What Kind of Transparency?*

Alex Frankel† Navin Kartik‡

September 16, 2015

Abstract

How transparent should a central bank be about (i) its objectives and (ii) its information on how monetary policy maps into economic outcomes? We consider a version of the Barro-Gordon framework in which monetary policy signals an inflation-biased Bank's private information on both dimensions. We find that a commitment to transparency about how policy affects outcomes is desirable while transparency about objectives need not be. Public uncertainty about a Bank's inflation or output target can mitigate the Bank's inability to commit to achieving that target.

*We thank Ryan Chahrour, Alessandro Lizzeri, Emi Nakamura, Jón Steinsson, Richard Van Weelden, Mike Woodford, and participants at the Columbia Macro lunch for helpful comments. Sergey Kolbin and Enrico Zanardo provided able research assistance.

†University of Chicago Booth School of Business; afrankel@chicagobooth.edu.

‡Columbia University, Department of Economics; nkartik@columbia.edu.

"A given [monetary] policy action…can have very different effects on the economy, depending (for example) on what the private sector infers…about the information that may have induced the policymaker to act, about the policymaker's objectives in taking the action…"

— Ben Bernanke (2003)

1. Introduction

Greater central bank transparency has come to be viewed as desirable over the past two decades. Among policy makers, the worldwide shift is illustrated by the two previous U.S. Fed Chairmen. Early in his tenure, Alan Greenspan told reporters in 1987, "Since I've become a central banker, I've learned to mumble with great incoherence. If I seem unduly clear to you, you must have misunderstood what I said." Yet, as recounted by Poole (2005), "The evolution to greater transparency proceeded step by step during the Greenspan years," and Greenspan's successor, Ben Bernanke has reflected that "one of [his] priorities was to make the Federal Reserve more transparent—and, in particular, to make monetary policy as transparent and open as reasonably possible." (Bernanke, 2013)

In this paper, following a substantial body of academic literature, we will view "more transparency" as greater disclosure of a central bank's *private* information—put differently, a reduction of asymmetric information between the central bank and the public. Our interest is in contrasting transparency about two kinds of private information: (i) policy objectives and (ii) information about how policy affects macroeconomic outcomes. We refer to these respectively as *preference transparency* and *transparency about policy effects*.[1] In a nutshell, our analysis suggests that transparency about policy effects is indeed desirable, but preference transparency need not be.

In Section 2, we develop a model of a CB that faces a credibility problem à la Kydland and Prescott (1977) and Barro and Gordon (1983a). The CB has private information about an economic variable (e.g., a cost push or money demand shock) that, together with monetary policy, determines inflation. The public forms its inflation expectation after observing monetary policy but not directly the shock. Output is determined by an expectational Phillips curve, and the CB seeks to raise output beyond the natural level; consequently, it is inflation biased.

The monetary instrument serves a dual role in this context: it allows the CB to stabilize the economy in response to shocks, but also acts as a *signal* to the public about the

[1] Geraats (2002) refers to closely related issues as political transparency and economic transparency. Our terminology recognizes that policy objectives may—even should—themselves be grounded in economic tradeoffs.

CB's information, thereby affecting inflation expectations (and hence output). Romer and Romer (2000), Nakamura and Steinsson (2013), and Melosi (2014) provide evidence of this signaling channel in the U.S.[2] As a benchmark, we establish the existence of a separating equilibrium in which the public perfectly predicts inflation—hence output remains at the natural level—but with excess inflation. This is a static signaling-game version of the familiar "time inconsistency" problem.

Our main contribution is to augment this model with private information for the CB about its inflation target.[3] (While we focus on the inflation target for concreteness, our points also hold when the private information about objectives concerns the output target; see Subsection 5.1 and Appendix C.) As in our epigraph quoting Ben Bernanke, the public is now faced with an "identification problem": is a monetary easing a response to a fundamental policy-effects shock with unchanged objectives (in which case inflation expectations would not change), or does it reflect a tolerance for higher inflation (which would alter inflation expectations)?

Section 3 establishes that reducing uncertainty about the CB's objectives causes the public to infer that policy changes are more likely due to private information about policy effects. Inflation expectations are therefore less responsive to the monetary instrument, which exacerbates the CB's temptation to produce surprise inflation. But then, owing to rational expectations, greater preference transparency simply ends up generating higher average excess inflation without affecting average output. On the flip side, transparency about policy effects makes the public's inflation expectation more responsive to the monetary instrument—changes are attributed more to the CB's preferences rather than its information about policy effects—which, in equilibrium, leads to less excess inflation.

The mechanism underlying our findings points to a downside of "well-anchored expectations", when that term refers to the public's inflation expectations not being very sensitive to what it observes.[4] In our framework, it is precisely greater sensitivity of inflation expectations to the CB's policies that is ex-ante desirable, because that increases the CB's cost of producing surprise inflation, ultimately leading to less excess inflation.

[2] Faust et al. (2004) argue that Romer and Romer's (2000) conclusions should be qualified, in part based on subsequent data.

[3] Gürkaynak et al. (2005) suggest that private-sector uncertainty and learning about the inflation target is important to reconcile certain empirical facts about the U.S. yield curve.

[4] In the words of Bernanke (2007): "I use the term 'anchored' to mean relatively insensitive to incoming data ... if the public is modeled as being confident in its current estimate of the long-run inflation rate, so that new information has relatively little effect on that estimate, then the essential idea of well-anchored expectations has been captured." Bernanke goes on to say that well-anchored expectations are desirable because they make actual inflation less responsive to economic fluctuations.

The European Central Bank states that it "aims at inflation rates of below, but close to, 2% over the medium term"; while some view this communication as insufficiently precise (e.g. Geraats, 2008), our results suggest a possible welfare justification for such opacity. We emphasize, however, that preference opacity is beneficial in our model only in a second-best sense: public uncertainty about the CB's objectives mitigates the CB's lack of commitment to achieving its preferred inflation level. If the CB could commit—or if that outcome could be obtained through other means[5]—then preference opacity would not be beneficial. But we doubt that such commitments can be obtained, at least fully, in a setting where the optimal inflation target is stochastic and is the CB's private information to begin with.

Our paper's argument is not that maximal preference opacity is always desirable; rather, we aim to elucidate a mechanism underlying the benefits of opacity. In general, these benefits should be weighed against benefits of transparency. Subsection 5.1 explores a tradeoff using a richer specification of the CB's preferences: a quadratic loss around an output target rather than a linear benefit of raising output. Here, the benefit that preference opacity provides from reducing excess inflation can be outweighed by greater output and inflation instability. We show that an interior level of preference transparency can be optimal. Preference opacity also makes it more difficult for the public to form accurate inflation expectations: inflation is less predictable. Predictable inflation may be desirable *per se*. Our results imply that generating predictability through CB preference transparency entails a tradeoff with excess inflation. Leveraging the quadratic-loss output-target extension, we discuss in Subsection 5.2 how preference transparency may or may not be desirable when balancing these two factors.

We also study transparency about the monetary policy itself (Subsection 4.2). Consistent with Faust and Svensson (2001) and Sibert (2009), we find that providing the public with less information about the CB's policy exacerbates the CB's credibility problem, because policy opacity reduces the responsiveness of inflation expectations to monetary policy. In equilibrium, though, the public anticipates this effect; policy opacity leads, on average, to more excess inflation without more output, which lowers welfare. By contrast, the influential work of Cukierman and Meltzer (1986) can be viewed as arguing that opacity about monetary policy can be desirable.[6] The difference arises because the ability to produce surprise inflation is not always self-defeating in Cukierman and Meltzer (1986). In their

[5] E.g., by suitable contracts with the central banker (Walsh, 1995) or by appointing suitably biased central bankers (Rogoff, 1985; Svensson, 1997).

[6] Cukierman and Meltzer (1986) model the CB as choosing "planned monetary growth" m^P while actual monetary growth, which is what the public observes, is m^P plus noise. As they assume that the CB's payoff depends only on m^P, it is natural to interpret the level of noise as the (inverse of) the degree of transparency about monetary policy.

specification, discretion can be valuable in tailoring surprise inflation to times in which it is most beneficial. Transparency about monetary policy makes it more costly for the CB to exploit this channel.

We are not the first to argue that CB preference transparency can be welfare reducing. To some extent, the idea is present in the literature on reputation building by CBs, starting with Barro and Gordon (1983b) and Backus and Driffill (1985). In these models preference uncertainty combined with dynamic considerations create incentives for CBs to pretend to be more inflation averse than they actually are, which leads to lower excess inflation. Our mechanism is somewhat different: while those papers emphasize the dynamics of the public's updating over time, our (one-period) setting can be viewed as a stationary or steady-state situation.

More similar to our main themes are points made by Faust and Svensson (2001), Geraats (2007), Mertens (2011), and Tang (2013). Both Faust and Svensson (2001) and Mertens (2011) illustrate numerically that direct observability of a CB's objectives (modeled as output targets) can lead to higher average inflation and reduce social welfare. Our closed-form analysis clarifies the mechanism underlying their findings, as elaborated at the end of Section 3. Tang (2013) derives related analytical results; like us, she argues that opacity about objectives can sometimes be beneficial.[7]

The unpublished work of Geraats (2007), which we only became aware of after circulating a previous draft of the current paper, has much in common with our own, both in modeling and analysis. Like us, she shows in her baseline model that (full) "economic transparency"—transparency about policy effects, in our terminology—is beneficial while (full) preference transparency is harmful.[8] Geraats (2007) only considers the extremes of full and no transparency in each case; we also consider intermediate levels.Furthermore, we study welfare consequences beyond just the level of excess inflation; this emphasis owes to our view that the inflation target reflects social tradeoffs rather than idiosyncratic goals of the central bank. The welfare distinction and intermediate levels of transparency are especially important in Subsection 5.1; among other things, we show there that an interior level of preference transparency can be optimal. We also study various interventions

[7] Unlike us, Tang (2013) only considers full and no transparency about each dimension of the CB's private information (objectives and "demand shocks"). Her models are not entirely comparable to ours and her results are couched somewhat differently. Besides her theoretical analysis, she undertakes an empirical exercise to demonstrate the signaling effect of interest rates.

[8] Ellingsen and Söderström (2001) is another theoretical study in which the CB has private information about either its preferences or about economic shocks (but not both, in their analysis). They are concerned with a very different issue than we are: they seek to explain why monetary policy can have contrasting effects on the yield curve depending on the nature of the CB's private information.

and extensions (Section 4 and Section 5) that Geraats (2007) does not; conversely, she covers certain issues that we don't.

Besides the mechanism elucidated in this paper, there are other reasons why (full) information disclosure by CBs may not be desirable, e.g. Morris and Shin (2002) and the literature it has spawned. Baeriswyl and Cornand (2010), Walsh (2007), and Tamura (2014) study the interaction of central bank transparency and monetary policy signaling while emphasizing different issues than we do; dispersed information among firms is central in their models.[9] Geraats (2002) and Blinder et al. (2008) provide useful surveys of the theoretical and empirical literature on central bank communication and transparency.

2. A Signaling Model of Monetary Policy

We consider a version of the Barro and Gordon (1983b) monetary policy game; following Canzoneri (1985), we incorporate private information for the Central Bank about the state of the economy. Formally, we study a one-shot signaling game between two agents: a Central Bank, CB, and the Public (or private sector), depicted in Figure 1 below.

The CB first observes the realization of a shock $\eta \in \mathbb{R}$ and an inflation target $\pi^* \in \mathbb{R}$. These random variables are drawn from a given joint distribution. The CB then chooses the value of a monetary instrument, $m \in \mathbb{R}$. The public observes m—but nothing directly about η or π^*, unless specified otherwise—and forms its inflation expectation $\hat{\pi} \equiv \mathbb{E}[\pi|m]$. Inflation π and output y are then determined according to

$$\pi = m - \eta, \tag{1}$$

$$y = \pi - \hat{\pi}, \tag{2}$$

where Equation 2 is based on an expectational Phillips curve with a natural rate of output normalized to 0. The CB's objective is to maximize the expected value of

$$\gamma y - \frac{(\pi - \pi^*)^2}{2}. \tag{3}$$

Equation 1 reflects η being a *policy-effects* shock; η determines how monetary policy translates into realized inflation. One can view η as either a real or monetary shock. The inflation target π^* is a preference parameter that determines the desired level of inflation,

[9] The signaling role of policy is also the focus of Angeletos et al. (2006), but in a distinct context. We should further note that others have argued that only certain kinds of transparency are desirable, for example Prat (2005); these mechanisms are different from our focus and they concern settings outside of central banking.

with quadratic losses for inflation above or below the target. We refer to $\pi - \pi^*$ as *excess inflation*. Finally, $\gamma > 0$ parameterizes the tradeoff between boosting output and producing inflation different from the target. The objective in Barro and Gordon (1983b) coincides with that in (3) when $\pi^* = 0$. While (3) represents a linear value of output, our main themes extend to a specification with a quadratic loss from an output target (see Subsection 5.1), which incorporates a stabilization role for monetary policy.

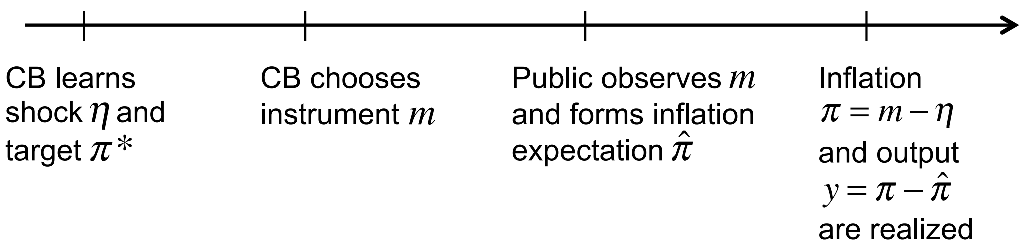

Figure 1 – Timeline.

