Ambiguity, Monetary Policy and Trend Inflation[*]

Riccardo M. Masolo   Francesca Monti

Bank of England and CfM

September 9, 2016

Abstract

We build a model that can explain the evolution of trend inflation in the US as a function of the changes in the private sector's confidence in their understanding of the monetary policymaker's behavior, measured by the dispersion in the individual SPF nowcasts of the 3-month T-Bill rate. Rather than resorting to exogenous changes in the inflation target process, the model features ambiguity-averse agents and ambiguity regarding the conduct of monetary policy. Ambiguity turns out to be key in explaining the difference between the inflation target pursued by the central bank and the inflation trend measured in the data.

JEL Classification: D84, E31, E43, E52, E58

Keywords: Ambiguity aversion, monetary policy, trend inflation

1   Introduction

The dynamics of inflation and, in particular, its persistence are driven in large part by a low-frequency, or trend, component, as documented for example in Stock and Watson (2007) and Cogley and Sbordone (2008). Most of the macroeconomic literature relies, however, on models that are approximated around a zero-inflation steady state[1] and that, consequently, cannot capture the persistent dynamic properties of inflation. Positive

---

[1]Or alternatively, and equivalently from this perspective, a full indexation steady-state.

steady-state inflation also has far-reaching effects on the static and dynamic properties of the model (see Ascari and Sbordone, 2014, for an overview), including a dramatic impact on the lowest degree of inflation responsiveness that can ensure equilibrium determinacy, as highlighted in Coibion and Gorodnichenko (2011). It is therefore crucial to have models that can account for trend inflation and its dynamics.

In order to capture the dynamics of inflation, Del Negro and Eusepi (2011) and Del Negro, Giannoni and Schorfheide (2015) propose replacing the constant inflation target with a time-varying very persistent but otherwise exogenous inflation target process. While this assumption helps match the data better, it does not provide a compelling explanation for the low frequency movements in inflation. We argue, in fact, that the assumption of an exogenous slow-moving inflation target is counterfactual. While the Federal Reserve only explicitly adopted an inflation target in 2012, we have reasons to believe, as we illustrate in Section 2, that inflation targets had been more or less implicitly in place long before that. Moreover, if inflation targets vary at all, they do so in a discrete and infrequent fashion and do not represent the same concept as trend inflation.

In this paper, we present a simple model that can explain the persistent differences between *target* and *trend* inflation. The key driver of this wedge is a time-varying degree of private sector's confidence in the conduct of monetary policy, which we model as ambiguity about the monetary policy rule. Ambiguity describes a situation in which there is uncertainty about the probability distribution over states of the world: agents entertain as possible not one but a set of beliefs and they are unable to assign probabilities to each of them. We augment a prototypical new-Keynesian model by 1) introducing ambiguity about the monetary policy rule and 2) assuming that agents are averse to ambiguity. We introduce ambiguity aversion in the model using a recursive version of the multiple prior preferences (see Gilboa and Schmeidler, 1989, and Epstein and Schneider, 2003), pioneered in business cycle models by Ilut and Schneider (2014).

Ambiguity-averse agents will base their decisions on the worst-case scenario consistent with their belief set. In practice, in our model agents base their decisions on a distorted belief about the interest rate, which differs from the one actually set by the central bank. As a result, inflation in steady state does not necessarily coincide with the target. Our model delivers a characterization of the relationship between inflation trend and target, which depends on the degree of ambiguity about monetary policy – that is on how wide the private sector's set of beliefs about the interest is – and the responsiveness of the policy rate to inflation. To keep things simple and implications stark, we work under the assumption that the policy response function followed by the Fed *never* changed over the 35-year period we consider. The observed changes in the behavior of inflation (both at high and low frequency) are, therefore, the result of differing degrees of confidence the private sector has in the Fed's policy.

With a very standard calibration and using data on the dispersion of interest rate nowcasts available in the *Survey of Professional Forecasters* as a measure for ambiguity about monetary policy, our model can explain the dynamics of trend inflation in the US since the early 1980s (the beginning of our sample) without resorting to exogenous changes in the inflation target. The model can match the dynamics of trend inflation

in the period before the Great Recession, when trend inflation was mostly above target but falling, as well as the low trend inflation of the post-crisis period.[2] We show that, before the crisis, the worst-case scenario was one in which the private sector feared that the interest rate would be lower than implied by the actual policy rule, thus resulting in above-target inflation. Our model explains the recent low level of trend inflation as a consequence of the proximity of the policy rate to the zero lower bound, which sets a floor to the distortion of the beliefs about the interest rate. If the asymmetry generated by the zero lower bound is large enough, then agents will make their consumption-saving decision based on a distorted belief that the interest rate is above the one actually set by the central bank, and this will generate lower inflation.

Our work can also explain the switch from indeterminacy to determinacy during the early years of the Volcker chairmanship without resorting to changes in the responsiveness of policy to inflation, in line with Coibion and Gorodnichenko (2012), or even to changes in the target.

The rest of the paper is organized as follows. In section 2 we discuss whether assuming a constant inflation target is a reasonable characterisation of the definition of the price stability part of the Fed's dual mandate and we present evidence regarding changes in confidence about the conduct of monetary policy and the plausible link with increases in transparency. Section 3 provides a description of the model we use for our analysis, characterizes the steady state of our economy as a function of the degree of ambiguity and studies optimal monetary policy in the presence of ambiguity. In Section 4 we show how our simple model can match the dynamics of trend inflation, while Section 5 concludes.

2   Empirical and narrative evidence

In this Section we first discuss evidence on the low-frequency component of inflation. We then present evidence in support of our claim that these low frequency movements do not seem to be driven by variations in the inflation target. Finally we discuss how measures of confidence about the conduct of monetary policy instead have been increasing, plausibly linked to the increases in transparency over the last three decades.

Trend inflation and the inflation target   A vast number of papers propose estimates of trend inflation based on different models.[3] And invariably, they have all declined since the early 1980s, as a result of the fall in headline inflation. We take as benchmark the estimate of trend inflation we obtain with a bayesian vector autoregression model with drifting coefficients, along the lines of Cogley and Sargent (2002). In particular we use the specification presented in Cogley and Sbordone (2008), using the same four data series: implicit GDP deflator, real GDP growth, unit labor cost to approximate a measure of marginal cost, and the fed funds rate on a discount basis. Figure 1 reports it, along with the mean inflation over the sample and the 90% confidence bands. Clearly, inflation is

---

[2]There is some evidence that trend inflation might have fallen below 2% since 2009-2010, according to several measures shown in Garciga (2016).

