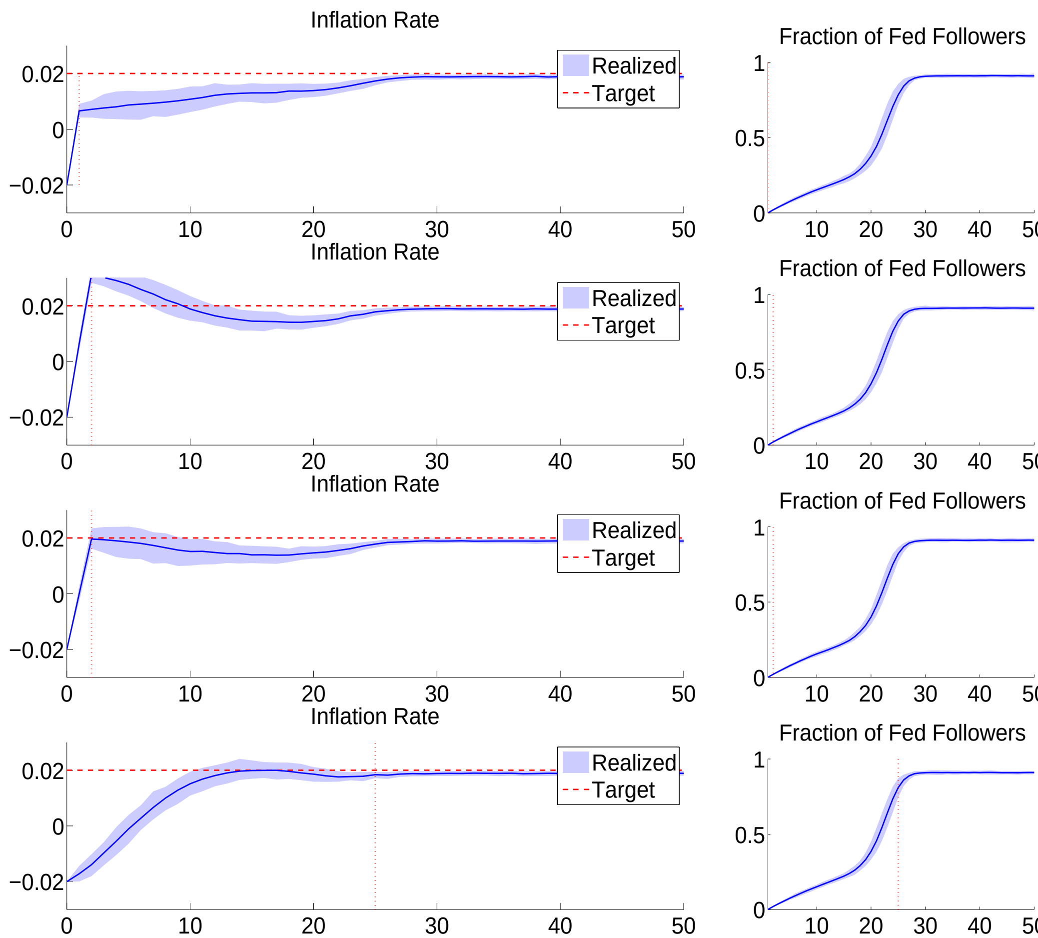

Notes: Different specifications of (Rounds, Intensity). The first row depicts (1,1), the second row depicts (2,1), the third row depicts (2,0.75), and the fourth row depicts (25,0.1).

Appendix - Proofs

Proof of Proposition 1

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:

$$\int_{x \geq c} \exp(rx)\,\phi\left(x,\sigma^2\right) dx = \exp\left(\tfrac{r^2\sigma^2}{2}\right)\Phi\left(r\sigma - \tfrac{c}{\sigma}\right) \tag{6}$$

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

$$\varepsilon_{it} \geq \ln\left(\tfrac{1}{\alpha\rho}\left(\tfrac{U}{1-\alpha\rho}\right)^{\frac{1-\alpha\rho}{\alpha\rho}} \tfrac{w_t^*}{P_{t-1}^*}\right) - \mu \equiv X \tag{7}$$

Combining equations (6) and (7) with the wage equation in (4) then yields an implicit definition of X which is independent of μ:

$$X = \tfrac{1-\alpha\rho}{\alpha\rho}\ln\left(\tfrac{U}{1-\alpha\rho}\right) + \tfrac{\sigma^2}{2(1-\alpha\rho)} + (1-\alpha\rho)\ln\Phi\left(\tfrac{\sigma}{1-\alpha\rho} - \tfrac{X}{\sigma}\right) \tag{8}$$