Even though the CB is effectively choosing inflation π through its choice of m (since $\pi = m - \eta$ and the CB knows η), the public does not directly observe π. Instead, the public tries to infer π (or, equivalently, η) from the CB's choice of m. Consequently, the CB's monetary instrument has both a direct effect on its payoff—by determining inflation—and an indirect effect in how it affects the public's inflation expectation. The latter is the signaling effect of monetary policy.

All aspects of the model except the realization of the CB's private information, η and π^*, are common knowledge. We will study pure strategy Perfect Bayesian equilibria, which for our purposes can be described entirely by the CB's strategy $m(\eta, \pi^*)$. For any on-the-equilibrium path m, the public's inflation expectation, $\hat{\pi}(m) \equiv \mathbb{E}[\pi|m]$, is determined by Bayes rule.[10] Given any $\hat{\pi}(m)$, we can substitute Equation 1 and Equation 2 into (3) and rewrite the CB's objective as choosing m to maximize

$$\underbrace{\gamma(m - \eta - \hat{\pi}(m))}_{\text{output benefit}} - \underbrace{\frac{(m - \eta - \pi^*)^2}{2}}_{\text{excess inflation cost}}. \tag{4}$$

We define *welfare* as the CB's ex-ante expected utility. It bears emphasis that interpreting this object as social welfare presumes that changes in π^* reflect socially optimal tradeoffs between output and inflation. This perspective takes both η and π^* as economic condi-

[10] Off path beliefs will not play any material role in our analysis.

tions.[11] Regardless, the key distinction between η and π^* is that given any policy choice m, inflation π is only affected by the variable η—and hence, inflation expectation $\hat{\pi}$ only depends on beliefs about η. Naturally, the choice of m will depend on both η and π^*; indeed, it is evident from (4) that $\eta + \pi^*$ is a sufficient statistic for the CB's choice of m.

2.1. Benchmarks

Our first result concerns the properties of equilibria when the public knows η or π^*.

Proposition 1. *The following benchmarks hold:*

1. *Assume η is known to the public. Then there is a unique equilibrium; the CB plays $m(\eta, \pi^*) = \eta + \pi^*$ and the public forms the expectation $\hat{\pi}(m) = m - \eta$. For any η and π^*, output is 0 and there is no excess inflation; welfare is 0.*

2. *Assume π^* is known to the public but η is not. Then there is a separating equilibrium in which the CB plays $m(\eta, \pi^*) = \eta + \pi^* + \gamma$ and the public forms the expectation $\hat{\pi}(m) = \pi^* + \gamma$. For any η and π^*, output is 0 and there is excess inflation of γ; welfare is $-\gamma^2/2$.*

(All proofs are in Appendix A.)

Proposition 1 is an analog of the time-inconsistency or credibility problem that has received much attention by monetary economists since the work of Kydland and Prescott (1977) and Barro and Gordon (1983b). As rational expectations implies that $\mathbb{E}[y] = 0$ under any CB policy (under any assumptions about the public's information about η and π^*), the CB cannot do better than committing to $m(\eta, \pi^*) = \eta + \pi^*$. This policy would produce no output benefit but also no cost from excess inflation. Part 1 of Proposition 1 says that if the CB could not fool the public about inflation (because the public knows η and observes m), it would induce $\pi = \pi^*$ and replicate the commitment outcome. Part 2 shows that when the CB has private information on how policy maps into inflation, lack of commitment reduces welfare. There is an equilibrium in which the CB still achieves no output benefit, but now bears an inflationary cost. The intuition is familiar: in the absence of commitment, if the public thought the CB were playing the commitment strategy, the CB could always profitably deviate by raising m slightly above $\eta + \pi^*$; such a deviation would produce only a second order cost from excess inflation, but yield a first order output benefit through the surprise inflation. Due to this credibility problem, equilibrium excess inflation is constant

[11] The equilibrium outcomes are the same, but with different welfare interpretations, if π^* is taken to be a preference parameter of the CB that is not tied to social welfare.

at γ, so that inflation is always above the CB's preferred level.[12]

We pause here to comment on our model's timing assumption. If we had instead assumed that the public's inflation expectation is formed prior to observing m—as in Barro and Gordon (1983a,b)—then there would be excess inflation even when the public knows η, contrary to part 1 of Proposition 1. We make our assumption because, as already discussed in the introduction, there is evidence that monetary policy has a signaling role and affects public expectations, and we wish to focus on this aspect of the problem. It is not important for our message that *all* firms/consumers be able to adjust their expectations (and/or their behavior) in response to monetary policy; only that *some* do.

2.2. Unknown Objectives

Now let the CB have private information about both η and its inflation target, π^*. To make our main points most transparently, we maintain the following distributional assumption hereafter, using the notation $\mathcal{N}(\mu, \sigma^2)$ to denote a normal distribution with mean μ and variance σ^2:

Assumption 1. *The variables η and π^* are independent, with $\eta \sim \mathcal{N}(\mu_\eta, \sigma_\eta^2)$ and $\pi^* \sim \mathcal{N}(\mu_{\pi^*}, \sigma_{\pi^*}^2)$.*

Normality implies that the realized inflation target, π^*, may be negative. The parameters μ_{π^*} and $\sigma_{\pi^*}^2$ can be chosen to make this event have arbitrarily small probability.

We consider *linear equilibria* in which the public's expectations $\hat{\eta} \equiv \mathbb{E}[\eta|m]$ and $\hat{\pi} \equiv \mathbb{E}[\pi|m]$ are given by

$$\hat{\eta}(m) = Lm + K, \tag{5}$$

$$\hat{\pi}(m) = m - \hat{\eta}(m) = (1 - L)m - K, \tag{6}$$

for some constants L and K.

It follows from (4) and (6) that given some L and K, the CB's optimal response for any realized (η, π^*) is

$$m(\eta, \pi^*) = \eta + \pi^* + \gamma L. \tag{7}$$

Using standard results about normal information (DeGroot, 1970), we obtain:

[12] As is familiar in signaling games, the separating equilibrium leads to a discontinuity in actions and payoffs as we go from no uncertainty about η to any uncertainty about η. In Subsection 2.2 we show that this discontinuity is an artifact that only arises when there is no uncertainty about π^*; adding some uncertainty about π^* recovers continuity of the equilibria we study with respect to beliefs about η.

Lemma 1. *Suppose the CB uses the strategy in Equation 7. Then, conditional on m, the public's posterior belief on η is normally distributed with mean*

$$\mu_{\eta|m} = \frac{\sigma_\eta^2}{\sigma_\eta^2 + \sigma_{\pi^*}^2} m + \mu_\eta - \frac{\sigma_\eta^2}{\sigma_\eta^2 + \sigma_{\pi^*}^2} \left(\gamma L + \mu_\eta + \mu_{\pi^*} \right). \tag{8}$$

Noting that $\hat{\eta}(m)$ is just $\mu_{\eta|m}$, we can match coefficients in Equation 5 and Equation 8 to determine the equilibrium L and K:

$$L = \frac{\sigma_\eta^2}{\sigma_\eta^2 + \sigma_{\pi^*}^2}, \tag{9}$$

$$K = \frac{\mu_\eta \sigma_{\pi^*}^2 - \sigma_\eta^2 \mu_{\pi^*}}{\sigma_\eta^2 + \sigma_{\pi^*}^2} - \gamma \left(\frac{\sigma_\eta^2}{\sigma_\eta^2 + \sigma_{\pi^*}^2} \right)^2. \tag{10}$$

Recall that in the equilibrium of Proposition 1 part 2, with uncertainty about η but not π^*, the public's inflation expectation $\hat{\pi}$ was independent of the monetary instrument m. This led to a commitment problem in which the CB was tempted to increase m to increase output until the (constant) marginal benefit of extra output was equalized with the (increasing) marginal cost of excess inflation. Excess inflation was $\pi - \pi^* = \gamma$, and society bore the cost of this excess inflation without any output benefit.

Adding uncertainty about π^* as well, new dynamics take hold. The public's inflation expectation now increases in m: $\hat{\pi}(m) = (1 - L)m - K$, and therefore, using Equation 9, $\hat{\pi}'(m) = 1 - L = \frac{\sigma_{\pi^*}^2}{\sigma_\eta^2 + \sigma_{\pi^*}^2}$. Increasing m thus becomes less attractive to the CB: there is a smaller marginal output benefit for the same increase in inflation. So while the CB still cannot increase average output, the commitment problem becomes less severe and inflation goes down. Excess inflation is constant at

$$\pi - \pi^* = m - \eta - \pi^* = \frac{\sigma_\eta^2}{\sigma_\eta^2 + \sigma_{\pi^*}^2} \gamma,$$

which is lower than γ, the excess inflation when π^* is known. This decrease in excess inflation gives a corresponding increase in welfare from $-\gamma^2/2$ to $-\left(\frac{\sigma_\eta^2}{\sigma_\eta^2 + \sigma_{\pi^*}^2}\right)^2 \gamma^2/2$.

Proposition 2. *There is a unique linear equilibrium; it is defined by the coefficients (9) and (10), with the CB's strategy given by (7) and public expectations given by (5) and (6). For any η and π^*, excess inflation is $\frac{\sigma_\eta^2}{\sigma_\eta^2 + \sigma_{\pi^*}^2}\gamma$; average output is 0; and welfare is $-\left(\frac{\sigma_\eta^2}{\sigma_\eta^2 + \sigma_{\pi^*}^2}\right)^2 \gamma^2/2$.*

We see that welfare is decreasing in σ_η^2 and increasing in $\sigma_{\pi^*}^2$. A higher variance in the policy-effects shock, η, is harmful, but a higher variance in the CB's objectives, π^*, actually improves welfare. These effects are driven by how monetary policy m affects the inflation expectation $\hat{\pi}$ in equilibrium. More ex-ante public uncertainty about π^*, or less uncertainty about η, results in the public attributing changes in m more to variation in π^* rather than to η. Consequently, changes in m have amplified impact on inflation expectation $\hat{\pi}$, which means that the effect of m on output is dampened. This effect reduces the marginal benefit for the CB of raising m beyond the commitment level, mitigating the CB's temptation to create surprise inflation.

3. The Consequences of Transparency about Policy Effects and Preferences

This section presents our main results on how different forms of transparency affect welfare. In all cases, we take the viewpoint of the CB committing ex-ante to disclosure about the respective variable. Such commitments are more easily obtained, we believe, than commitments to monetary policy functions. For example, more or less disclosure can be induced by regulating the frequency of official central bank communication.

Subsection 2.2 showed that decreasing the variance of η reduces inflation and improves welfare, while decreasing the variance of π^* has the reverse effect of increasing inflation and worsening welfare. While these comparative statics are on the underlying distributions of these variables, they suggest corresponding comparative statics on transparency. Giving out information about η or π^* reduces the residual variance of the public's beliefs on these variables, and intuitively should have the same inflation and welfare implications.

Indeed, from Proposition 1 and Proposition 2, we can already look at the effect of revealing full information about η, π^*, or both. Fully revealing η—whether or not π^* is revealed—gets us to the equilibrium in Proposition 1 part 1 in which the commitment solution is recovered. This reduces inflation and improves welfare relative to the environment of Proposition 2 where both are unknown. On the other hand, revealing π^* but not η gets us to the equilibrium of Proposition 1 part 2, which has higher inflation and lower welfare than where both are unknown. In other words: *transparency on policy effects is beneficial, while transparency on preferences is harmful*.

We can extend these observations to more general information policies, not just the limiting cases of revealing everything or nothing. Formally, take the underlying variances of π^* and η in the economy as fixed, but suppose the CB may reveal information to the public that reduces the residual uncertainty in the following way. Prior to observing η and π^*, the CB publicly commits to revealing public signals $\eta' = \eta + \varepsilon_\eta$ and $\pi'^* = \pi^* + \varepsilon_{\pi^*}$, for

$\varepsilon_\eta \sim \mathcal{N}(0, \sigma^2_{\varepsilon_\eta})$ and $\varepsilon_{\pi^*} \sim \mathcal{N}(0, \sigma^2_{\varepsilon_{\pi^*}})$, independent of each other as well as the fundamentals. The variances $\sigma^2_{\varepsilon_\eta}$ and $\sigma^2_{\varepsilon_{\pi^*}}$ are each chosen from some subset of feasible alternatives in $\mathbb{R}_+ \cup \{\infty\}$. On each dimension, higher variance corresponds to less transparency: a variance of 0 provides the public with perfect information, while a variance of ∞ provides the public with no information. The signals are revealed after the CB observes η and π^*, but before the CB chooses m.[13]

Standard results imply that the posterior beliefs of the public after observing the signal realizations become

$$\eta \sim \mathcal{N}\left(\frac{\mu_\eta \sigma^2_{\varepsilon_\eta} + \eta' \sigma^2_\eta}{\sigma^2_\eta + \sigma^2_{\varepsilon_\eta}}, \frac{\sigma^2_\eta \sigma^2_{\varepsilon_\eta}}{\sigma^2_\eta + \sigma^2_{\varepsilon_\eta}}\right), \tag{11}$$

$$\pi^* \sim \mathcal{N}\left(\frac{\mu_{\pi^*} \sigma^2_{\varepsilon_{\pi^*}} + \pi'^* \sigma^2_{\pi^*}}{\sigma^2_{\pi^*} + \sigma^2_{\varepsilon_{\pi^*}}}, \frac{\sigma^2_{\pi^*} \sigma^2_{\varepsilon_{\pi^*}}}{\sigma^2_{\pi^*} + \sigma^2_{\varepsilon_{\pi^*}}}\right). \tag{12}$$

In particular, in the continuation game, the public's updated means of η and π^* depend on the signal realizations, but the variances do not. So the welfare and the excess inflation in the continuation game will be as in Proposition 2, with σ^2_η and $\sigma^2_{\pi^*}$ replaced by the respective variances from (11) and (12). Hence we can extend our earlier observation on the role of different types of transparency into a formal comparative statics result:

Proposition 3. *Welfare is increased by more transparency about η, i.e., lower $\sigma^2_{\varepsilon_\eta}$, and less transparency about π^*, i.e., higher $\sigma^2_{\varepsilon_{\pi^*}}$.*

It is appropriate to compare Proposition 3 with a finding of Faust and Svensson (2001); see also Mertens (2011). Faust and Svensson (2001) model a dynamic interaction between the CB and the public and consider the effects of transparency about the CB's preferences (its "goal," modeled as the CB's idiosyncratic employment target) and its monetary policy ("intention"). Using a numerical analysis, they deduce that welfare can be lower when goals are transparent. Our analytical result on the adverse effects of preference transparency clarifies the mechanism underlying their finding. In Faust and Svensson (2001), inflation is stochastically determined by monetary policy, but because the CB's goals are autocorrelated over time, today's monetary policy also affects future inflation expectations. Thus, when goals are unobservable, the CB faces a reputational cost of inflationary policy: more inflation today increases the public's beliefs about the CB's goal, which leads to higher future inflation expectation. This channel is precluded when goals are observable.