[3]See for example Clark and Doh (2013) for a review of the main models in use and their relative forecasting performance.

Figure 1: Inflation, mean inflation and trend inflation
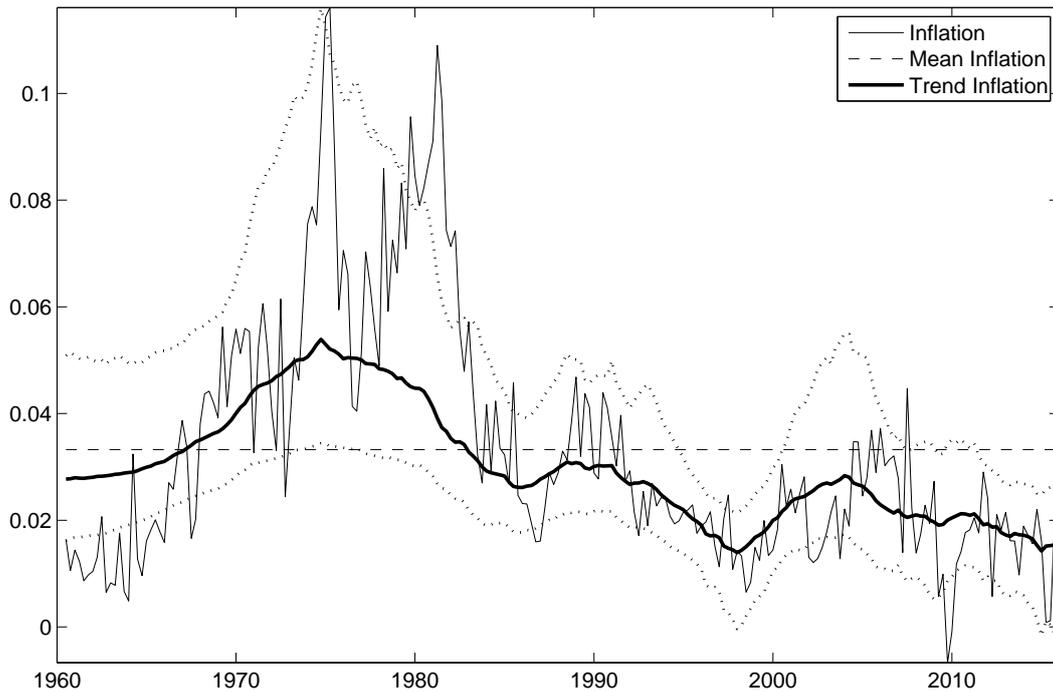
characterized by a trend component, which has fallen since the early 1980s and is currently estimated to be slightly below target. The aim of our paper is to show how it is possible to match its dynamics without resorting to an exogenous inflation target shock.

Models that explain the dynamics of inflation with exogenous variation in the inflation target seem at odds with evidence from various sources, including the *Blue Book*, a document about monetary policy alternatives presented to the committee by Fed staff before each FOMC meeting. While the Federal Reserve officially did not have a target value for inflation until 2012, the *Blue Book* simulations have been produced assuming targets of 1.5% and 2% since at least 2000. In his book "A Modern History of FOMC Communication: 1975-2002", Lindsey (2003) states that, as early as July 1996, numerous FOMC committee members had indicated at least an informal preference for an inflation rate in the neighborhood of 2%, as indicated by FOMC transcripts.[4] Goodfriend (2003) provides detailed evidence that the Fed had *implicitly* adopted an inflation target in the 1980s, as this excerpt makes clear:

Chairman Greenspan testified in 1989 in favor of a qualitative zero inflation objective for the Fed defined as a situation in which "the expected rate of change of the general level of prices ceases to be a factor in individual and business decisionmaking. Thus, it is reasonable to think that the Greenspan Fed set out to achieve low enough inflation to make that definition of price stability a reality. [...] This is the first sense in which it is plausible to think that the Greenspan Fed has adopted an implicit form of inflation targeting.
Goodfriend (2003), p. 11

Indeed, Orphanides (2002) supports the view that the policy goal has not funda-

---

[4]See transcripts of the July 2-3, 1996, FOMC meeting for a statement by Chairman Greenspan.

mentally changed since at least World War II, and that low and stable inflation plays a prominent role in its definition:

Did not the policymakers of the 1970s make a systematic effort to guide the economy to its non-inflationary full employment potential? This, after all, had been and remains the underlying macroeconomic policy objective of government policies in the United States since at least the end of World War II.
**Orphanides (2002), p. 1**

In sum, it seems clear that the concept of inflation trend differs from that of inflation target. We see the latter as changing, if at all, at few and relatively distant points in time, while the former moves, albeit slowly, with every change in headline inflation. The rest of this paper is primarily dedicated to explaining the wedge between trend inflation and the target. In so doing we maintain a very conservative assumption that the target has never changed over the period covered by our sample, which starts in the early 1980s.

Private sector confidence and transparency. The discussion above highlights how the explicit announcement of an inflation target in 2012 can be seen as the culmination of a process geared towards greater transparency that started much earlier. The main milestones in the Federal Reserve's progress toward greater openness, according to Lindsey's (2003) detailed account of FOMC communication, include: in 1979, the first release of semiannual economic projections; in 1983, the first publication of the Beige Book, which summarizes information about economic conditions received from the Federal Reserve System's business contacts; in 1994, the decision to release a postmeeting statement when policy actions had been taken; in 2000, the beginning of the practice of issuing a statement after each meeting of the Federal Open Market Committee (FOMC) and including in the statement an assessment of the balance of risks to the Committee's objectives; and in 2002, adding the FOMC roll call vote to the postmeeting statements.