Turn now to inflation. Substitute the firm pricing equation (1) into the price aggregator (5) and simplify to get:

$$\pi_t^* = \tfrac{1-\rho}{\rho}\ln\left(\tfrac{1-\alpha\rho}{U}\right) + \tfrac{\alpha(1-\rho)(X+\mu)}{1-\alpha\rho} + \ln\left(\frac{\int_{\varepsilon_{it}\geq X} \exp\left(\min\left\{\frac{\widehat{\pi}_t^i}{1-\rho}, \frac{\pi_t^* + \ln(\gamma_t)}{1-\rho}\right\} - \frac{\rho(1-\alpha)\widehat{\pi}_t^i}{(1-\rho)(1-\alpha\rho)}\right) di}{\int_{\varepsilon_{it}\geq X} \exp\left(\min\left\{\frac{\widehat{\pi}_t^i}{1-\rho}, \frac{\pi_t^* + \ln(\gamma_t)}{1-\rho}\right\} - \frac{(1-\alpha)\widehat{\pi}_t^i}{(1-\rho)(1-\alpha\rho)}\right) di}\right) \tag{9}$$

Combining equations (8) and (9) then yields:

$$\begin{aligned}\pi_t^* \quad = \quad & \tfrac{\alpha(1-\rho)}{1-\alpha\rho}\left(\mu + \tfrac{\sigma^2}{2(1-\alpha\rho)}\right) + \alpha(1-\rho)\ln\Phi\left(\tfrac{\sigma}{1-\alpha\rho} - \tfrac{X}{\sigma}\right) \\ & + \ln\left(\frac{\int_{\varepsilon_{it}\geq X} \exp\left(\min\left\{\frac{\widehat{\pi}_t^i}{1-\rho}, \frac{\pi_t^* + \ln(\gamma_t)}{1-\rho}\right\} - \frac{\rho(1-\alpha)\widehat{\pi}_t^i}{(1-\rho)(1-\alpha\rho)}\right) di}{\int_{\varepsilon_{it}\geq X} \exp\left(\min\left\{\frac{\widehat{\pi}_t^i}{1-\rho}, \frac{\pi_t^* + \ln(\gamma_t)}{1-\rho}\right\} - \frac{(1-\alpha)\widehat{\pi}_t^i}{(1-\rho)(1-\alpha\rho)}\right) di}\right)\end{aligned} \tag{10}$$

Now use $\widehat{\pi}_t^i = \mu + \varepsilon_{it}$ and $\varepsilon_{it} \sim N\left(0,\sigma^2\right)$ with (6) to simplify (10). It will help to define:

$$\Upsilon(X,\sigma) \equiv \frac{\Phi\left(-\frac{(1-\alpha)\sigma}{(1-\rho)(1-\alpha\rho)} - \frac{X}{\sigma}\right)\left[\Phi\left(\frac{\sigma}{1-\alpha\rho} - \frac{X}{\sigma}\right)\exp\left(\frac{[2-\alpha(1+\rho)]\sigma^2}{2(1-\rho)(1-\alpha\rho)^2}\right)\right]^{1-\alpha}\left(\frac{U}{1-\alpha\rho}\right)^{\frac{1-\alpha\rho}{\alpha\rho}}}{\Phi\left(-\frac{\rho(1-\alpha)\sigma}{(1-\rho)(1-\alpha\rho)} - \frac{X}{\sigma}\right)} \tag{11}$$

If $\gamma_t \leq \Upsilon(X,\sigma)$, then:

$$\pi_t^* - \mu = \frac{(\alpha-\rho)(1-\alpha\rho)-(1-\alpha)^2}{(1-\rho)(1-\alpha\rho)^2}\frac{\sigma^2}{2} + \alpha(1-\rho)\ln\Phi\left(\frac{\sigma}{1-\alpha\rho}-\frac{X}{\sigma}\right) + \ln\left(\frac{\Phi\left(-\frac{\rho(1-\alpha)\sigma}{(1-\rho)(1-\alpha\rho)}-\frac{X}{\sigma}\right)}{\Phi\left(-\frac{(1-\alpha)\sigma}{(1-\rho)(1-\alpha\rho)}-\frac{X}{\sigma}\right)}\right) \tag{12}$$