[13] It turns out not to matter whether the CB itself observes the realization of the signals η' and π'^* before choosing its policy m. The reason is that what matters to the CB is the public's residual variance about the CB's information, which, as seen in (11) and (12), is independent of the signals' realizations.

Similarly, in our one-period signaling model, a more transparent inflation target makes inflation expectations less responsive to monetary policy. In both Faust and Svensson (2001) and in our model, greater preference transparency lowers the cost of producing higher inflation, which leads to greater excess inflation in equilibrium and lower welfare. Tang (2013) also highlights this mechanism.

In Section 5 we revisit the analysis of transparency on policy effects and preferences for alternative specifications of the objective functions or of the inflation dynamics. Subsection 5.1 considers a setting with quadratic preferences around an output target, rather than a linear benefit of output. Subsection 5.2 considers welfare losses due to uncertainty in inflation, as distinct from those due to predictable excess inflation. Subsection 5.3 shows that our results are robust to allowing the public's inflation expectations to directly affect realized inflation, in the spirit of a new Keynesian framework.

4. Other Policy Interventions

This section discusses some interventions beyond transparency on policy effects and preferences.

4.1. Design of Objectives

Take the value of γ and the joint distribution of π^* and η in the economy as given. Return to the baseline assumptions that the CB observes the realizations of π^* and η, and it cannot credibly report any information about their realizations. However, now suppose a planner has some power in designing the CB's objectives. Let welfare be given by (3), as before: $\gamma y - (\pi - \pi^*)^2/2$. But consider a CB that optimizes the possibly distinct objective

$$\gamma_{CB} y - \frac{(\pi - \pi^*_{CB})^2}{2}. \tag{13}$$

Following Rogoff (1985), we can observe that endowing the CB with objectives that are misaligned with society's can be used to allay the CB's credibility problems and ultimately benefit society. In particular, Proposition 1 part 2 implies that the commitment solution can be obtained by taking $\gamma_{CB} \to 0$ while keeping $\pi^*_{CB} = \pi^*$. When the CB cares only about inflation and puts no weight on output, it will not have any incentive to juice output by increasing inflation above the ideal level. Similarly, keeping $\gamma_{CB} = \gamma$, we can recover the commitment solution by setting $\pi^*_{CB} = \pi^* - \gamma$ for every realization of π^*; this method corresponds to the inflation targeting of Svensson (1997).

Thus, if the planner has enough flexibility to appoint "biased" central bankers, or can use contractual incentives to appropriately (mis)align incentives, the underlying commitment problem goes away. There are many reasons however, including those outside our model, why such solutions are likely to be imperfect. Our analysis of transparency addresses the case of interest in which interventions to design the CB's objectives cannot be perfectly executed.

4.2. Transparency about Monetary Policy

Another dimension of transparency concerns monetary policy itself: we can study the effects of revealing or obscuring information about the CB's action, m, rather than its information. Suppose that instead of observing m directly, the public only observes a noisy signal about m. Specifically, when the CB chooses m, the public observes $m' = m + \varepsilon_m$, where $\varepsilon_m \sim \mathcal{N}(0, \sigma^2_{\varepsilon_m})$ independent of m. Given a linear conjecture by the public,

$$\hat{\pi}(m') = (1 - L)m' - K,$$

a CB of type (η, π^*) chooses m to maximize

$$\int_{-\infty}^{\infty} \gamma(m - \eta - ((1 - L)(m + \varepsilon_m) - K))\varphi(\varepsilon_m)\mathrm{d}\varepsilon_m - \frac{(m - \eta - \pi^*)^2}{2},$$

where $\varphi(\cdot)$ is the density of ε_m. Since ε_m is independent of m, the solution is precisely the same as in the noiseless case: $m(\eta, \pi^*) = \eta + \pi^* + \gamma L$.

The following is a generalization of Lemma 1.

Lemma 2. *Suppose the CB uses the strategy $m(\eta, \pi^*) = \eta + \pi^* + \gamma L$. Then, conditional on observing $m' = m + \varepsilon_m$ with $\varepsilon_m \sim \mathcal{N}(0, \sigma^2_{\varepsilon_m})$ independent of m,*

1. *the public's posterior belief on η is is normally distributed with mean*

$$\mu_{\eta|m'} = \frac{\sigma^2_\eta}{\sigma^2_\eta + \sigma^2_{\pi^*} + \sigma^2_{\varepsilon_m}} m' + \mu_\eta - \frac{\sigma^2_\eta}{\sigma^2_\eta + \sigma^2_{\pi^*} + \sigma^2_{\varepsilon_m}} (\mu_\eta + \mu_{\pi^*} + \gamma L); \tag{14}$$

2. *the public's posterior belief on m is normally distributed with mean*

$$\mu_{m|m'} = \frac{\sigma^2_\eta + \sigma^2_{\pi^*}}{\sigma^2_\eta + \sigma^2_{\pi^*} + \sigma^2_{\varepsilon_m}} m' + \frac{\sigma^2_{\varepsilon_m}}{\sigma^2_\eta + \sigma^2_{\pi^*} + \sigma^2_{\varepsilon_m}} \left(\mu_\eta + \mu_{\pi^*} + \gamma L\right); \tag{15}$$

3. *and hence the public's posterior belief on $\pi = m - \eta$ is normally distributed with mean*

$$\mu_{\pi|m'} = \frac{\sigma^2_{\pi^*}}{\sigma^2_\eta + \sigma^2_{\pi^*} + \sigma^2_{\varepsilon_m}} m' - \left(\mu_\eta - \frac{\sigma^2_{\varepsilon_m} + \sigma^2_\eta}{\sigma^2_\eta + \sigma^2_{\pi^*} + \sigma^2_{\varepsilon_m}} \left(\mu_\eta + \mu_{\pi^*} + \gamma L \right) \right). \tag{16}$$

Matching coefficients, there is a unique linear equilibrium, with

$$L = \frac{\sigma^2_\eta + \sigma^2_{\varepsilon_m}}{\sigma^2_\eta + \sigma^2_{\pi^*} + \sigma^2_{\varepsilon_m}},$$

$$K = \frac{\mu_\eta \sigma^2_{\pi^*} - \left(\sigma^2_\eta + \sigma^2_{\varepsilon_m}\right)\mu_{\pi^*}}{\sigma^2_\eta + \sigma^2_{\pi^*} + \sigma^2_{\varepsilon_m}} - \gamma \left(\frac{\sigma^2_\eta + \sigma^2_{\varepsilon_m}}{\sigma^2_\eta + \sigma^2_{\pi^*} + \sigma^2_{\varepsilon_m}} \right)^2.$$

Comparing these solutions with those in Equation 9 and Equation 10, we see that increasing $\sigma^2_{\varepsilon_m}$ is isomorphic to increasing σ^2_η; excess inflation, $\pi - \pi^*$, is now $\frac{\sigma^2_\eta + \sigma^2_{\varepsilon_m}}{\sigma^2_\eta + \sigma^2_{\pi^*} + \sigma^2_{\varepsilon_m}} \gamma$ and welfare is $-\left(\frac{\sigma^2_\eta + \sigma^2_{\varepsilon_m}}{\sigma^2_\eta + \sigma^2_{\pi^*} + \sigma^2_{\varepsilon_m}} \right)^2 \gamma^2/2$.

Corollary 1. *Welfare is increased by more transparency about m, i.e., lower $\sigma^2_{\varepsilon_m}$.*

The intuition is as follows. Although a noisier observation of m reduces the responsiveness of the public's expectation of η to its observation, it reduces the responsiveness of the expectation of m by even more—the derivative of the coefficient of m' with respect to $\sigma^2_{\varepsilon_m}$ is larger in magnitude in (15) than in (14). Consequently, as seen in (16), the expectation of $\pi = m - \eta$ becomes less responsive to the public's observation of monetary policy when the observation is noisier. This increases the CB's marginal benefit of producing inflation, which ultimately is self-defeating: it leads to greater excess inflation without any average output benefit.

Note that preference uncertainty is essential for transparency about monetary policy to matter: were $\sigma^2_{\pi^*} = 0$, the public's inflation expectation would be independent of its policy observation and welfare would be $-\gamma^2/2$ (the low welfare in the benchmark separating equilibrium of Proposition 1 part 2) independent of $\sigma^2_{\varepsilon_m}$. Moreover, as monetary policy becomes completely opaque, i.e., $\sigma^2_{\varepsilon_m} \to \infty$, welfare converges down to $-\gamma^2/2$; the benefit of preference uncertainty is completely undone.

The message in Corollary 1 is consistent with Faust and Svensson (2001). Their argument in favor of "observable intentions" can be viewed as an argument in favor of (full) transparency about m. Their mechanism is somewhat different, however. Due to their dynamic model and preference persistence, it would be as if the CB in our context gets an additional benefit when the public infers a low π^* from its observation about monetary policy. Intuitively, greater transparency about m would then better harness the CB's incentive to signal lower π^* by choosing lower m.

4.3. Controlling the CB's Information

So far we have assumed that the CB observes η and π^* perfectly. More generally, the CB may only get only noisy information about these variables.[14] First consider the extreme cases where the CB either gets no information about or perfectly learns each variable.

1. *No Info about η, No Info about π^*.* The public expects $\hat{\eta} = \mu_\eta$, and therefore $\hat{\pi} = m - \mu_\eta$. The CB will choose $m = \mu_\eta + \mu_{\pi^*}$, and inflation is $\mu_{\pi^*} + \mu_\eta - \eta$, which on average is μ_{π^*} with a variance of σ_η^2. Welfare is $-\frac{\sigma_{\pi^*}^2 + \sigma_\eta^2}{2}$.

2. *No Info about η, Full Info about π^*.* The public expects $\hat{\eta} = \mu_\eta$, and therefore $\hat{\pi} = m - \mu_\eta$. The CB will choose $m = \mu_\eta + \pi^*$, and inflation is $\pi^* + \mu_\eta - \eta$. Conditional on the CB's information, inflation is on average π^* but with variance of σ_η^2. Welfare is $-\frac{\sigma_\eta^2}{2}$.

3. *Full Info about η, No Info about π^*.* This is similar to the case of Proposition 1 part 2, with a separating equilibrium in which the CB chooses $m = \eta + \mu_{\pi^*} + \gamma$ and the public forms expectation $\hat{\pi}(m) = \mu_{\pi^*} + \gamma$. Inflation is $\mu_{\pi^*} + \gamma$, and welfare is $-\frac{\gamma^2 + \sigma_{\pi^*}^2}{2}$.

4. *Full Info about η, Full Info about π^*.* This is the case of Proposition 2, with inflation of $\pi^* + \frac{\sigma_\eta^2}{\sigma_\eta^2 + \sigma_{\pi^*}^2}\gamma$ and welfare of $-\left(\frac{\sigma_\eta^2}{\sigma_\eta^2 + \sigma_{\pi^*}^2}\right)^2 \gamma^2/2$.

We see that by remaining uninformed about η, the CB effectively commits not to generate surprise inflation, which is beneficial. If the CB were to try and generate any excess inflation with its monetary policy, this inflation would be fully anticipated by the public. It would have a cost but no output benefit, and so would be wasted. There is a tradeoff, however: remaining uninformed about η means that the CB is unable to set inflation exactly at the preferred level. So we may or may not want the CB to become informed about η, depending on the parameters. In particular, remaining uninformed about η increases welfare by $\gamma^2/2$ or $\left(\frac{\sigma_\eta^2}{\sigma_\eta^2 + \sigma_{\pi^*}^2}\right)^2 \gamma^2/2$ through the commitment channel, depending on whether the CB is informed about π^*, but gives a loss of σ_η^2 from the added noise in realized inflation. So the CB prefers to stay uninformed on η when the weight γ on output is sufficiently large (and thus the lack of commitment is very costly, because the incentive to surprise the public

[14] To interpret why the CB may only have limited information about π^*, recall that we view π^* as reflecting economic conditions that determine the socially optimal level of inflation.

is very large), or when σ_η^2 is sufficiently small (in which case not tailoring policy m to the shock η is not very costly).

There is no corresponding benefit from remaining uninformed on π^*. When the CB does not have private information about η, the CB faces no commitment problem and sets $m = \mu_\eta + \mathbb{E}[\pi^*]$; forgoing information about π^* simply adds variance to $\pi - \pi^*$. When the CB does have private information about η, forgoing information about π^* entails not only the loss from the variance on $\pi - \pi^*$, but moreover a loss from exacerbating the commitment problem as the public now infers more about equilibrium inflation from monetary policy.

Naturally, it would be difficult for a CB that is capable of getting information to commit to remaining uninformed; holding fixed the public's conjecture about its information precision, it is always better off with more information. An alternative interpretation is in terms of *competence*. Viewing a CB that is informed about η or π^* as a competent CB—as is appropriate when π^* is determined by social tradeoffs—the above discussion implies that competence on π^* is indirectly welfare enhancing because "competence provides credibility" (given some private information about η). On the other hand, competence about η may exacerbate the CB's credibility problem. In a related but distinct framework, Moscarini (2007) argues that a CB is more credible in its cheap-talk announcements when it is more competent.