Transparency is key to boost the private sector's confidence in the conduct of monetary policy, as various papers demonstrate. Swanson (2006) shows that, since the late 1980s, U.S. financial markets and private sector forecasters have become better able to forecast the federal funds rate at horizons out to several months. Moreover, his work shows that the cross-sectional dispersion of their interest rate forecasts shrank over the same period and, importantly, also provides evidence that these phenomena can be traced back to increases in central bank transparency. Ehrmann *et al.* (2012) also find that increased central bank transparency lowers disagreement among professional forecasters. This seems natural. For a given degree of uncertainty about the state of the economy, improved knowledge about the policymakers' objectives and model will help the private sector anticipate policy responses more accurately. In other words, it translates into a reduction of ambiguity about monetary policy. More recently Boyarchenko *et al.* (2016) use high-frequency data on a wide range of asset yields and find evidence that FOMC announcements affect the prices of risky assets not only directly by announcing changes in Fed funds targets, but also indirectly by influencing the risk premium.

These findings are supported by statements by policymakers, clearly conveying the idea that they saw increased transparency and communication as both necessary and ef-

Figure 2: A measure of disagreement about the interest rate
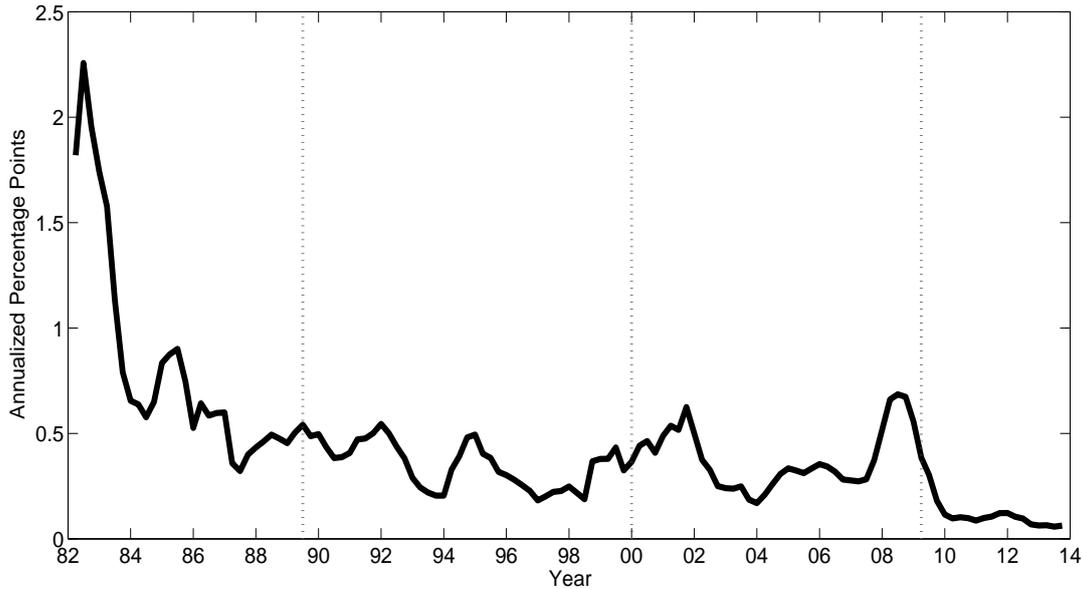
The solid line is the interdecile dispersion of SPF nowcasts of the 3-month T-Bill rate. The three dotted vertical lines indicate the various policy eras highlighted in Lindsey (2003) - 1982:1989 targeting borrowed reserves; 1990-1999 targeting thefed funds rate; 2000-2008 the "era of communication" - and the period since the Great Recession.

fective in helping the private sector anticipate policy moves:

... the faulty estimate was largely attributable to misapprehensions about the Fed's intentions. [...] Such misapprehensions can never be eliminated, but they can be reduced by a central bank that offers markets a clearer vision of its goals, its 'model' of the economy, and its general strategy.

Blinder (1998)[5]

Our model will capture exactly this mechanism by explicitly considering that private sector agents entertain multiple priors on the monetary policy rule, as a way of formalizing the *mishapprehensions* Blinder (1998) refers to.

A separate question is how to measure ambiguity. We follow the bulk of the literature, which takes forecast dispersion as a proxy for ambiguity (see for example Drechsler, 2013, and Ilut and Schneider, 2014). Works such as those by Swanson (2006) and Ehrmann *et al.* (2012) give us confidence that measuring ambiguity this way captures the gist of the effects of increased transparency on the private sector's understanding of monetary policy. Our measure of changes in confidence, reported in Figure 2, is the dispersion of survey of professional forecasters (SPF) nowcasts of the 3-month T-Bill rate, which are available from 1981Q3 onward[6]. We take a 4-quarter moving average to smooth out very high-frequency variations which would have not much to say about trends. But, clearly, the scale of the numbers, which is really what we are primarily concerned with, is unaffected.

---

[5]Also reported by Coibion and Gorodnichenko (2012).

[6]We experimented with Consensus data as well. Its dispersion is very similar to the one of the SPF in size and dynamics, but Consensus data is available only from 1993 onward, while we are interested in extending the sample back as much as possible.

Unfortunately SPF data on the federal funds rate is, to our knowledge, only available since the late 1990s, so we use the 3-month T-Bill rate as a proxy for the policy rate. We evaluate the dispersion of the nowcasts rather than forecasts at further horizons, because we aim to isolate the disagreement about monetary policy itself, rather than about the realizations of key macroeconomic variables. The nowcasts are produced in the middle of a quarter and thus incorporate a lot of information about the current state of the economy. We compared our measure of ambiguity with a commonly used measure of policy uncertainty, which is not directly based on forecasters' disagreement, *i.e.* the measure of policy uncertainty proposed by Baker, Bloom and Davis (2015). The latter series is only available starting in 1985, but the correlation over the common sample is around .51 and as high as .68, if we limit the sample to the 1980s and 1990s.

Narrative evidence from the 2007Q4 SPF survey. Evidence from a set of special questions that were included in the 2007Q4 SPF Survey also supports our thesis. Respondents were asked if they thought the Fed followed a numerical target for long-run inflation and, if so, what that value was. Respondents provided, at the same time, their expectations for inflation over the next 10 years, which we can consider as a proxy for their estimate of trend[7]. About half of the respondents, thought that the Fed had a numerical target (top row of Table 1) and, remarkably, the average numerical value provided, 1.74%, was almost exactly half way between the two values (1.5% and 2%) routinely used for *Blue Book* simulations.