Otherwise, $\pi_t^* - \mu$ solves:

$$\begin{aligned}\pi_t^* - \mu \;=\; & \frac{1-\alpha^2+\alpha(1-\rho)}{(1-\alpha\rho)^2}\frac{\sigma^2}{2} + \alpha(1-\rho)\ln\Phi\left(\frac{\sigma}{1-\alpha\rho}-\frac{X}{\sigma}\right) \\ & + \ln\left(\frac{\Phi\left(\frac{\sigma}{1-\alpha\rho}-\frac{X}{\sigma}\right)-\Phi\left(\frac{\sigma}{1-\alpha\rho}-\frac{\pi_t^*-\mu+\ln(\gamma_t)}{\sigma}\right)+\exp\left(\frac{\pi_t^*-\mu+\ln(\gamma_t)}{1-\rho}-\frac{(1+\alpha\rho-2\rho)\sigma^2}{2(1-\rho)^2(1-\alpha\rho)}\right)\Phi\left(-\frac{\rho(1-\alpha)\sigma}{(1-\rho)(1-\alpha\rho)}-\frac{\pi_t^*-\mu+\ln(\gamma_t)}{\sigma}\right)}{\Phi\left(\frac{\alpha\sigma}{1-\alpha\rho}-\frac{X}{\sigma}\right)-\Phi\left(\frac{\alpha\sigma}{1-\alpha\rho}-\frac{\pi_t^*-\mu+\ln(\gamma_t)}{\sigma}\right)+\exp\left(\frac{\pi_t^*-\mu+\ln(\gamma_t)}{1-\rho}+\frac{(1+\alpha\rho-2\alpha)\sigma^2}{2(1-\rho)^2(1-\alpha\rho)}\right)\Phi\left(-\frac{(1-\alpha)\sigma}{(1-\rho)(1-\alpha\rho)}-\frac{\pi_t^*-\mu+\ln(\gamma_t)}{\sigma}\right)}\right)\end{aligned} \tag{13}$$

Either way, we have a definition of $\pi_t^* - \mu$ which is independent of μ. □

Part 2 Impose $\alpha = 1$ on equations (1), (2), and (3) to get $p\left(w_t; \widehat{P}_t^i\right) = \frac{w_t}{\rho}$, $O_t(w_t) = [0,1]$, and $\ell\left(w_t; \widehat{P}_t^i\right) = \left(\frac{\rho \widehat{P}_t^i}{w_t}\right)^{\frac{1}{1-\rho}}$. Substituting $p\left(w_t^*; \widehat{P}_t^i\right)$ into equation (5) gives $P_t^* = \frac{w_t^*}{\rho}$ and substituting $\ell\left(w_t^*; \widehat{P}_t^i\right)$ into equation (4) 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 $\widehat{\pi}_t^i \sim N(\mu, \sigma^2)$, we can use the moment generating function of the normal distribution to simplify the preceding integral. The integral is taken over the entire set so the moment generating function just yields $\pi_t^* = \mu + \frac{\sigma^2}{2(1-\rho)}$. □

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 (8) from the proof of Part 1 above produces:

$$\frac{dX}{d\sigma} = \frac{\sigma}{1-\alpha\rho} + \frac{\frac{X}{\sigma}}{1 + \frac{\Phi\left(\frac{\sigma}{1-\alpha\rho}-\frac{X}{\sigma}\right)}{\phi\left(\frac{\sigma}{1-\alpha\rho}-\frac{X}{\sigma}\right)}\frac{\sigma}{1-\alpha\rho}} \tag{14}$$

The desired inequality is thus $\left(\frac{\sigma}{1-\alpha\rho}-\frac{X}{\sigma}\right)\frac{\Phi\left(\frac{\sigma}{1-\alpha\rho}-\frac{X}{\sigma}\right)}{\phi\left(\frac{\sigma}{1-\alpha\rho}-\frac{X}{\sigma}\right)} > -1$. Showing $x\Phi(x) > -\phi(x)$ completes the proof: $x\Phi(x) = x\int_{-\infty}^{x}\phi(t)\,dt > \int_{-\infty}^{x} t\phi(t)\,dt = -\int_{-\infty}^{x}\phi'(t)\,dt = -\phi(x)$. □

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

$$X = \frac{1-\rho}{\rho}\ln\left(\frac{U}{1-\rho}\right) + \frac{\sigma^2}{2(1-\rho)} + (1-\rho)\ln\Phi\left(\frac{\sigma}{1-\rho}-\frac{X}{\sigma}\right) \tag{15}$$

$$\pi_t^* - \mu = \frac{\sigma^2}{2(1-\rho)} + (1-\rho)\ln\Phi\left(\frac{\sigma}{1-\rho} - \frac{X}{\sigma}\right) \tag{16}$$

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 (15). 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 thus at most one $\sigma_0 > 0$ such that $f(\sigma_0) = 0$. To show exactly one such $\sigma_0 > 0$, it will suffice to show $\lim_{\sigma\to 0^+} f(\sigma) < 0$ and $\lim_{\sigma\to\infty} f(\sigma) > 0$. Equation (15) 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 (15) 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, $f'(\sigma_0) > 0$.