In Appendix B we extend these themes to the CB getting noisy signals about η and π that are neither fully informative nor uninformative. In follow-up work to our own, Tamura (2015) considers even more general information acquisition policies.

5. Alternative Specifications

Our baseline specification of the CB's objective in (3) and of how inflation is determined in Equation 1 illustrate our main points as simply and clearly as possible. In this section we extend the analysis to some other specifications of interest, in order to both qualify our earlier results and also to show the robustness of the intuitions.

5.1. Quadratic Output Target

Suppose now that instead of the original objective (3), the CB no longer values output linearly. Instead, additional output due to surprise inflation is only valuable up to a point. Specifically, let the CB's payoff (and welfare) be given by

$$-\upsilon\frac{(y-\overline{y})^2}{2} - \frac{(\pi-\pi^*)^2}{2} + \frac{1}{2}\upsilon\overline{y}^2, \tag{17}$$

where $\overline{y} > 0$ is a commonly-known output target and $\upsilon > 0$ is a commonly-known parameter representing how the CB trades off output and inflation. The last term, $\upsilon\overline{y}^2/2$, is a constant that normalizes welfare to 0 when excess inflation is 0 and output is 0; this is the welfare that would be obtained were η commonly known. If $\upsilon = \gamma/\overline{y}$, the marginal benefit of output due to surprise inflation at $y = 0$ would be γ—the same as under (3). Hence, under $\upsilon = \gamma/\overline{y}$, the original linear objective (3) can be thought of as a first-order approximation of (17) about $y = 0$.[15]

When π^* is commonly known, there is a separating equilibrium which induces excess inflation of $\upsilon\overline{y}$. Under the specification $\upsilon = \gamma/\overline{y}$, excess inflation is γ and the equilibrium is identical to that of Proposition 1 part 2.

Lemma 3. *Let the CB's objective be (17) and suppose π^* is commonly known. Then, for any prior density of η, there is a separating equilibrium given by strategy $m(\eta) = \eta + \upsilon\overline{y} + \pi^*$ and beliefs $\hat{\eta}(m) = m - \upsilon\overline{y} - \pi^*$, $\hat{\pi}(m) = \pi^* + \upsilon\overline{y}$. For any η and π^*, output is 0 and there is excess inflation of $\upsilon\overline{y}$; welfare is $-(\upsilon\overline{y})^2/2$.*

Accordingly, we view higher υ and higher $\overline{y}$ as exacerbating the CB's credibility problem in this setting. For the rest of the section, we maintain the assumption that the CB privately learns both η and π^*, with the normal information structure as before. When the public holds a linear conjecture $\hat{\eta}(m) = Lm + K$, the CB's first-order condition from (17)—plugging in $y = m - \eta - \hat{\pi}$, $\hat{\pi} = m - \hat{\eta}$, and $\pi = m - \eta$—yields

$$m(\eta, \pi^*) = \eta\frac{1 + \upsilon L}{1 + \upsilon L^2} + \pi^*\frac{1}{1 + \upsilon L^2} + \frac{\upsilon L(\overline{y} - K)}{1 + \upsilon L^2}. \tag{18}$$

The second-order condition $-\upsilon L^2 - 1 < 0$ is satisfied.

Lemma 4. *Suppose the CB uses the strategy in Equation 18. Then, conditional on m, the public's posterior belief on η is normally distributed with mean*

$$\mu_{\eta|m} = \frac{\sigma_\eta^2(1 + \upsilon L)(1 + \upsilon L^2)}{\sigma_\eta^2(1 + \upsilon L)^2 + \sigma_{\pi^*}^2}m + \mu_\eta - \frac{\sigma_\eta^2(1 + \upsilon L)}{\sigma_\eta^2(1 + \upsilon L)^2 + \sigma_{\pi^*}^2}\left(\mu_\eta(1 + \upsilon L) + \mu_{\pi^*} + \upsilon L(\overline{y} - K)\right).$$

Matching coefficients from the conjectured beliefs $\hat{\eta}(m) = Lm + K$ to the corresponding formula in Lemma 4, in a linear equilibrium L must satisfy $L = \frac{\sigma_\eta^2(1+\upsilon L)\left(1+\upsilon L^2\right)}{\sigma_\eta^2(1+\upsilon L)^2+\sigma_{\pi^*}^2}$. There are

[15] The approximation becomes perfect as we take $\overline{y} \to \infty$ while maintaining $\upsilon = \gamma/\overline{y}$. (Perfect in the sense that for any given y, the marginal benefit of output tends to γ.)

two solutions, only one of which is positive:

$$L = \frac{\upsilon - 1 - \sigma_{\pi^*}^2/\sigma_\eta^2 + \sqrt{4\upsilon + \left(\upsilon - 1 - \sigma_{\pi^*}^2/\sigma_\eta^2\right)^2}}{2\upsilon} > 0. \tag{19}$$

As it is natural to focus on *increasing equilibria* ($\hat{\eta}(m)$ and $\hat{\pi}(m)$ are non-decreasing), we focus on the above solution. Under this equilibrium value for L, one can again match coefficients to solve for the constant K in the beliefs and strategy functions as

$$K = -\upsilon L^2 \overline{y} + \mu_\eta(1 - L) - \mu_{\pi^*} L. \tag{20}$$

Lemma 5. *L defined by Equation 19 is decreasing in $\sigma_{\pi^*}^2/\sigma_\eta^2$ with range $(0, 1)$.*

When $\sigma_{\pi^*}^2/\sigma_\eta^2 \to 0$, Equations 19 and 20 yield $L = 1$ and $K = -\upsilon\overline{y} - \mu_{\pi^*}$. This gives

$$\begin{aligned}
m(\eta, \pi^*) &= \eta + \frac{1}{1+\upsilon}(\pi^* + \upsilon\mu_{\pi^*}) + \upsilon\overline{y}, \\
\hat{\eta}(m) &= m - \upsilon\overline{y} - \mu_{\pi^*}, \\
\hat{\pi}(m) &= \mu_{\pi^*} + \upsilon\overline{y}.
\end{aligned}$$

For instance, in the limit when $\sigma_{\pi^*}^2 \to 0$, π^* becomes concentrated at μ_{π^*}, $m \to \eta + \mu_{\pi^*} + \upsilon\overline{y}$, and we reproduce the benchmark of Lemma 3. On the other hand, when $\sigma_{\pi^*}^2/\sigma_\eta^2 \to \infty$, we have $L = 0$ and $K = \mu_\eta$, yielding

$$\begin{aligned}
m(\eta, \pi^*) &= \eta + \pi^*, \\
\hat{\eta}(m) &= \mu_\eta, \\
\hat{\pi}(m) &= m - \mu_\eta.
\end{aligned}$$

Proposition 4. *Let the CB's objective be (17). Then, there is a unique increasing linear equilibrium; the public uses $\hat{\eta}(m) = Lm + K$ and the CB plays (18), with L and K given by Equation 19 and Equation 20. Welfare is*

$$-\frac{\upsilon}{2}\left[\frac{(1-L)^2\sigma_\eta^2 + L^2\sigma_{\pi^*}^2}{1+\upsilon L^2} + \upsilon L^2\overline{y}^2\right], \tag{21}$$

which can be decomposed as

$$\underbrace{-\frac{\upsilon}{2}\left(\frac{1-L}{\upsilon L^2+1}\right)^2\sigma_\eta^2-\frac{\upsilon}{2}\left(\frac{L}{\upsilon L^2+1}\right)^2\sigma_{\pi^*}^2}_{\text{loss from output variance}}\;\underbrace{-\frac{1}{2}L^2\upsilon^2\overline{y}^2}_{\substack{\text{loss from}\\ \text{avg. excess inflation}}}\;\underbrace{-\frac{1}{2}\left(\frac{L\upsilon}{1+\upsilon L^2}\right)^2\left((1-L)^2\sigma_\eta^2+L^2\sigma_{\pi^*}^2\right)}_{\text{loss from excess inflation variance}}. \tag{22}$$

The public's inflation expectation, $\hat{\pi}(m)=m-\eta(m)=(1-L)m-K$ is more sensitive to m when L is lower, which by Lemma 5 occurs when $\sigma_{\pi^*}^2/\sigma_\eta^2$ is higher. A more sensitive inflation expectation makes it more costly for the CB to create surprise inflation. Consequently, as seen in the "loss from avg. excess inflation" term in (22), the same mechanism as in Section 2 is at work in the current setting. However, besides average excess inflation, there are some new components of welfare. First, excess inflation is no longer constant across the CB's private information (see Equation 25 in Appendix A), which leads to a welfare loss from the associated variance. Second, because of the quadratic loss in deviations from the output target, there is a loss from output variance. The welfare consequence of average output being below the target (by rational expectations, average output is zero) simply cancels out with the constant $\frac{1}{2}\upsilon\overline{y}^2$ in our welfare specification.

The direct effect of increasing either $\sigma_{\pi^*}^2$ or σ_η^2 is to reduce welfare because these add variance to both output and excess inflation; note that in Section 2 there was no variance in excess inflation, while variance in output did not affect welfare. But there are also indirect effects: increasing $\sigma_{\pi^*}^2$ (σ_η^2) reduces (increases) L, by Lemma 5. Such a reduction (increase) in L increases (reduces) welfare, adding up the effects on average excess inflation as well as the variances of output and excess inflation.[16] Thus, increasing σ_η^2 unambiguously reduces welfare while increasing $\sigma_{\pi^*}^2$ has ambiguous effects. We can establish the following:

Proposition 5. *Let the CB's objective be* (17). *Then, in the equilibrium of Proposition 4:*

1. *More uncertainty about policy effects reduces welfare: welfare is decreasing in σ_η^2.*

2. *A little variation in objectives improves welfare if and only if the credibility problem is large relative to policy-effects uncertainty : welfare is higher with small $\sigma_{\pi^*}^2>0$ than with $\sigma_{\pi^*}^2=0$ if and only if $2\upsilon\overline{y}^2>\sigma_\eta^2$.*

3. *As $\sigma_{\pi^*}^2\to\infty$ welfare goes to $-\frac{\upsilon}{2}\sigma_\eta^2$. This welfare is lower than that which obtains at $\sigma_{\pi^*}^2=0$ if and only if there is sufficient policy-effects uncertainty relative to the credibility problem, i.e., if and only if $\sigma_\eta^2>\upsilon\overline{y}^2$.*

[16] See the proof of Proposition 5.

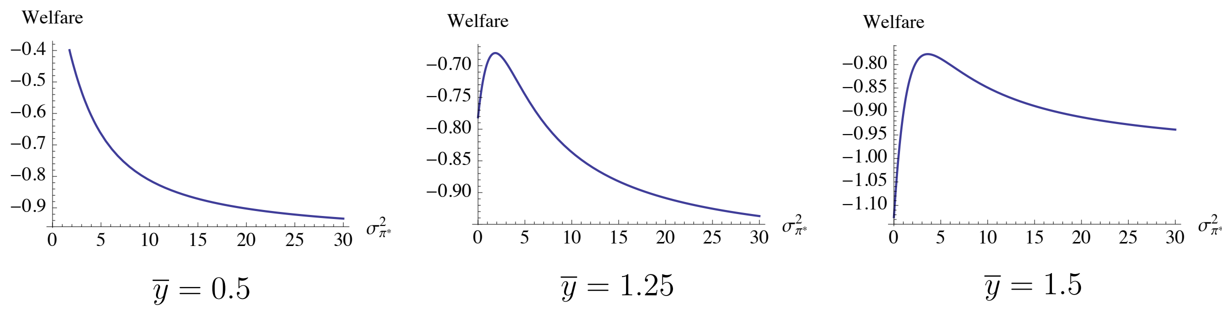

Figure 2 – Welfare as a function of $\sigma^2_{\pi^*}$, with a quadratic output target. $\sigma^2_\eta = 2$ and $\upsilon = 1$.

4. *Hence, for a range of parameters ($\upsilon\overline{y}^2 < \sigma^2_\eta < 2\upsilon\overline{y}^2$), welfare is non-monotonic in $\sigma^2_{\pi^*}$.*

Figure 2 illustrates parts 2–4 of Proposition 5. In the left-most panel, the CB's credibility problem is sufficiently weak that welfare is globally decreasing in $\sigma^2_{\pi^*}$. In the middle panel, the credibility problem is moderate, and so moderate values of $\sigma^2_{\pi^*}$ improve welfare but large $\sigma^2_{\pi^*}$ are worse than $\sigma^2_{\pi^*} = 0$. In the right-most panel, the credibility problem is so strong than any $\sigma^2_{\pi^*} > 0$ yields higher welfare than $\sigma^2_{\pi^*} = 0$. Notice that even in the right-most panel, welfare is eventually decreasing in $\sigma^2_{\pi^*}$; numerically, we find this to be a general conclusion.[17]

Transparency about policy effects and preferences. Consider the effects of transparency about η or π^*, using the same formalization as in Section 3 of providing the public normally distributed signals about these variables. As given by (11) and (12), observing the signal about each variable reduces the variance of the public's belief about that variable; crucially, the residual variances are independent of the signal realizations. Since the means of the beliefs about the variables (which are affected by the signal realizations) do not enter the welfare formula (21)—neither directly nor indirectly through the effect on L (Equation 19)—it follows from iterated expectations that welfare under any degree of transparency is simply given by formula (21) with the adjustment of replacing σ^2_η and $\sigma^2_{\pi^*}$ with the respective variances from (11) and (12), both directly and indirectly through the calculation of L.

Consequently, the effect of greater transparency is identical to that of reducing the primitives, $\sigma^2_{\pi^*}$ or σ^2_η. A number of implications now follow directly from Proposition 5:

Corollary 2. *Let the CB's objective be* (17). *Then,*

[17] Indeed, it appears numerically that welfare is always quasi-concave in $\sigma^2_{\pi^*}$ and decreasing for large enough $\sigma^2_{\pi^*}$; by part 2 of Proposition 5, this would imply that welfare is non-monotonic in $\sigma^2_{\pi^*}$ if and only if $2\overline{y}^2\upsilon > \sigma^2_\eta$.