Table 1: 2007 Q4 SPF Special Survey

	Targeters	Non-Targeters
Percentage of Responders	48	46
Average Target	1.74	n.a.
10-yr PCE Inflation Expectation	2.12	2.25
Short-rate Dispersion	.49	.61

Note: 6 percent of responders did not answer this question.

The comments of some respondents also shed further light on their views. It seems fair to conclude that, while the fact that policymakers aimed at low and stable inflation was well understood, there emerged varying degrees of confidence in the extent to which this goal could be achieved. The group of forecasters who thought the Fed was indeed following a numerical target (whom we dub "Targeters") expected inflation over the next 10 years to be on average .4 percent above said target which, once more, highlights the inherent difference between the two concepts of target inflation and trend inflation. From our perspective, Targeters represent the economic agents displaying the greater degree of confidence in the conduct of monetary policy. It is thus interesting to notice that non-Targeters expected inflation to be even higher than Targeters (third row of Table 1) over

---

[7]See Clark and Nakata (2008), p. 19.

the next ten years. Moreover, Non-Targeters displayed a higher degree of disagreement on the short rate than Targeters (bottom row of Table 1).

Unfortunately, the 2007Q4 set of special questions was a one-off event[8] and the limited number of respondents makes it difficult to find statistically significant differences. Yet, this survey's results support the view that lower degrees of confidence in the conduct of monetary policy tend to associate with greater degree of dispersion regarding the short-term interest and, ultimately, a higher level of trend inflation.

We will now turn to presenting our model, after which we will return to these stylized fact and illustrate how our model can explain them.

3 The Model

3.1 Setup

We modify a textbook New-Keynesian model (Galí, 2008) by assuming that the agents face ambiguity about the expected future policy rate. In so doing, our setup is related to Ilut and Schneider (2014), whose model features agents having multiple priors on the TFP process. It also relates to Benigno and Paciello (2012) who consider a robust monetary policy model[9], the main difference residing in how multiple-priors can deliver first-order effects.

While we do not include time-t subscripts on steady state variables to save on notation, we think of agents as *anticipated utility* decision makers as proposed by Kreps (1998) and as used in a number of macroeconomic applications, including Ascari and Sbordone (2014), which constitutes our main reference in terms of definition and estimation of an inflation trend.[10]

To isolate the effects of ambiguity, we set up our model so that, absent ambiguity, the first-best allocation is attained thanks to a sufficiently strong response of the central bank to inflation (the so-called Taylor principle) and to a government subsidy that corrects the distortion introduced by monopolistic competition.

Ambiguity, however, will cause steady-state, or trend, inflation to deviate from its target. For expositional simplicity the derivation of the model is carried out assuming the inflation target is zero. But the model is equivalent to one in which the central bank targets a positive constant level of inflation to which firms index their prices. The steady-state level of inflation we find, should then be interpreted as a deviation from the target.

---

[8]Another one-off set of questions were included in a 2012 SPF survey, which was administered right after the announcement of the target. The questions do not map exactly into those we discuss here, but the idea that different agents display varying degrees of confidence in the conduct of monetary policy emerges from that survey as well.

[9]In the tradition of Hansen and Sargent (2007).

[10]In previous versions of the paper we have also considered the log-linear approximation around a given steady state, which remains a viable option. The anticipated-utility specification, however, does not restrict the trend to be mean-reverting and allows us to capture how the first-order approximation coefficients may have changed over time, which is particularly important for the discussion of the determinacy region.

Households. Let $s_t \in S$ be the vector of exogenous states. We use $s^t = (s_1, ..., s_t)$ to denote the history of the states up to date t. A consumption plan $\overrightarrow{C}$ says, for every history s^t, how many units of the final good $C_t(s^t)$ a household consumes and for how many hours $N_t(s^t)$ a household works. The consumer's felicity function is:

$$u(\overrightarrow{C}_t) = \log(C_t) - \frac{N_t^{1+\psi}}{1+\psi}$$

Utility conditional on history s^t equals felicity from the current consumption and labour mix plus discounted expected continuation utility, *i.e.* the households' utility is defined recursively as

$$U_t(\overrightarrow{C}; s^t) = \min_{p \in \mathcal{P}_t(s^t)} \mathbb{E}^p \left[u(\overrightarrow{C}_t) + \beta U_{t+1}(\overrightarrow{C}; s_t, \ s_{t+1}) \right] \tag{1}$$

where $\mathcal{P}_t(s^t)$ is a set of conditional probabilities about next period's state $s_{t+1} \in S$. The recursive formulation ensures that preferences are dynamically consistent. The multiple priors functional form (1) allows modeling agents that have a set of multiple beliefs and also captures a strict preference for knowing probabilities (or an aversion to not knowing the probabilities of outcomes), as discussed in Ilut and Schneider (2014)[11]. A non-degenerate belief set $\mathcal{P}_t(s^t)$ means that agents are not confident in probability assessments, while the standard rational expectations model can be obtained as a special case of this framework in which the belief set contains only one belief. The main difference with respect to the multiplier preferences popularised by Hansen and Sargent (2007) to represent ambiguity aversion is that multiple priors utility is not smooth when belief sets differ in mean, as in our model. So, by following the multiple priors approach, we can characterize the effects of ambiguity on the steady state, while the multiplier preferences approach implies that the effect of ambiguity can be identified only by approximating the model at higher orders.