The exit channel behind this result depends on heterogeneity, not on $\widehat{\pi}_t^i$ being a point expectation. To see this, consider two types of firms $j \in \{1, 2\}$. The fraction of type 1 firms is τ and the fraction of type 2 firms is $1-\tau$. A type j firm takes inflation expectations over the entire distribution $N\left(\mu, \sigma_j^2\right)$, with the resulting CDF for its price expectations denoted by $F_j(\cdot)$. The price-setting problem with $\alpha = 1$ still yields $P_t^* = \frac{w_t^*}{\rho}$ but the labor demand of a type j firm is now $\left(\frac{\rho}{w_t^*}\right)^{\frac{1}{1-\rho}} \int \widehat{P}_t^{\frac{1}{1-\rho}} dF_j\left(\widehat{P}_t\right)$ and operation requires:

$$\frac{w_t^*}{\rho} \leq \left[\frac{1-\rho}{U}\int \widehat{P}_t^{\frac{\rho}{1-\rho}} dF_j\left(\widehat{P}_t\right)\right]^{\frac{1-\rho}{\rho}} \tag{17}$$

If only one type operates (say $j = 1$), then labor market clearing yields:

$$\frac{w_t^*}{\rho} = \left[\tau\int \widehat{P}_t^{\frac{1}{1-\rho}} dF_1\left(\widehat{P}_t\right)\right]^{1-\rho} \tag{18}$$

To ensure that only type 1 firms operate, we need (17) with $\frac{w_t^*}{\rho}$ as per (18) to hold at $j = 1$ but not at $j = 2$. Stated otherwise, we need:

$$\frac{1-\rho}{\tau^\rho \exp\left(\frac{\rho\left(\sigma_1^2 - \rho\sigma_2^2\right)}{2(1-\rho)^2}\right)} < U \leq \frac{1-\rho}{\tau^\rho \exp\left(\frac{\rho\sigma_1^2}{2(1-\rho)}\right)}$$

A necessary condition is $\frac{\rho\left(\sigma_1^2 - \rho\sigma_2^2\right)}{2(1-\rho)^2} > \frac{\rho\sigma_1^2}{2(1-\rho)}$ or, equivalently, $\sigma_1 > \sigma_2$. Combining $P_t^* = \frac{w_t^*}{\rho}$ with (18), we can now write:

$$\pi_t^* = \mu + \frac{\sigma_1^2}{2(1-\rho)} + (1-\rho)\ln\tau$$

If $\tau = 1$, then all firms use the same prior so everyone operates and we again have $\pi_t^* > \mu$. If $\tau \in (0,1)$, then only type 1 firms operate and $\pi_t^* = \mu$ provided $\sigma_1 = (1-\rho)\sqrt{2\ln\left(\frac{1}{\tau}\right)}$. $\square$

Part 5 Define $\overline{U} \equiv 1 - \alpha\rho$. It will suffice to establish the result for some subset of $\left(0, \overline{U}\right)$. At $\sigma = 0$, equations (12) and (13) both reduce to:

$$f(0) \equiv \lim_{\sigma\to 0^+} f(\sigma) = \alpha(1-\rho)\ln\Phi\left(\lim_{\sigma\to 0^+} \frac{-X(\sigma)}{\sigma}\right) \tag{19}$$

If $\lim\limits_{\sigma\to 0^+} \frac{-X(\sigma)}{\sigma} = \infty$, then $f(0) = 0$. Moreover, (12) and (13) will both also produce:

$$f'(0) \equiv \lim_{\sigma\to 0^+} f'(\sigma) = \alpha(1-\rho)\lim_{\sigma\to 0^+} \frac{\phi\left(-\frac{X(\sigma)}{\sigma}\right)}{\Phi\left(-\frac{X(\sigma)}{\sigma}\right)}\frac{1}{\sigma}\left(\frac{X(\sigma)}{\sigma} - X'(\sigma)\right) = \frac{\alpha(1-\rho)}{1-\alpha\rho}\lim_{\sigma\to 0^+} X'(\sigma) \tag{20}$$

where the last equality follows from using equation (14). Turn now to $X(\cdot)$. At $\sigma = 0$, equation (8) yields:

$$X(0) \equiv \lim_{\sigma\to 0^+} X(\sigma) = \overline{U}\left[\frac{1}{\alpha\rho}\ln\left(\frac{U}{\overline{U}}\right) + \ln\Phi\left(\lim_{\sigma\to 0^+}\frac{-X(\sigma)}{\sigma}\right)\right] \tag{21}$$