1. *Greater transparency about policy effects increases welfare.*

2. *Welfare is decreasing (increasing) in transparency about preferences when $\sigma^2_{\pi^*}$ is small and the CB's credibility problem is large (small) relative to policy-effects uncertainty ($2\upsilon\overline{y}^2 > \sigma^2_\eta$ or $2\upsilon\overline{y}^2 < \sigma^2_\eta$ respectively).*

3. *Welfare is higher (lower) under full preference transparency than full preference opacity when $\sigma^2_{\pi^*}$ is large and the CB's credibility problem is mild (severe) relative to policy-effects uncertainty ($\sigma^2_\eta > \upsilon\overline{y}^2$ or $\sigma^2_\eta < \upsilon\overline{y}^2$ respectively).*

4. *For a range of parameters ($\upsilon\overline{y}^2 < \sigma^2_\eta < 2\upsilon\overline{y}^2$), when $\sigma^2_{\pi^*}$ is sufficiently large the optimal degree of transparency about preferences is neither full transparency nor full opacity.*

Part 1 of Corollary 2 extends the conclusion about policy-effects transparency from our baseline specification, as is intuitive given the discussion preceding Proposition 5. Part 2 extends our earlier conclusion about the harm of greater preference transparency, but there is a qualification now that the CB's credibility problem must be large and preference uncertainty be limited. The other two parts provide new insights: part 3 says that full preference transparency can be preferred to full opacity in some cases; while part 4 provides conditions assuring an interior level of optimal preference transparency. The logic behind these findings stems from points noted earlier: with a quadratic loss from an output target, greater uncertainty about the CB's preferences harms welfare through the variance of output and excess inflation; this by itself favors preference transparency. On the other hand, greater preference transparency increases average excess inflation, which is undesirable. This latter effect becomes dominant when the credibility problem is large and there isn't much uncertainty about the CB's preferences, which explains part 4 of Corollary 2.

Transparency about monetary policy. Suppose m is unobservable to the public. The equilibrium is now given by $\hat{\pi} = \upsilon\overline{y} + \mu_{\pi^*}$, $m(\eta, \pi^*) = \eta + \frac{\pi^* + \upsilon\mu_{\pi^*}}{1+\upsilon} + \upsilon\overline{y}$, and therefore $y = \frac{\pi^* - \mu_{\pi^*}}{1+\upsilon}$.[18] Welfare is $-\frac{\upsilon}{2}\frac{\sigma^2_{\pi^*}}{1+\upsilon} - \frac{1}{2}(\upsilon\overline{y})^2$. When $\sigma^2_{\pi^*} = 0$, this welfare is the same as that under observable m (Lemma 3), because a known π^* leads to a level of inflation that is invariant to η even when m is unobservable. On the other hand, because $-\frac{\upsilon}{2}\frac{\sigma^2_{\pi^*}}{1+\upsilon} - \frac{1}{2}(\upsilon\overline{y})^2 < -\frac{\upsilon}{2}\sigma^2_\eta$ if and only if $\sigma^2_{\pi^*} > (1+\upsilon)(\sigma^2_\eta - \upsilon\overline{y}^2)$, we see using Proposition 5 part 3 that preventing the public from observing m would lead to lower welfare whenever $\sigma^2_{\pi^*}$ is sufficiently high. While it is difficult to obtain a general analytical result, numerical simulations suggest that for any $\sigma^2_{\pi^*} > 0$, unobservable m leads to lower welfare than observable m, just as in the baseline specification. See Figure 3.

[18] To confirm this, note that given an arbitrary $\hat{\pi}$, the CB optimally chooses chooses $m(\eta, \pi^*) = \eta + \frac{\upsilon(\hat{\pi}+\overline{y})+\pi^*}{1+\upsilon}$. Hence, the equilibrium $\hat{\pi}$ must solve $\hat{\pi} = \frac{\upsilon(\hat{\pi}+\overline{y})+\mu_{\pi^*}}{1+\upsilon}$.

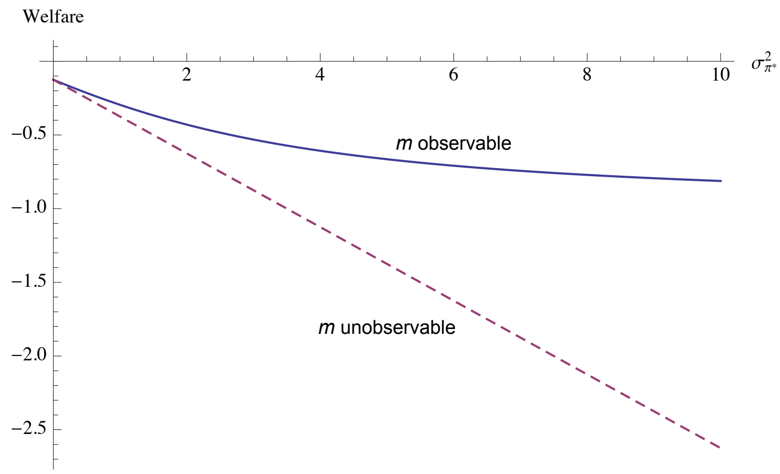

Figure 3 – Welfare under observable and unobservable monetary policy, as a function of $\sigma^2_{\pi^*}$ with a quadratic output target. $\sigma^2_\eta = 2$, $\upsilon = 1$, and $\overline{y} = 0.5$.

Uncertainty about the output target. We have also analyzed a model where the CB's objective is (17), but the two dimensions of private information are η and $\overline{y}$ rather than η and π^*. In other words, the public's preference uncertainty is about the output target rather than the inflation target. We find that full preference opacity always yields higher welfare than full preference transparency. The mechanism should be clear by this point: preference opacity makes the public's inflation expectation more sensitive to the CB's instrument, which mitigates the commitment problem. Appendix C contains a fuller discussion.

5.2. Penalizing Unexpected Inflation

Return to the baseline model with a linear value of output (objective (3)) and uncertainty about η and π^*. Suppose now the CB faces an additional loss due to unexpected inflation, so that it's objective is given by

$$\gamma y - \frac{(\pi - \pi^*)^2}{2} - \zeta\frac{(\pi - \hat{\pi})^2}{2}, \tag{23}$$

where $\zeta > 0$ is commonly known and parameterizes the payoff loss of unexpected inflation. In a reduced form way, this captures the misallocation that occurs when consumers and firms predict inflation incorrectly and therefore "get prices wrong". A given distribution of inflation is more costly when the public does not know its realization in advance.

As the expectational Phillips curve entails $y = \pi - \hat{\pi}$, unexpected inflation translates into excess output. So the objective function (23) simplifies to $\gamma y - \frac{\zeta}{2}y^2 - \frac{(\pi-\pi^*)^2}{2}$; plugging in $\overline{y} = \frac{\gamma}{\zeta}$ and $\upsilon = \zeta$, (23) reduces to $-\upsilon\frac{(y-\overline{y})^2}{2} - \frac{(\pi-\pi^*)^2}{2} + \frac{1}{2}\upsilon\overline{y}^2$. This function is exactly the objective with a quadratic loss from output deviations, (17). So the analysis here just follows the analysis in Subsection 5.1. Likewise, adding a welfare loss of unexpected inflation to (17) effectively just reduces the output target $\overline{y}$ while increasing the coefficient υ; it remains in the form of (17) with different parameters.[19]

5.3. Inflation Expectations Affect Inflation

Finally, we consider a variation in the spirit of a new Keynesian framework. Rather than inflation being determined solely by monetary policy and exogenous conditions, as in Equation 1, we now suppose that inflation expectations also have a direct influence on realized inflation. In particular, inflation is determined by a weighted average of policy-cum-conditions and of the public's beliefs about inflation:

$$\pi = (1-r)(m-\eta) + r\hat{\pi} \tag{24}$$

for some parameter $r \in [0, 1)$. The original specification, Equation 1, corresponds to $r = 0$, while $r \to 1$ means that expectations are perfectly self-fulfilling regardless of CB policy or underlying conditions. Note that even in the limit of $r \to 1$, monetary policy need not be non-influential despite becoming "cheap talk"; the public's belief $\hat{\pi}$ is an equilibrium object, which may be affected by m even if m has no direct effect on inflation.

Aside from this change to how inflation is determined, the model is as in Section 2. Output is determined by Equation 2 as $y = \pi - \hat{\pi}$. The CB maximizes the linear objective (3) given by $\gamma y - \frac{(\pi-\pi^*)^2}{2}$. In the interests of space, we defer a detailed analysis of this model to Appendix D, providing only a summary of the conclusions here.

First, Proposition 1 goes through unchanged for any $r \in [0, 1)$. When η is known to the public, the CB achieves the commitment outcome of $\pi = \pi^*$ and $y = 0$ by playing $m(\eta, \pi^*) = \eta + \pi^*$. The public predicts $\hat{\pi}(m) = m - \eta$ after observing m; this prediction is correct as realized inflation is $\pi = (1-r)(m-\eta) + r\hat{\pi}(m) = (1-r)(m-\eta) + r(m-\eta) = m - \eta$. On the other hand, when π^* is known to the public but η is not, there is a separating equilibrium in which the CB plays $m(\eta, \pi^*) = \eta + \pi^* + \gamma$, and inflation is $\pi^* + \gamma$, i.e., there is excess inflation. Here the public predicts inflation of $\hat{\pi} = \pi^* + \gamma$ for any m, and again the prediction is fulfilled as $\pi = (1-r)(m(\eta, \pi^*) - \eta) + r\hat{\pi}(m) = (1-r)(\eta + \pi^* + \gamma - \eta) + r(\pi^* + \gamma) = \pi^* + \gamma$.

In both the above cases the equilibrium is unchanged because the public perfectly predicts inflation given its information about the economy and its observation of m. When the public has uncertainty about both η and π^*, it can no longer predict inflation perfectly. Hence the new equilibrium will not be identical to that of the baseline model.

[19] When we add a term $\zeta(\pi-\hat{\pi})^2/2$ to (17), the new output target replacing $\overline{y}$ is $\frac{\upsilon}{\upsilon+\zeta}\overline{y}$, and the new coefficient replacing υ is $\upsilon + \zeta$.

We follow the approach of Subsection 2.2. Taking η and π^* to be independent normal, we seek a linear equilibrium of the form $\hat{\eta}(m) = Lm + K$, for some constants L and K.[20] The first-order condition for the CB's objective function yields a corresponding policy of

$$m(\eta, \pi^*) = \eta \frac{1-r}{1-Lr} + \pi^* \frac{1}{1-Lr} + \frac{Kr(1-Lr) + L(1-r)\gamma}{(1-Lr)^2}.$$

Proposition 6. *Let the CB's objective be* (3) *and inflation be determined by* (24). *Then, there is a unique linear equilibrium. For any η and π^*, excess inflation is* $\frac{(1-r)^2\sigma_\eta^2\gamma}{(1-r)^2\sigma_\eta^2+\sigma_{\pi^*}^2}$; *welfare is* $-\frac{1}{2}\left(\frac{(1-r)^2\sigma_\eta^2\gamma}{(1-r)^2\sigma_\eta^2+\sigma_{\pi^*}^2}\right)^2$.

Thus, excess inflation is constant, as was the case in our analysis in Section 2. For $r = 0$ we recover the result of Proposition 2. For other values of r the actual value of excess inflation is different—in particular, excess inflation is decreasing in r—but the comparative statics on the variances are unchanged. Excess inflation is increasing in σ_η^2 and decreasing in $\sigma_{\pi^*}^2$, and hence welfare is decreasing in σ_η^2 and increasing in $\sigma_{\pi^*}^2$. So we have the same qualitative conclusions as before, previously summarized in Proposition 3. Transparency on policy effects (η) is beneficial, while transparency on preferences (π^*) is harmful.

6. Conclusion

This paper has developed a simple signaling model of the central-bank credibility problem first pointed out by Kydland and Prescott (1977) and Barro and Gordon (1983a). A (benevolent) CB benefits from increasing output through surprise inflation; the public anticipates the CB's behavior, leading to excess inflation without any output benefit on average. The credibility problem arises because the CB has private information about *policy effects*, i.e., how the observable monetary policy translates into inflation. Our main contribution is to analyze the consequences of adding additional public uncertainty about the CB's *preferences*, i.e., its inflation target. Greater preference uncertainty makes the public's inflation expectation more responsive to monetary policy. This mitigates the credibility problem and leads to lower average inflation. Consequently, although transparency about policy effects is desirable, transparency about preferences can be detrimental.

The tractability of our model makes it amenable to a number of extensions besides those we have discussed. For example, we conjecture that it would fairly straightforward to (i) incorporate "control errors" by taking $\pi = m - \eta + \varepsilon$, where ε is noise that is realized only after m is chosen; (ii) modify the source of preference uncertainty to be about the CB's

[20] It continues to hold in equilibrium that $\hat{\pi} = m - \hat{\eta}$. To see this, take expectations on both sides of (24) to get $\hat{\pi} = (1-r)(m-\eta) + r\hat{\pi}$ and then rearrange to solve for $\hat{\pi}$.

rate of substitution between output and excess inflation (recall that uncertainty about the output target was discussed in Subsection 5.1); or (iii) augment the CB's preferences with reduced-form reputational concerns for being perceived as having a low inflation target. Tamura (2015) extends our model of a quadratic output target to allow for endogenous information acquisition by the CB, and shows that the optimal public signal should induce positive correlation between beliefs on π^*, y^*, and η even when the underlying realizations of these shocks are independent.