As discussed in more detail below, we parametrise the belief set with an interval $[-\overline{\mu}_t, \ \overline{\mu}_t]$ of means centered around zero, so we can think of a loss of confidence as an increase in the width of that interval. That is, a wider interval at history s^t describes an agent who is less confident, perhaps because he has only poor information about what will happen at $t+1$. The preferences above then take the form:

$$U_t(\overrightarrow{C}; s^t) = \min_{\mu \in [-\overline{\mu}_t, \ \overline{\mu}_t]} \mathbb{E}^\mu \left[u(\overrightarrow{C}_t) + \beta U_{t+1}(\overrightarrow{C}; s_t, \ s_{t+1}) \right] \tag{2}$$

The households' budget constraint is:

$$P_t C_t + B_t = R_{t-1} B_{t-1} + W_t N_t + T_t \tag{3}$$

where T_t includes government transfers as well as a profits, W_t is the hourly wage, P_t is the price of the final good and B_t are bonds with a one-period nominal return R_t. There is no heterogeneity across households, because they all earn the same wage in the competitive labor market, they own a diversified portfolio of firms, they consume the same Dixit-Stiglitz consumption bundle and face the same ambiguity. The only peculiarity of

---

[11]More details and axiomatic foundations for such preferences are in Epstein and Schneider (2003).

households in this setup is their uncertainty about the return to their savings R_t. As we describe in more detail below, R_t is formally set by the Central Bank after the consumption decision is made, while the agents make their decisions based on their perceived interest rate $\tilde{R}_t$, which is a function of the ambiguity μ. The Central Bank sets R_t based on current inflation and the current level of the natural rate, so absent ambiguity, the private sector would know its exact value and it would correspond to the usual risk-free rate. In this context, however, agents do not fully trust the Central Bank's response function and so they will consider a range of interest rates indexed by μ.

The household's intertemporal and intratemporal Euler equation are:

$$\frac{1}{C_t} = \mathbb{E}_t^\mu \left[\frac{\beta R_t}{C_{t+1}\Pi_{t+1}} \right] \tag{4}$$

$$N_t^\psi C_t = \frac{W_t}{P_t} \tag{5}$$

While they both look absolutely standard the expectation for the intertemporal Euler equation reflects agents' ambiguous beliefs.

In particular, we assume that ambiguity manifests itself in a potentially distorted value of the policy rate:

$$\mathbb{E}_t^\mu \left[\frac{\beta R_t}{C_{t+1}\Pi_{t+1}} \right] \equiv \mathbb{E}_t \left[\frac{\beta \tilde{R}_t}{C_{t+1}\Pi_{t+1}} \right]$$

Note that this is convenient for expositional purposes but not critical for our results. Since our solution, following Ilut and Schneider (2014), focuses on the worst-case steady state and a linear approximation around it, distorting beliefs about future inflation, future consumption or, indeed, any combination thereof (e.g. the real rate), would be equivalent.

The intertemporal Euler equation thus becomes:

$$\frac{1}{C_t} = \mathbb{E}_t \left[\frac{\beta \tilde{R}_t}{C_{t+1}\Pi_{t+1}} \right] \tag{6}$$

where $\tilde{R}_t \equiv R_t e^{\mu_t}$ and $\mathbb{E}_t$ is the rational-expectations operator.

The government and the central bank. The Government runs a balanced budget and finances the production subsidy with a lump-sum tax. Out of notational convenience, we include the firms' profits and the deadweight loss resulting from price dispersion Δ_t, which is defined in the next section, in the lump-sum transfer:

$$
\begin{aligned}
T_t &= P_t \left(-\tau \frac{W_t}{P_t} N_t + Y_t \left(1 - (1-\tau)\frac{W_t \Delta_t}{P_t A_t} \right) \right) \\
&= P_t Y_t \left(1 - \frac{W_t \Delta_t}{P_t A_t} \right).
\end{aligned}
$$

The first expression explicitly shows that we include in T_t the financing of the subsidy, the second refers to the economy-wide profits, which include the price-dispersion term Δ_t. The Central Bank follows a very simple Taylor rule:

$$R_t = R_t^n \left(\Pi_t \right)^\phi, \tag{7}$$

here R_t is the gross nominal interest rate paid on bonds maturing at time $t+1$ and $R_t^n = \mathbb{E}_t \frac{A_{t+1}}{\beta A_t}$ is the gross natural interest rate[12]. The Central Bank formally sets rates after the private sector makes their economic decisions, but it does so based on variables such as the current natural rate and current inflation, which are known to the private sector as well. At this stage we are trying to characterize an optimal rule so we do not include monetary policy shocks, which would be inefficient in this economy. Therefore if the private sector were to fully trust the Central Bank, *i.e.* $\overline{\mu}_t = 0$:

$$\tilde{R}_t = R_t^n \left(\Pi_t\right)^\phi$$

which is the nominal rate that implements first-best allocations (together with the subsidy). In the context of our analysis, however, ambiguity about the policymaker's response function ($\mu_t \neq 0$) will cause agents to base their decision on the interest-rate level that would hurt their welfare the most if it was to prevail - within the range they entertain:

$$\tilde{R}_t = R_t^n \left(\Pi_t\right)^\phi e^{\mu_t} \tag{8}$$

In this case, even in the presence of the production subsidy, the first-best allocation cannot be achieved, despite the Central Bank following a Taylor rule like that in equation (7) that would normally implement it, because the private sector will use a somewhat different interest rate for their consumption-saving decision. In this stylized setup, we thus capture a situation in which, despite the policymakers actions, the first-best allocation fails to be attained because of a lack of confidence and/or understanding on the part of the private sector, which sets the stage for studying the benefits resulting from making the private sector more aware and confident about the implementation of monetary policy.

Firms. The final good Y_t is produced by final good producers who operate in a perfectly competitive environment using a continuum of intermediate goods $Y_t(i)$ and the standard CES production function

$$Y_t = \left[\int_0^1 Y_t(i)^{\frac{\varepsilon-1}{\varepsilon}} di\right]^{\frac{\varepsilon}{\varepsilon-1}}. \tag{9}$$

Taking prices as given, the final good producers choose intermediate good quantities $Y_t(i)$ to maximize profits, resulting in the usual Dixit-Stiglitz demand function for the intermediate goods

$$Y_t(i) = \left(\frac{P_t(i)}{P_t}\right)^{-\varepsilon} Y_t \tag{10}$$

and in the aggregate price index

$$P_t = \left[\int_0^1 P_t(i)^{1-\varepsilon} di\right]^{\frac{1}{1-\varepsilon}}.$$

---

[12]While there is an expectation in the definition of the natural rate, under rational expectations the expectations of the Central Bank will coincide with those of the private sector, hence the natural rate will be known by both sides and there will be no ambiguity about it.