Notice from (21) that $X(0) > 0$ is impossible while $X(0) < 0$ is only possible if $U < \overline{U}$. Therefore, $U = \overline{U}$ implies $X(0) = 0$ and thus $\lim\limits_{\sigma\to 0^+}\frac{-X(\sigma)}{\sigma} = \infty$. Since $X'(0) \equiv \lim\limits_{\sigma\to 0^+} X'(\sigma) \cong \frac{X(h)-X(0)}{h-0} = \frac{X(h)}{h} \xrightarrow{h\to 0^+} -\infty$, it now follows that $f(0) = 0$ and $f'(0) < 0$ when $U = \overline{U}$. Taken together, $f(0) = 0$ and $f'(0) < 0$ imply existence of a $\sigma > 0$ such that $f(\sigma) < 0$. Combined with $\lim\limits_{\sigma\to\infty} f(\sigma) = \infty$, 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 $\epsilon > 0$ such that there also exists a $\sigma_B > 0$ satisfying $f(\sigma_B) = 0$ and $f'(\sigma_B) > 0$ when $U \in \left(\overline{U} - \epsilon, \overline{U}\right)$. To show existence of a $\sigma_A \in (0, \sigma_B)$ satisfying $f(\sigma_A) = 0$ and $f'(\sigma_A) < 0$, it will suffice to show $f(0) = 0$ and $f'(0) > 0$ when $U \in \left(\overline{U} - \epsilon, \overline{U}\right)$. If $f'(0) = 0$, then it will suffice to show $f(0) = 0$ and $f''(0) \equiv \lim\limits_{\sigma\to 0^+} f''(\sigma) > 0$. For any $U \in \left(0, \overline{U}\right)$, we have $X(0) \in (-\infty, 0)$ from (21) and thus $f(0) = 0$ from (19). We also have $X'(0) = \frac{1-\alpha\rho}{X(0)}\lim\limits_{\sigma\to 0^+}\left(\frac{X(\sigma)}{\sigma}\right)^2\phi\left(-\frac{X(\sigma)}{\sigma}\right) = 0$ from (14) and the property $\lim\limits_{z\to -\infty} z^2\phi(-z) = 0$. Therefore, equation (20) yields $f'(0) = 0$ and it remains to show $f''(0) > 0$. With $\sigma = 0$ and $\lim\limits_{\sigma\to 0^+}\frac{-X(\sigma)}{\sigma} = \infty$, equation (11) implies $\Upsilon(X(0), 0) = \left(\frac{U}{\overline{U}}\right)^{\frac{1-\alpha\rho}{\alpha\rho}} < 1$ so, for small taste shocks, the behavior of $f''(\cdot)$ around zero is dictated by equation (13). After some algebra (available upon request), we obtain $f''(0) =$

$\frac{(1-\alpha)(\alpha-\rho)+\alpha(1-\rho)^2}{(1-\alpha\rho)^2(1-\rho)}$ which is positive for $\alpha \in \left(1-\frac{\rho(1-\rho)}{2}-\sqrt{\left(1-\frac{\rho(1-\rho)}{2}\right)^2-\rho},1\right) \equiv (\underline{\alpha},1)$. Therefore, there must exist a $\sigma_A \in (0,\sigma_B)$ satisfying $f(\sigma_A)=0$ and $f'(\sigma_A)<0$ when $U \in \left(\overline{U}-\epsilon,\overline{U}\right)$ and $\alpha \in (\underline{\alpha},1)$. ■

Proof of Proposition 2

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:

$$\exp(y_t)=\xi_t\exp\left(\frac{\upsilon_A^2}{2}\right)\Phi\left(\upsilon_A-\frac{y_t+\lambda}{\upsilon_A}\right)+(1-\xi_t)\exp\left(\frac{\upsilon_B^2}{2}\right)\Phi\left(\upsilon_B-\frac{y_t+\lambda}{\upsilon_B}\right) \tag{22}$$

Under $\xi_t=0$, equation (22) 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 (22) 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 (22) implies $y_t<0$ which is a contradiction. Therefore, $y_t=0$ for $\xi_t\in(0,1)$. ■