We close by mentioning that our central themes are relevant beyond monetary policy. In various signaling environments, agents have private information on multiple dimensions but only wish to affect market beliefs on a subset of them (Fischer and Verrecchia, 2000; Bénabou and Tirole, 2006; Frankel and Kartik, 2015). Mandating the disclosure of information that the market does not directly care about will tend to improve equilibrium information on the dimensions of interest. However, the improvement in information could be accompanied by exacerbated signaling, which can lead to a net reduction of welfare when signaling is dissipative.[21]

[21] Maggi (1999) shows that adding "purely private information" mitigates the effects of noisy observation of the first-mover's action in a sequential move setting (cf. Bagwell, 1995). Our inflation target, π^*, is purely private information in Maggi's sense, and we too find that it can benefit the first mover, which is the CB here. Maggi's (1999) model is not one of signaling, however. Indeed, he writes (p. 557) that his analysis combined with standard signaling models suggest that "in the presence of imperfect observability, the value of commitment is increased by purely private information (e.g., information about own cost), but tends to be decreased by follower-relevant private information (e.g., information about demand)." One can view our results as validating this hypothesis even without imperfect observability.

Appendices

A. Proofs

Proof of Proposition 1. Part 1: For any m, the public infers that $\hat{\pi} = m - \eta$. So realized output is $y = m - \eta - \hat{\pi} = 0$. The CB's objective (4) now reduces to choosing m to minimize the cost of excess inflation, $(m - (\eta + \pi^*))^2$. Minimization yields $m = \eta + \pi^*$ and therefore $\pi = m - \eta = \pi^*$.

Part 2: If the CB uses the strategy $m(\eta) = \eta + \pi^* + \gamma$, then realized inflation is $\pi = m - \eta = \pi^* + \gamma$. Hence, the public's expectation must be $\hat{\pi} = \pi^* + \gamma$ for all m. Under these beliefs, the CB's objective (4) becomes

$$\gamma(m - \eta - (\pi^* + \gamma)) - \frac{(m - (\eta + \pi^*))^2}{2}.$$

Maximizing over m gives the conjectured strategy $m = \eta + \pi^* + \gamma$ as optimal. *Q.E.D.*

Proof of Lemma 1. The prior on η is $\mathcal{N}(\mu_\eta, \sigma_\eta^2)$. Conditional on η (but not π^*), the distribution of $m = \eta + \pi^* + \gamma L$ is $\mathcal{N}(\eta + \mu_{\pi^*} + \gamma L, \sigma_{\pi^*}^2)$; equivalently, the observable value $m - \gamma L - \mu_{\pi^*}$ is normally distributed with mean η and variance $\sigma_{\pi^*}^2$. The lemma's conclusion now follows from standard results on updating normal priors with normal signals (DeGroot, 1970, p. 167). *Q.E.D.*

Proof of Proposition 2. Follows from the discussion in the text preceding the proposition. *Q.E.D.*

Proof of Proposition 3. Follows from the discussion in the text preceding the proposition. *Q.E.D.*

Proof of Lemma 2. The public observes $m' = m + \varepsilon_m = \eta + \pi^* + \gamma L + \varepsilon_m$.

Part 1: Conditional on η, the distribution of m' is $\mathcal{N}(\eta + \mu_{\pi^*} + \gamma L, \sigma_{\pi^*}^2 + \sigma_{\varepsilon_m}^2)$; from here, the proof follows that of Lemma 1.

Part 2: The prior distribution of m is $\mathcal{N}(\mu_\eta + \mu_{\pi^*} + \gamma L, \sigma_{\pi^*}^2 + \sigma_{\varepsilon_m}^2)$ and, conditional on m, the distribution of m' is $\mathcal{N}(m, \sigma_{\varepsilon_m}^2)$. The conclusion now follows from standard results on updating normal priors with normal signals.

Part 3: This expression is simply (15) minus (14). *Q.E.D.*

Proof of Lemma 3. As Bayes rule is obviously satisfied, we need only show that the CB's behavior is optimal given the public's conjecture. Since $\hat{\pi}(m) = \pi^* + \upsilon\overline{y}$ is independent of m, the CB's program given any η can be written as

$$\max_{\pi}\left[-\upsilon\frac{(\pi-\pi^*-\gamma\overline{y})^2}{2}-\frac{(\pi-\pi^*)^2}{2}\right].$$

This is a concave objective; the first order condition

$$-\upsilon(\pi-\pi^*-\upsilon\overline{y}-\overline{y})-(\pi-\pi^*)=0$$

solves for $\pi = \pi^* + \upsilon\overline{y}$, which corresponds to choosing $m = \eta + \pi^* + \upsilon\overline{y}$. *Q.E.D.*

Proof of Lemma 4. The prior on η is $N(\mu_\eta, \sigma_\eta^2)$. Conditional on η (but not π^*), the distribution of $m(\eta, \pi^*)$ from (18) is $\mathcal{N}\left(\eta\frac{1+\upsilon L}{1+\upsilon L^2}+\mu_{\pi^*}\frac{1}{1+\upsilon L^2}+\frac{\upsilon L(\overline{y}-K)}{1+\upsilon L^2}, \sigma_{\pi^*}^2\left(\frac{1}{1+\upsilon L^2}\right)^2\right)$; equivalently, the observable value $m\frac{1+\upsilon L^2}{1+\upsilon L}-\mu_{\pi^*}\frac{1}{1+\upsilon L}-\frac{\upsilon L(\overline{y}-K)}{1+\upsilon L}$ is distributed normally with mean η and variance $\sigma_{\pi^*}^2\frac{1}{(1+\upsilon L)^2}$. The lemma's conclusion follows from standard results on updating normal priors with normal signals. *Q.E.D.*

Proof of Lemma 5. From Equation 19, $L \to 1$ as $\sigma_{\pi^*}^2/\sigma_\eta^2 \to 0$ and $L \to 0$ as $\sigma_{\pi^*}^2/\sigma_\eta^2 \to \infty$. Moreover,

$$\frac{\partial L}{\partial(\sigma_{\pi^*}^2/\sigma_\eta^2)} = -\frac{L}{\sqrt{4\upsilon+(-\upsilon+\sigma_{\pi^*}^2/\sigma_\eta^2+1)^2}} < 0. \qquad \textit{Q.E.D.}$$

Proof of Proposition 4. We compute

$$\begin{aligned}\mathbb{E}[(y-\overline{y})^2] &= (\mathbb{E}[y]-\overline{y})^2+\mathrm{Var}(y)\\ &= \overline{y}^2+\mathrm{Var}(\hat{\eta}-\eta)\\ &= \overline{y}^2+\left(\frac{L-1}{\upsilon L^2+1}\right)^2\sigma_\eta^2+\left(\frac{L}{\upsilon L^2+1}\right)^2\sigma_{\pi^*}^2,\end{aligned}$$

where the second equality is because $\mathbb{E}[y]=0$ and $y=\hat{\eta}-\eta$, and the last equality is because for any realized η, π^*,

$$\begin{aligned}\hat{\eta}-\eta &= m(\eta,\pi^*)L+K-\eta\\ &= L\left[\eta\frac{1+\upsilon L}{1+\upsilon L^2}+\pi^*\frac{1}{1+\upsilon L^2}+\frac{L\upsilon(\overline{y}-K)}{1+\upsilon L^2}\right]+K-\eta\\ &= \eta\frac{L-1}{\upsilon L^2+1}+\pi^*\frac{L}{\upsilon L^2+1}+\frac{\upsilon L^2\overline{y}+K}{\upsilon L^2+1}.\end{aligned}$$

Moreover, for any realized η, π^*,

$$\begin{aligned}\pi - \pi^* &= m(\eta, \pi^*) - \eta - \pi^* \\ &= \eta\frac{1+\upsilon L}{1+\upsilon L^2} + \pi^*\frac{1}{1+\upsilon L^2} + \frac{L\upsilon(\overline{y}-K)}{1+\upsilon L^2} - \eta - \pi^* \\ &= \frac{L\upsilon}{1+L^2\upsilon}\left[(1-L)\eta - L\pi^* + \overline{y} - K\right],\end{aligned} \tag{25}$$

so we also compute

$$\begin{aligned}\mathbb{E}[(\pi-\pi^*)^2] &= \left(\mathbb{E}[\pi-\pi^*]\right)^2 + \mathrm{Var}(\pi-\pi^*) \\ &= \left(\frac{L\upsilon}{1+L^2\upsilon}\right)^2\left[\left((1-L)\mu_\eta - L\mu_{\pi^*} + \overline{y} - K\right)^2 + (1-L)^2\sigma_\eta^2 + L^2\sigma_{\pi^*}^2\right] \\ &= \left(\frac{L\upsilon}{1+L^2\upsilon}\right)^2\left(\overline{y}^2(1+\upsilon L^2)^2 + L^2\sigma_{\pi^*}^2 + (1-L)^2\sigma_\eta^2\right). \qquad \textit{Q.E.D.}\end{aligned}$$

Proof of Proposition 5. Part 1: That (21) is decreasing in L can be confirmed by partially differentiating it, substituting in (19), and performing some algebraic manipulations. Further, (21) is obviously decreasing in σ_η^2 (using $L \in [0,1]$ by Lemma 5). Finally, L is increasing in σ_η^2 by Lemma 5.

Part 2: A routine computation using (21) and (19) establishes that the total derivative of (22) with respect to $\sigma_{\pi^*}^2$ evaluated at $\sigma_{\pi^*}^2 = 0$ (at which point $L = 1$) is $\frac{\upsilon\left(2\upsilon\overline{y}^2 - \sigma_\eta^2\right)}{2(\upsilon+1)\sigma_\eta^2}$.

Part 3: Let $\sigma_{\pi^*}^2 \to \infty$. Since $L \to 0$ (Lemma 5), welfare goes to $-\frac{\upsilon}{2}\sigma_\eta^2$. The result follows by comparing this expression with the welfare expression obtained in Lemma 3.

Part 4: Implied by the two preceding parts of the result. *Q.E.D.*

Proof of Corollary 2. The result follows from Proposition 5 and the discussion in the text preceding the corollary. In more detail: consider a two-stage game in which the CB and the public first observe noisy signals $\eta' = \eta + \varepsilon_\eta$ and $\pi'^* = \pi^* + \varepsilon_{\pi^*}$, with $\varepsilon_\eta \sim \mathcal{N}(0, \sigma_{\varepsilon_\eta}^2)$ and $\varepsilon_{\pi^*} \sim \mathcal{N}(0, \sigma_{\varepsilon_{\pi^*}}^2)$ independent of each other as well as the fundamentals. In the second stage, the CB privately observes the true η and π^*, and the rest of the game proceeds as usual. In any "subgame" following realizations η' and π'^*, the unique increasing linear equilibrium is given by (19) and (20), with σ_η^2 and $\sigma_{\pi^*}^2$ replaced by the corresponding variances from (11) and (12). Hence, the CB's expected utility conditional on any realizations η' and π'^* is given by (21), using the new computation of L and substituting σ_η^2 and $\sigma_{\pi^*}^2$ with the corresponding variances from (11) and (12). As this "interim" expected utility is independent of η' and π'^* (because the variances (11) and (12) do not depend on η' and π'^*), welfare—ex-ante expected utility, which simply takes an expectation of interim expected utility over both η' and π'^*—is just the interim expected utility. *Q.E.D.*

Proof of Proposition 6. See Appendix D. *Q.E.D.*

B. Controlling the CB's information

Here we extend the analysis of Subsection 4.3 to intermediate information policies. Specifically, suppose that rather than observing η and π^* directly, the CB observes noisy signals $\eta' = \eta + \delta_\eta$ and $\pi'^* = \pi^* + \delta_{\pi^*}$, with δ_η and δ_{π^*} independently normally distributed with means of 0 and respective variances of $\sigma^2_{\delta_\eta}$, $\sigma^2_{\delta_{\pi^*}}$. The CB's posteriors will be exactly those described in Equation 11 and Equation 12, with the obvious adjustment of replacing $\sigma^2_{\varepsilon_\eta}$ and $\sigma^2_{\varepsilon_{\pi^*}}$ with $\sigma^2_{\delta_\eta}$ and $\sigma^2_{\delta_{\pi^*}}$ respectively.

Let $\tilde{\eta}$ and $\tilde{\pi}^*$ denote the posterior means of the CB's beliefs on η and π^*, and let $\tilde{\sigma}^2_\eta$ and $\tilde{\sigma}^2_{\pi^*}$ denote the posterior variances.[22] The ex-ante variances of these posterior means are

$$\mathrm{Var}(\tilde{\eta}) = \frac{\sigma^4_\eta}{\sigma^2_\eta + \sigma^2_{\delta_\eta}}, \ \mathrm{Var}(\tilde{\pi}^*) = \frac{\sigma^4_{\pi^*}}{\sigma^2_{\pi^*} + \sigma^2_{\delta_{\pi^*}}},$$

which are lower than the ex-ante variances of η and π^*. This is intuitive: if and only if the CB puts non-zero weight on its prior when forming its posterior (about either η or π^*), the posterior mean will vary less than the variance in the prior.

In a linear equilibrium with public beliefs as in Equation 5 and Equation 6, the CB maximizes the expectation of its objective (4) over choice of m given its information. We get the following strategy analagous to Equation 7:

$$m(\eta', \pi'^*) = \tilde{\eta} + \tilde{\pi}^* + \gamma L, \tag{26}$$

where we suppress the dependence of $\tilde{\eta}$ and $\tilde{\pi}^*$ on η' and π'^* respectively.