Intermediate goods are produced by a continuum of monopolistically competitive firms with the following linear technology:

$$Y_t(i) = A_t N_t(i), \tag{11}$$

where A_t is a stationary technology process. Prices are sticky in the sense of Calvo (1983): only a random fraction of firms $(1 - \theta)$ can re-optimise their price at any given period, while the others must keep the nominal price unchanged[13]. Whenever a firm can re-optimise, it sets its price maximising the expected presented discounted value of future profits

$$\max_{P_t^*} E_t \left[\sum_{s=0}^{\infty} \theta^s Q_{t+s} \left(\left(\frac{P_t^*(i)}{P_{t+s}} \right)^{1-\epsilon} Y_{t+s} - \Psi \left(\left(\frac{P_t^*(i)}{P_{t+s}} \right)^{-\epsilon} Y_{t+s} \right) \right) \right] \tag{12}$$

where Q_{t+s} is the stochastic discount factor, Y_{t+s} denotes aggregate output in period $t+s$ and $\Psi(\cdot)$ is the net cost function. Given the simple linear production function in one input the (real) cost function simply takes the form $\Psi(Y_t(i)) = (1-\tau) \frac{W_t}{P_t} \frac{Y_t(i)}{A_t}$, where τ is the production subsidy. The firm's price-setting decision in characterised by the following first-order condition:

$$\frac{P_t^*(i)}{P_t} = \frac{E_t \sum_{j=0}^{\infty} \theta^j Q_{t+s} \left(\frac{P_{t+j}}{P_t} \right)^{\epsilon} \frac{\epsilon}{\epsilon-1} MC_{t+j}}{E_t \sum_{j=0}^{\infty} \theta^j Q_{t+s} \left(\frac{P_{t+j}}{P_t} \right)^{\epsilon-1}},$$

which ultimately pins down inflation, together with the following equation derived from the law of motion for the price index:

$$\frac{P_t^*(i)}{P_t} = \left(\frac{1 - \theta \Pi_t^{\epsilon-1}}{1 - \theta} \right)^{\frac{1}{1-\epsilon}}, \tag{13}$$

and would result in the usual purely forward-looking Phillips Curve if it wasn't for the persistent deviation from its target.

Market clearing. The firm's problem is entirely standard, hence we relegate it to the appendix and turn to the market-clearing conditions. Market clearing in the goods markets requires that

$$Y_t(i) = C_t(i)$$

for all firms $i \in [0,1]$ and all t. Given aggregate output Y_t is defined as in equation (9), then it follows that

$$Y_t = C_t.$$

---

[13]Or indexed to the inflation when we consider it to be non-zero.

Market clearing on the labour market implies that

$$
\begin{aligned}
N_t &= \int_0^1 N_t(i)di. \\
&= \int_0^1 \frac{Y_t(i)}{A_t}di \\
&= \frac{Y_t}{A_t} \int_0^1 \left(\frac{P_t(i)}{P_t}\right)^{-\varepsilon} di
\end{aligned}
$$

where we obtain the second equality substituting in the production function (11) and then use the demand function (10) to obtain the last equality. Let us define $\Delta_t \equiv \int_0^1 \left(\frac{P_t(i)}{P_t}\right)^{-\epsilon} di$ as the variable that measures the relative price dispersion across intermediate firms. Δ_t represents the inefficiency loss due to relative price dispersion under the Calvo pricing scheme: the higher Δ_t, the more labor is needed to produce a given level of aggregate output.

3.2   The Worst-Case Steady State

Following Ilut and Schneider (2014), we study our model economy around the worst-case steady state, because ambiguity-averse agents will make their decisions *as if* that were the steady state. Therefore, we must first identify the worst-case scenario and characterise it. We derive the steady state of the agents' first-order conditions as a function of a generic constant level of μ and we then rank the different steady states (indexed by the level of distortion induced by ambiguity) to characterize the worst-case steady state. Some derivations are reported in the Appendix.

Steady State Inflation and the Policy Rate.   Steady state is characterised by a constant consumption stream. As a result, the intertemporal Euler equation pins down the perceived *real* interest rate, *i.e.* the rate that determines the intertemporal substitution of consumption. Combining this with our simple Taylor rule then delivers the steady state level of inflation consistent with the distortion and the constant consumption stream, as the following result states.

Result 3.1. *In a steady state with no real growth, inflation depends on the ambiguity distortion parameter as follows:*

$$
\Pi(\mu, \cdot) = e^{-\frac{\mu}{\phi - 1}}, \tag{14}
$$

while the policy rate is:

$$
R(\mu, \cdot) = \frac{1}{\beta} e^{-\frac{\phi \mu}{\phi - 1}}. \tag{15}
$$

Hence, $\phi > 1$ implies that for any $\mu > 0$:

$$
\Pi(\mu, \cdot) < \Pi(0, \cdot) = 1 \qquad R(\mu, \cdot) < R(0, \cdot) = \frac{1}{\beta},
$$

and the opposite for $\mu < 0$.

13

Proof. Proof in Appendix B. □

Result 3.1 clearly shows that inflation is a decreasing function of μ as long as $\phi > 1$. The mapping from μ to $\Pi(\mu, \cdot)$ implies that the steady state of the model and its associated welfare, can be equivalently characterised in terms of inflation or in terms of the level of belief distortion μ, since μ does not enter any other steady-state equation, except via the steady-state inflation term. To build some intuition on the steady-state formula for inflation and the interest rate, let us consider the case in which household decisions are based on a level of the interest rate that is systematically lower than the true policy rate $(\mu < 0)$[14]. Other things equal, this will induce a high demand pressure, causing an increase in inflation. In the end, higher inflation will be matched by higher nominal interest rate so that constant consumption in steady state is attained. The result of this is that the policy rate will end up being higher than in the first-best steady state.

Worst-Case Characterization So far we have considered the optimal behaviour of consumers and firms for a given μ - *i.e.* for a given distortion in the agents' beliefs about the expected policy rate. To pin down the worst-case scenario we need to consider how the agents' welfare is affected by different values of the belief distortion μ and find the μ that minimises their welfare.

In our simple model, the presence of the production subsidy ensures that monetary policy implements the first-best allocation. Therefore, any belief distortion $\mu \neq 0$ will generate a welfare loss. However, it is not a priori clear if a negative μ is worse than a positive one of the same magnitude, *i.e.* if underestimation the interest rate is worse than overestimating it by the same amount.

The following result rules out the presence of interior minima for sufficiently small ambiguity ranges, given the weakest restrictions on parameter values implied by economic theory.

Result 3.2. *For $\beta \in [0,1)$, $\epsilon \in (1, \infty)$, $\theta \in [0,1)$, $\phi \in (1, \infty)$, $\psi \in [0, \infty)$, $\mathbb{V}(\mu, \cdot)$ is continuously differentiable around $\mu = 0$ and:*

$$\frac{\partial \mathbb{V}(0, \cdot)}{\partial \mu} = 0 \quad and \quad \frac{\partial^2 \mathbb{V}(0, \cdot)}{\partial \mu^2} < 0$$

As a consequence, for small enough $\overline{\mu}$, there are no minima in $\mu \in (-\overline{\mu}, \overline{\mu})$.

Proof. Proof in Appendix B. □

Result 3.2 illustrates that the welfare function is locally concave around the first-best (see Figure 3, drawn under our baseline calibration described above). Realistic calibrations show that the range of μ for which the value function is concave is in practice much larger than any plausible range for the ambiguity.

Result 3.2 rules out interior minima, but it remains to be seen which of the two extremes is worse from a welfare perspective. Graphically and numerically it is immediate

---

[14]The zero-lower bound poses a restriction on the range of μ of the form: $\mu < -\frac{\phi-1}{\phi} \log(\beta)$.

Figure 3: Steady-state welfare as a function of μ (measured in annualized percentage points).
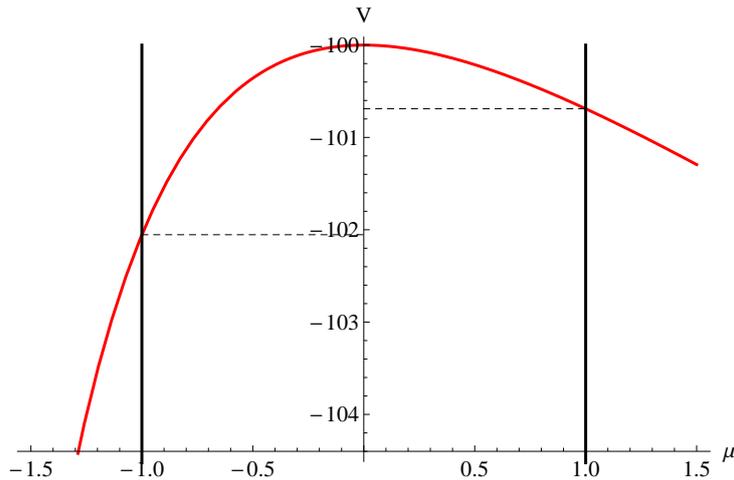
to see from Figure 3 that the welfare function is concave but not symmetric with respect to μ. More generally, it is possible to establish this sufficient condition for the characterisation of the worst-case scenario, under the assumption of a symmetric range of ambiguity around the interest rate.

Result 3.3. *For β sufficiently close but below 1 and all the other parameters in the intervals defined in Result 3.2, $\mu = -\overline{\mu}$ minimizes $\mathbb{V}(\mu, \cdot)$ over $[-\overline{\mu}, \overline{\mu}]$, for any sufficiently small $\overline{\mu} > 0$.*

Proof. Proof in Appendix B. □

In practice, this sufficient condition is not at all restrictive because any sensible calibration would set β to a value that is well within the range for which our result holds. The intuition for the asymmetry of the welfare function is the following. While the effect of μ on inflation is symmetric (in logs) around zero, the impact of inflation on welfare is not. In particular, positive steady-state inflation - associated with negative levels of μ as shown in Result 3.1 - leads to a bigger welfare loss than a corresponding level of negative inflation. This results from the fact that positive inflation tends to lower the relative price of firms who do not get a chance to re-optimise. These firms will face a very high demand, which in turn will push up their labour demand. On the other hand, negative inflation will reduce the demand for firms which do not re-optimise and this will reduce their demand for labour and their marginal costs[15]. In the limit, as the relative price goes to zero, firms will incur huge marginal costs while as their relative price goes to infinity their demand goes to zero.

Having characterised the worst-case steady state, we can now use Result 3.1 to directly infer that, in an ambiguity-ridden economy, inflation will be higher than in the first-best allocation and so will be the policy rate. When agents base their decisions on a perceived interest rate $\tilde{R}_t$ that is lower than the actual policy rate, the under-estimation of the policy rate tends to push up consumption and generate inflationary pressures, which,

---

[15]Once more, this shows that first best is attained in the absence of ambiguity, when the marginal costs equals the inverse of the markup.

in turn, lead to an increase in the policy rate. In particular, so long as $\phi > 1$, the policy rate will increase more than one-for-one with inflation, hence not only the actual but also the perceived rate will be higher than its first-best value, in nominal terms. In sum, because the policy rate responds to the endogenously determined inflation rate, distortions in expectations have a *feedback effect* via their impact on the steady state level of inflation. The combined effects of higher inflation, higher policy rates and negative μ make the perceived *real* rate of interest equal to $\frac{1}{\beta}$, which is necessary to deliver a constant consumption stream, *i.e.* a steady state. The level of this constant consumption stream (and ultimately welfare) depends, in turn, on the price dispersion generated by the level of inflation that characterises the steady state. The following Result summarises Results 3.1 and 3.3 and establishes more formally the effects of ambiguity on the worst-case steady state levels of inflation and the policy rate.

Result 3.4. *For β sufficiently close but below one and all the other parameters in the intervals defined in Result 3.2, for any small enough $\overline{\mu} > 0$:*

$$\mathbb{V}_w(\overline{\mu}') > \mathbb{V}_w(\overline{\mu}) \qquad \Pi_w(\overline{\mu}') < \Pi_w(\overline{\mu}) \qquad R_w(\overline{\mu}') < R_w(\overline{\mu}) \qquad \forall\, 0 \leq \overline{\mu}' < \overline{\mu} \qquad (16)$$

where the w subscript refers to welfare-minimizing steady state value of each variable over the interval $[-\overline{\mu}, \overline{\mu}]$.