Lemma 6. *Suppose the CB uses the strategy in Equation 26. Conditional on m, the public's posterior belief on η is normally distributed with mean*

$$\mu_{\eta|m} = \frac{\mathrm{Var}(\tilde{\eta})}{\mathrm{Var}(\tilde{\eta}) + \mathrm{Var}(\tilde{\pi}^*)} m + \mu_\eta - \frac{\mathrm{Var}(\tilde{\eta})}{\mathrm{Var}(\tilde{\eta}) + \mathrm{Var}(\tilde{\pi}^*)} \left(\gamma L + \mu_\eta + \mu_{\pi^*}\right). \tag{27}$$

[22] Standard results about normal-normal updating yield

$$\tilde{\eta} = \frac{\mu_\eta \sigma^2_{\delta_\eta} + \eta' \sigma^2_\eta}{\sigma^2_\eta + \sigma^2_{\varepsilon_\eta}}, \ \tilde{\pi}^* = \frac{\mu_{\pi^*} \sigma^2_{\delta_{\pi^*}} + \pi'^* \sigma^2_{\pi^*}}{\sigma^2_{\pi^*} + \sigma^2_{\delta_{\pi^*}}}, \ \tilde{\sigma}^2_\eta = \frac{\sigma^2_\eta \sigma^2_{\delta_\eta}}{\sigma^2_\eta + \sigma^2_{\delta_\eta}}, \ \tilde{\sigma}^2_{\pi^*} = \frac{\sigma^2_{\pi^*} \sigma^2_{\delta_{\pi^*}}}{\sigma^2_{\pi^*} + \sigma^2_{\delta_{\pi^*}}}.$$

Proof. Omitted, as it is entirely analogous to that of Lemma 1. *Q.E.D.*

Matching coefficients in Equation 5 and Equation 27 gives us the equilibrium constant

$$L = \frac{\mathrm{Var}(\tilde{\eta})}{\mathrm{Var}(\tilde{\eta}) + \mathrm{Var}(\tilde{\pi}^*)}. \tag{28}$$

For any underlying η and π^*, excess inflation is $\pi - \pi^* = (\tilde{\eta} - \eta) + (\tilde{\pi}^* - \pi^*) + \gamma L$, which has a mean of γL and a variance of $\tilde{\sigma}_\eta^2 + \tilde{\sigma}_{\pi^*}^2$. As output is zero on average, this gives an expected welfare of $-(\gamma^2 L^2 + \tilde{\sigma}_\eta^2 + \tilde{\sigma}_{\pi^*}^2)/2$, which, by Equation 28, is equivalent to

$$-\frac{1}{2}\left[\gamma^2\left(\frac{\frac{\sigma_\eta^4}{\sigma_\eta^2+\sigma_{\delta_\eta}^2}}{\frac{\sigma_\eta^4}{\sigma_\eta^2+\sigma_{\delta_\eta}^2} + \frac{\sigma_{\pi^*}^4}{\sigma_{\pi^*}^2+\sigma_{\delta_{\pi^*}}^2}}\right)^2 + \frac{\sigma_\eta^2\sigma_{\delta_\eta}^2}{\sigma_\eta^2 + \sigma_{\varepsilon_\eta}^2} + \frac{\sigma_{\pi^*}^2\sigma_{\delta_{\pi^*}}^2}{\sigma_{\pi^*}^2 + \sigma_{\delta_{\pi^*}}^2}\right]. \tag{29}$$

Expression (29) is decreasing in $\sigma_{\delta_{\pi^*}}^2$. Hence:

Corollary 3. *Welfare is higher when the CB is more informed about π^*, i.e., when $\sigma_{\delta_{\pi^*}}^2$ is lower.*

As suggested by the discussion in Subsection 4.3 of the extreme cases, reducing information about η by increasing $\sigma_{\delta_\eta}^2$ has complicated effects on (29) depending on the parameters. As depicted in Figure 4, numerical simulations suggest that (i) when γ is low, increasing $\sigma_{\delta_\eta}^2$ seems unambiguously bad; (ii) as we raise γ, payoffs become nonmonotonic in $\sigma_{\delta_\eta}^2$: falling, increasing, then decreasing again; (iii) for sufficiently high γ, we get an interior solution in $\sigma_{\delta_\eta}^2$ on the increasing range; and (iv) taking γ even higher, it seems that payoffs eventually increase and then decrease in $\sigma_{\delta_\eta}^2$, or perhaps even increase over the whole range.

C. Uncertainty about the Output Target

As in Subsection 5.1, let the CB's objective (and welfare) be given by (17):

$$-\upsilon\frac{(y-\overline{y})^2}{2} - \frac{(\pi-\pi^*)^2}{2} + \frac{1}{2}\upsilon\overline{y}^2.$$

In this section we sketch the effects of transparency about an uncertain output target, $\overline{y}$. Take $\overline{y}$ and η to be independently normally distributed, with means $\mu_{\overline{y}}$, μ_η and variances $\sigma_{\overline{y}}^2 > 0$, $\sigma_\eta^2 > 0$. Take $\upsilon > 0$ and $\pi^* \in \mathbb{R}$ to be commonly known. The CB observes the realizations of $\overline{y}$ and η prior to choosing the policy m.

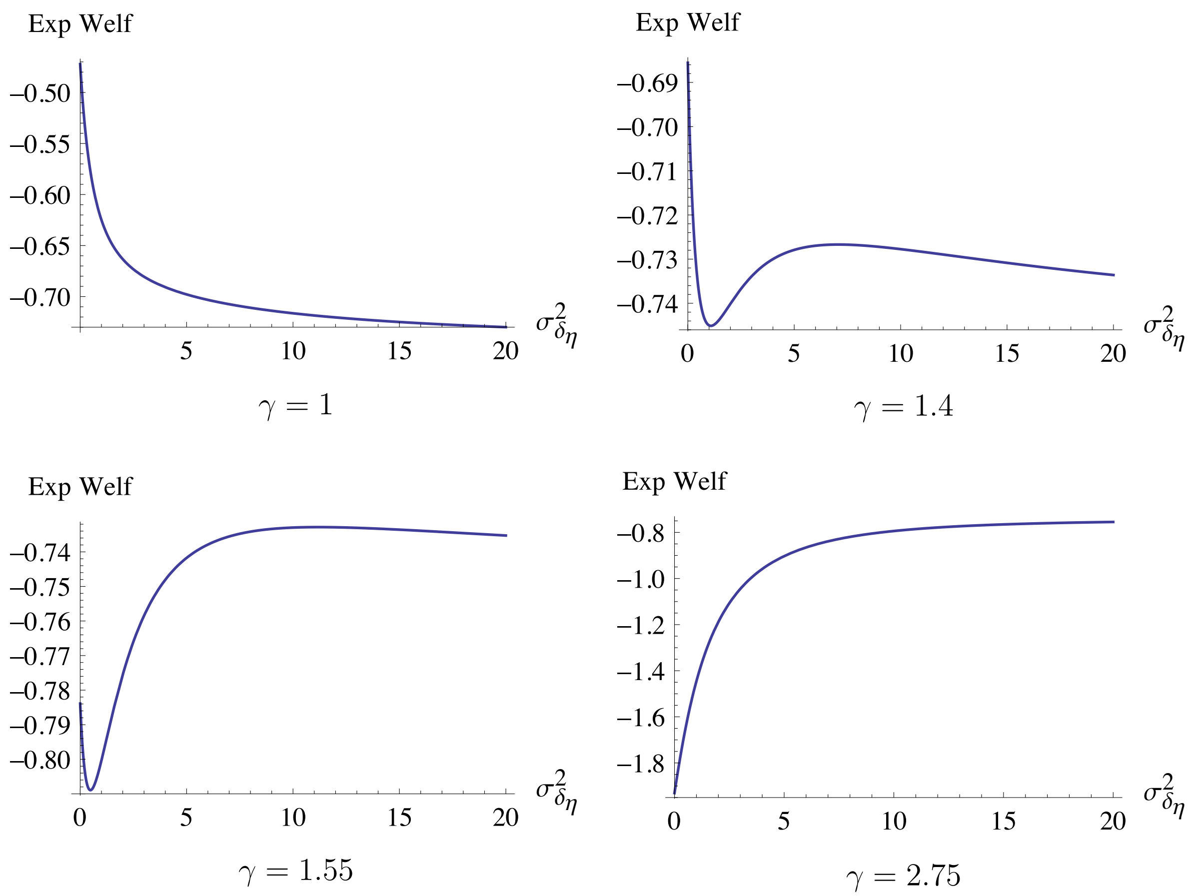

Here we plot expected welfare as a function of $\sigma^2_{\delta_\eta}$ for different values of γ. In all cases we take $\sigma^2_{\delta_{\pi^*}} = 1$, $\sigma^2_\eta = 1$, and $\sigma^2_{\pi^*} = 1$.

Figure 4 – Precision of the CB's observation of η has ambiguous welfare effects.

We consider two information structures for the public. Under (full) opacity about the output target, the public observes no additional signals. Under (full) transparency about the output target, the public observes the output target $\overline{y}$.

Preference opacity. First, let us consider opacity. Conjecture an increasing linear equilibrium in which $\hat{\eta}(m) = Lm + K$ for some $L > 0$ and K. The CB's best response is to choose a policy $m(\eta, \overline{y}) = \frac{1+\upsilon L}{1+\upsilon L^2}\eta + \frac{\upsilon L}{1+\upsilon L^2}\overline{y} + \frac{\pi^* - \upsilon LK}{1+\upsilon L^2}$. Going through a similar argument as in Subsection 5.1—solving for beliefs conditional on the CB's strategy and then matching coefficients to find a fixed point—one finds that the equilibrium value of L must solve

$$\frac{\sigma_{\overline{y}}^2}{\sigma_\eta^2} L^3 \upsilon - (1 + L\upsilon)(1 - L) = 0. \tag{30}$$

There is a unique positive solution; it satisfies (i) $L \in (0, 1)$, (ii) L is decreasing in $\sigma_{\overline{y}}^2/\sigma_\eta^2$, and (iii) $L \to 0$ as $\sigma_{\overline{y}}^2/\sigma_\eta^2 \to \infty$, while $L \to 1$ as $\sigma_{\overline{y}}^2/\sigma_\eta^2 \to 0$.

Equation 30 can be plugged back in to the CB's equilibrium strategy, and that can be used to compute welfare. With some straightforward algebraic manipulations (details of which are available from the authors), it can be shown that welfare under opacity is given by the following expression:

$$\frac{1 - L}{2L}\sigma_\eta^2 - \frac{L^2}{2}\upsilon^2 \mu_{\overline{y}}^2. \tag{31}$$

Preference transparency. Now consider transparency about the output target, in which $\overline{y}$ is revealed to the public. For any realization of $\overline{y}$, transparency corresponds to the game above with $\mu_{\overline{y}}$ taken as $\overline{y}$, and with $\sigma_{\overline{y}}^2$ taken as 0 (which implies an equilibrium value of $L = 1$). The CB's payoff under transparency given any $\overline{y}$ is therefore $-\frac{1}{2}\upsilon^2\overline{y}^2$. Taking ex-ante expectation over the distribution of $\overline{y}$, transparency yields a welfare of

$$-\frac{1}{2}\upsilon^2(\sigma_{\overline{y}}^2 + \mu_{\overline{y}}^2). \tag{32}$$

Since the L in expression (31) is in $(0, 1)$, we see that (31) is larger than (32).[23] Hence, we have shown:

Proposition 7. *When preference uncertainty concerns the output target, opacity about preferences yields higher welfare than transparency about preferences.*

[23] Note that the value of the expressions would be equal when $\sigma_{\overline{y}}^2 = 0$, in which case $L = 1$ in (31). The case of $\sigma_\eta^2 = 0$ requires more care, because then (i) $L = 0$ under opacity and (ii) the CB's payoff under transparency given any $\overline{y}$ is no longer $-\frac{1}{2}\upsilon^2\overline{y}^2$. One can show that when $\sigma_\eta^2 = 0$, welfare is 0 under both opacity and transparency.

D. Analysis for Subsection 5.3

This appendix provides details for Subsection 5.3, including a proof of Proposition 6 .

D.1. Preliminaries

Throughout this section inflation is determined by (24):

$$\pi = (1-r)(m-\eta) + r\hat{\pi}.$$

Taking expectations of both sides gives $\hat{\pi} = (1-r)(m-\eta) + r\hat{\pi}$. Rearranging:

$$\hat{\pi} = m - \hat{\eta}. \tag{33}$$

Substituting (33) back into (24) yields

$$\begin{aligned} \pi &= (1-r)(m-\eta) + r(m-\hat{\eta}) \\ &= m - (1-r)\eta - r\hat{\eta}. \end{aligned} \tag{34}$$

Output is determined by $y = \pi - \hat{\pi}$. Plugging in for π and $\hat{\pi}$ from (24) and (33):

$$\begin{aligned} y &= (1-r)(m-\eta) + r\hat{\pi} - \hat{\pi} \\ &= (1-r)(m-\eta-m+\hat{\eta}) \\ &= (1-r)(\hat{\eta}-\eta). \end{aligned} \tag{35}$$

The CB maximizes $\gamma y - (\pi - \pi^*)^2/2$ which, plugging in (35) and (34), reduces to

$$\gamma(1-r)(\hat{\eta}-\eta) - (m-(1-r)\eta - r\hat{\eta} - \pi^*)^2/2. \tag{36}$$

D.2. Common Knowledge of η or π^*

Common knowledge of η. Here, it holds that $\hat{\eta} = \eta$. Substituting in, the first-order condition (FOC) for maximizing (36) over m yields $m = \eta + \pi^*$. Hence, inflation expectations are $\hat{\pi} = \pi^*$, inflation is π^*, output is 0, and welfare is 0. This is the same as in Proposition 1 part 1.

Common knowledge of π^* but not η. We seek a linear separating equilibrium with $m = k\eta + l\pi^* + b$ for some $k \neq 0$. In that case $\hat{\eta} = \frac{m - l\pi^* - b}{k}$. Plugging $\hat{\eta}$ into (36) gives a maximization problem of

$$\max_m \gamma(1-r)\left(\frac{m-l\pi^*-b}{k}-\eta\right)-\left(m-(1-r)\eta-r\frac{m-l\pi^*-b}{k}-\pi^*\right)^2/2.$$

Taking a FOC and solving for m gives

$$m=\eta\frac{(1-r)k}{k-r}+\frac{k^2\pi^*-bkr-k\pi^*r-kl\pi^*r+br^2+l\pi^*r^2+k\gamma-kr\gamma}{(k-r)^2}.$$

Matching coefficients on η,

$$k=\frac{(1-r)k}{k-r} \implies k=1.$$

Plugging $k=1$ into the expression for m gives

$$m=\eta+\pi^*\frac{1-lr}{1-r}+\frac{-br+\gamma}{1-r}.$$

Matching coefficients on l and b gives

$$l=\frac{1-lr}{1-r} \implies l=1,$$

$$b=\frac{-br+\gamma}{1-r} \implies b=\gamma.$$

So, we have a linear equilibrium with $m=\eta+\pi^*+\gamma$. The belief is $\hat{\eta}=m-\pi^*-\gamma=\eta$, and so $\pi=\hat{\pi}=m-\eta=\pi^*+\gamma$. This is the same as in Proposition 1 part 2.

D.3. Uncertainty about η and π^*

Now assume η and π^* are independent, with $\eta\sim\mathcal{N}(\mu_\eta,\sigma_\eta^2)$ and $\pi^*\sim\mathcal{N}(\mu_{\pi^*},\sigma_{\pi^*}^2)$. We consider linear equilibria in which the public's expectations $\hat{\eta}\equiv\mathbb{E}[\eta|m]$ and $\hat{\pi}\equiv\mathbb{E}[\pi|m]$ are given by

$$\begin{aligned}\hat{\eta}(m)&=Lm+K,\\ \hat{\pi}(m)&=m-\hat{\eta}(m)=(1-L)m-K,\end{aligned}$$

for some constants L and K. Plugging in to (36), the CB's objective can be written as

$$\begin{aligned}&\gamma(1-r)(\hat{\eta}-\eta)-(m-(1-r)\eta-r\hat{\eta}-\pi^*)^2/2\\ &\quad=\gamma(1-r)(Lm+K-\eta)-(m-(1-r)\eta-r(Lm+K)-\pi^*)^2/2.\end{aligned}$$

Maximizing over m, the FOC yields

$$m(\eta, \pi^*) = \eta \frac{1-r}{1-Lr} + \pi^* \frac{1}{1-Lr} + \frac{Kr(1-Lr) + L(1-r)\gamma}{(1-Lr)^2}. \tag{37}$$

Lemma 7. *Suppose the CB uses the strategy in Equation 37. Conditional on m, the public's posterior belief on η is normally distributed with mean*

$$\mu_{\eta|m} = m \frac{(1-r)(1-Lr)}{\frac{\sigma^2_{\pi^*}}{\sigma^2_\eta} + (1-r)^2} + \frac{(1-Lr)\mu_\eta \sigma^2_{\pi^*} - (1-Lr)\mu_{\pi^*}(1-r)\sigma^2_\eta - (1-r)\sigma^2_\eta(Kr(1-Lr) + L(1-r)\gamma)}{(1-Lr)((1-r)^2\sigma^2_\eta + \sigma^2_{\pi^*})}.$$

Proof. Given the policy function (37), conditional on η, the distribution of m is normal with mean $\eta \frac{1-r}{1-Lr} + \mu_{\pi^*} \frac{1}{1-Lr} + \frac{Kr(1-Lr)+L(1-r)\gamma}{(1-Lr)^2}$ and variance $\frac{\sigma^2_{\pi^*}}{(1-Lr)^2}$; equivalently, $m' \equiv \frac{m(1-Lr)}{1-r} - \mu_{\pi^*}\frac{1}{1-r} - \frac{Kr(1-Lr)+L(1-r)\gamma}{(1-Lr)(1-r)}$ has mean η and variance $\frac{\sigma^2_{\pi^*}}{(1-r)^2}$. Standard results on updating normal beliefs with a normal signal imply a posterior mean on η conditional on m' of

$$\frac{\frac{\mu_\eta}{\sigma^2_\eta} + \frac{m'(1-r)^2}{\sigma^2_{\pi^*}}}{\frac{1}{\sigma^2_\eta} + \frac{(1-r)^2}{\sigma^2_{\pi^*}}}.$$

Substituting in for m' yields the lemma's conclusion. *Q.E.D.*

We can match coefficients from the equations for $\hat{\eta}$ and $\mu_{\eta|m}$ to solve for L and K:

$$L = \frac{(1-r)(1-Lr)}{\frac{\sigma^2_{\pi^*}}{\sigma^2_\eta} + (1-r)^2} \quad \implies L = \frac{(1-r)\sigma^2_\eta}{(1-r)\sigma^2_\eta + \sigma^2_{\pi^*}},$$

$$\begin{aligned} K &= \frac{(1-Lr)\mu_\eta \sigma^2_{\pi^*} - (1-Lr)\mu_{\pi^*}(1-r)\sigma^2_\eta - (1-r)\sigma^2_\eta(Kr(1-Lr) + L(1-r)\gamma)}{(1-Lr)((1-r)^2\sigma^2_\eta + \sigma^2_{\pi^*})} \\ \implies \quad K &= \frac{(1-Lr)(-\mu_{\pi^*}(1-r)\sigma^2_\eta + \mu_\eta\sigma^2_{\pi^*}) - L(1-r)^2\sigma^2_\eta\gamma}{(1-Lr)((1-r)\sigma^2_\eta + \sigma^2_{\pi^*})}. \end{aligned}$$

Now plug L and K back into the policy function (37) to get

$$m(\eta, \pi^*) = \eta \frac{(1-r)((1-r)\sigma^2_\eta + \sigma^2_{\pi^*})}{(1-r)^2\sigma^2_\eta + \sigma^2_{\pi^*}} + \pi^* \frac{(1-r)\sigma^2_\eta + \sigma^2_{\pi^*}}{(1-r)^2\sigma^2_\eta + \sigma^2_{\pi^*}} + \frac{(1-r)^2\sigma^2_\eta\gamma + \mu_\eta r\sigma^2_{\pi^*} - \mu_{\pi^*}(1-r)r\sigma^2_\eta}{(1-r)^2\sigma^2_\eta + \sigma^2_{\pi^*}}.$$

Plugging in $m = m(\eta, \pi^*)$ and $\hat{\eta} = Lm + K$ into (34), we get realized inflation of

$$\pi = m - (1-r)\eta - r\hat{\eta} = \pi^* + \frac{(1-r)^2\sigma_\eta^2\gamma}{(1-r)^2\sigma_\eta^2 + \sigma_{\pi^*}^2}.$$

Notice that excess inflation is independent of η and π^* in equilibrium. Welfare is thus

$$-\mathbb{E}[(\pi - \pi^*)^2/2] = -\left(\frac{(1-r)^2\sigma_\eta^2\gamma}{(1-r)^2\sigma_\eta^2 + \sigma_{\pi^*}^2}\right)^2/2.$$

All of the comparative statics on welfare with respect to $\sigma_{\pi^*}^2$ and σ_η^2 are unchanged with the addition of the r term: welfare increases in $\sigma_{\pi^*}^2$ and decreases in σ_η^2.

We can also take comparative statics on r. Excess inflation is positive for $r = 0$ (the benchmark case from the main text), and decreases to 0 as r goes to 1. Hence, welfare improves as inflation expectations rather than underlying policy start to drive inflation. In the limit as expectations become self-fulfilling, we approach the commitment outcome with no excess inflation.

Notice that in this limit as $r \to 1$, we have $m = \pi^* + \mu_\eta$: policy is pushing inflation towards $m - \eta = \pi^* + (\mu_\eta - \eta)$. So on average (taking expectation over η), policy pushes inflation towards π^*, but towards a higher actual value when η is small and a lower actual value when η is large. However, policy becomes irrelevant in the limit as $r \to 1$, and only expectations matter. The policy signals the value of π^*, and inflation expectations—and therefore realized inflation—become π^*. In other words, in the limit game with $r = 1$, policy m is just a cheap-talk message. Realized inflation is $\pi = \hat{\pi}$, and so m affects the CB's objective only through the implied beliefs on $\hat{\pi}$: the CB chooses m to maximize $-(\hat{\pi}(m) - \pi^*)^2/2$. There is no longer any credibility problem, as the CB is willing to choose m to signal $\hat{\pi} = \pi^*$. So under any proposed bijection of m into $\hat{\pi}(m)$, there would be a separating cheap talk equilibrium with m chosen so that $\hat{\pi}(m) = \pi^*$ for each realized π^*. The limit as $r \to 1$ selects one such separating equilibrium with beliefs $\hat{\pi}(m) = m - \mu_\eta$ and policy $m = \pi^* + \mu_\eta$.

References

ANGELETOS, G.-M., C. HELLWIG, AND A. PAVAN (2006): "Signaling in a global game: Coordination and policy traps," *Journal of Political Economy*, 114, 452–484.

BACKUS, D. AND J. DRIFFILL (1985): "Inflation and Reputation," *American Economic Review*, 75, 530–538.

BAERISWYL, R. AND C. CORNAND (2010): "The signaling role of policy actions," *Journal of Monetary Economics*, 57, 682–695.

BAGWELL, K. (1995): "Commitment and Observability in Games," *Games and Economic Behavior*, 8, 271–280.

BARRO, R. J. AND D. B. GORDON (1983a): "A Positive Theory of Monetary Policy in a Natural Rate Model," *Journal of Political Economy*, 91, 589–610.

——— (1983b): "Rules, Discretion and Reputation in a Model of Monetary Policy," *Journal of Monetary Economics*, 12, 101–121.

BÉNABOU, R. AND J. TIROLE (2006): "Incentives and Prosocial Behavior," *American Economic Review*, 96, 1652–1678.

BERNANKE, B. S. (2003): "A Perspective on Inflation Targeting," Remarks at the Annual Washington Policy Conference of the National Association of Business Economists, Washington, D.C.

——— (2007): "Inflation Expectations and Inflation Forecasting," Speech at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute.

——— (2013): "Communication and monetary policy," Speech at the National Economists Club Annual Dinner, Herbert Stein Memorial Lecture, Washington, DC.

BLINDER, A. S., M. EHRMANN, M. FRATZSCHER, J. DE HAAN, AND D.-J. JANSEN (2008): "Central Bank Communication and Monetary Policy: A Survey of Theory and Evidence," *Journal of Economic Literature*, 46, 910–45.

CANZONERI, M. B. (1985): "Monetary policy games and the role of private information," *American Economic Review*, 75, 1056–1070.

CUKIERMAN, A. AND A. H. MELTZER (1986): "A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information," *Econometrica*, 54, 1099–1128.

DEGROOT, M. H. (1970): *Optimal Statistical Decisions*, New York: McGraw-Hill Book Company.

ELLINGSEN, T. AND U. SÖDERSTRÖM (2001): "Monetary Policy and Market Interest Rates," *American Economic Review*, 1594–1607.

FAUST, J. AND L. E. O. SVENSSON (2001): "Transparency and Credibility: Monetary Policy With Unobservable Goals," *International Economic Review*, 42, 369–397.

FAUST, J., E. T. SWANSON, AND J. H. WRIGHT (2004): "Do Federal Reserve Policy Surprises Reveal Superior Information about the Economy?" *Contributions in Macroeconomics*, 4.

FISCHER, P. E. AND R. E. VERRECCHIA (2000): "Reporting bias," *The Accounting Review*, 75, 229–245.

FRANKEL, A. AND N. KARTIK (2015): "Muddled Information," Unpublished.

GERAATS, P. M. (2002): "Central Bank Transparency," *Economic Journal*, F532–F565.

——— (2007): "Commitment, Transparency and Monetary Policy," Unpublished.

——— (2008): "ECB Credibility and Transparency," Mimeo.

GÜRKAYNAK, R. S., B. SACK, AND E. SWANSON (2005): "The Sensitivity of Long-term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models," *American Economic Review*, 425–436.

KYDLAND, F. E. AND E. C. PRESCOTT (1977): "Rules rather than Discretion: The inconsistency of optimal plans," *Journal of Political Economy*, 473–491.

MAGGI, G. (1999): "The Value of Commitment with Imperfect Observability and Private Information," *The RAND Journal of Economics*, 555–574.

MELOSI, L. (2014): "Signaling Effects of Monetary Policy," Unpublished.

MERTENS, E. (2011): "Managing Beliefs about Monetary Policy under Discretion," Unpublished.

MORRIS, S. AND H. S. SHIN (2002): "Social value of Public Information," *American Economic Review*, 92, 1521–1534.

MOSCARINI, G. (2007): "Competence Implies Credibility," *American Economic Review*, 97, 37–63.

NAKAMURA, E. AND J. STEINSSON (2013): "High Frequency Identification of Monetary Non-Neutrality," Unpublished.

POOLE, W. (2005): "Remarks: panel on "After Greenspan: Whither Fed Policy?"," Remarks delivered at the Western Economics Association International conference, San Francisco, California, July 6, 2005.

PRAT, A. (2005): "The Wrong Kind of Transparency," *American Economic Review*, 95, 862–877.

ROGOFF, K. (1985): "The Optimal Degree of Commitment to an Intermediate Monetary Target," *Quarterly Journal of Economics*, 100, 1169–1189.

ROMER, C. D. AND D. H. ROMER (2000): "Federal Reserve Information and the Behavior of Interest Rates," *American Economic Review*, 90, 429–457.

SIBERT, A. (2009): "Is Transparency about Central Banks Desirable?" *Journal of the European Economic Association*, 7, 831–857.

SVENSSON, L. E. O. (1997): "Optimal Inflation Targets, "Conservative" Central Banks, and Linear Inflation Contracts," *American Economic Review*, 87, 98–114.

TAMURA, W. (2014): "Optimal Monetary Policy and Transparency under Informational Frictions," Unpublished.

——— (2015): "Endogenous Information and Central Bank Transparency," Unpublished.

TANG, J. (2013): "Uncertainty and the Signaling Channel of Monetary Policy," Unpublished.

WALSH, C. E. (1995): "Optimal Contracts for Central Bankers," *American Economic Review*, 85, 150–167.

——— (2007): "Optimal Economic Transparency," *International Journal of Central Banking*, 3.