Proof. Proof in Appendix B. □

3.3 Optimal Monetary Policy

So far, we have assumed that policymakers follow a rule that *would be* optimal in the absence of ambiguity (equation (7)). We now put ourselves in the shoes of a policymaker that, having followed that rule for a long time, realizes that inflation persistently deviates from its target and output from its potential. Our setting, however, does not readily lend itself to a standard application of Ramsey monetary policy. The reason lies in the fact that if a benevolent planner was to choose the interest rate, they would never select that which minimizes welfare. Not only that, in this class of models the Euler equation is not a binding constraint in the formulation of the Ramsey problem (see Yun, 2005). Another way of seeing this is the following. From a *timeless* perspective the steady state of the Ramsey problem corresponds, in this environment, to one in which there is no inflation, no price dispersion and nor welfare loss of sorts. This is not surprising because we even know a straightforward implementation of this equilibrium, which corresponds to our model when $\phi \to \infty$. We know this is a limit case, one that is not particularly interesting in practice so we will try to characterize optimal policy when ϕ is constrained by some finite value $\overline{\phi}$.

Optimal Policy Rule. Our main optimal monetary policy result characterizes the optimal monetary policy rule when there is a bound on the responsiveness of the policy rate to inflation. In their analysis, Schmitt-Grohé and Uribe (2007) discuss how values

of ϕ above around 3 are unrealistic and, in practice, it is hard to appeal to values much larger than that in light of the ZLB or just common wisdom.

Before turning to the result, it is necessary to point out that we distinguish *static* and *dynamic* optimality, the former referring to steady states and the latter to the stabilization of the economy around it, which we assess based on a welfare function approximation along the lines of that in Coibion, Gorodnichenko and Wieland (2012)[16]

Proposition 3.1. *Given the economy described in Section 3.1, a small $\overline{\mu} > 0$ and restricting $\underline{\phi}(-\overline{\mu}, \cdot) < \phi \leq \overline{\phi}$, the following rule is statically and dynamically optimal in its class:*

$$R_t = R_t^* \Pi_t^{\overline{\phi}} \tag{17}$$

where

$$R_t^* = R_t^n e^{\delta^*(\overline{\mu}, \overline{\phi}, \cdot)} \tag{18}$$

and

$$0 < \delta^*(\overline{\mu}, \overline{\phi}; \cdot) < \overline{\mu} \tag{19}$$

is implicitly defined by $\mathbb{V}\left(-\overline{\mu} + \delta^*(\overline{\mu}, \overline{\phi}; \cdot), \cdot\right) = \mathbb{V}\left(\overline{\mu} + \delta^*(\overline{\mu}, \overline{\phi}; \cdot), \cdot\right)$.

Proof. See Appendix C.6 □

We can summarize the result by saying that the central bank needs to be more *hawkish* because it will respond as strongly as it possibly can to inflation and will increase the Taylor rule's *intercept*. The fact that setting $\phi = \overline{\phi}$ is optimal should not be surprising at this point[17]. A common theme of our analysis is that the higher ϕ the better the outcome in this setting, so the fact that policymakers will set its value as high as they can is natural. The fact that the optimal intercept is higher than the natural rate would warrant is somewhat more novel but quite intuitive too. The central bank would like to tighten more if it had no bound on ϕ, because inflation is still inefficiently high. The only other way it can do so, in this setting, is by increasing the intercept of its Taylor rule. In doing so it runs a risk though. If it increases the intercept too much the worst case will switch. In particular, consider a *naïve* policymaker who realizes the private sector is systematically underestimating its policy rate by $\overline{\mu}$. Its response could amount to systematically setting rates higher than its standard Taylor rule would predict by the same amount $\overline{\mu}$. If these were just parameters and there was no minimization involved, this policy action would implement first best[18]. This would be naïve because, so long as there is some ambiguity lingering the first-best outcome will never be attained. Or, in

---

[16]The log-linear representation of the model can be found in the Appendix. For the purposes of our results the crucial thing to note is that the only shock is a technology shock.

[17]Note that lower bound $\underline{\phi}(-\overline{\mu}, \cdot)$ is simply meant to capture the lowest degree of responsiveness to inflation that ensures determinacy for a given degree of steady-state ambiguity.

[18]That is because, in this economy, steady-state inflation is $\Pi(-\overline{\mu}, \delta(\cdot), \cdot) = e^{-\frac{-\overline{\mu} + \delta(\cdot)}{\phi - 1}}$, so by setting $\delta = \overline{\mu}$, steady state inflation would be zero.

other words, if $\delta = \overline{\mu}$ the worst-case would no longer correspond to one in which the policy rate is under-estimated and positive inflation creeps up in steady state. Rather it would correspond to a situation in which the policy rate is over-estimated and deflation emerges. At this point it is obvious that the central bank can do better than setting $\delta = 0$ because a small positive δ would decrease steady state inflation. At the same time setting $\delta = \overline{\mu}$ is counterproductive. The optimal solution is the one in which the highest δ not causing the worst-case to switch sides is selected (which is the equality implicitly pinning down its value). This, in turn, highlights the fact that the policymaker is not simply facing some kind of a constant wedge, but rather it faces a distortion that can potentially respond to its policy-design efforts.

Another aspect of Proposition 3.1 is that the rule in equation (17) is optimal in *its class*, by which we mean rules including inflation and a measure of the natural rate. One could legitimately argue what would happen if a measure of the output gap or some other variable was included in the specification. Rather than trying to exhaust any possible combination we present the following corollary of Proposition 3.1 which basically states that there cannot be a rule which outperforms the one we proposed provided we are prepared to relax the constraint on ϕ. Or, equivalently, our functional form is only (potentially) restrictive in terms of practical implementability but is otherwise as good as any other could be.

Corollary 3.1. *Given any constrained-optimal monetary policy plan, a monetary policy rule with the same functional form as that in Proposition 3.1 can be made welfare equivalent for a suitably high level of $\overline{\phi}$.*

Proof. See appendix C.7 $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

Having demonstrated how our functional form is actually not really restrictive, one more issue needs to be addressed: what happens if there are tradeoff-inducing shocks on top of the TFP shock. Discussion of dynamic optimality would require a numerical exercise, but we can definitely draw some general conclusions about the efficiency of the steady state. In particular, as trend inflation approaches its first-best value, not only the steady state loss is reduced but also the dynamic tradeoff between inflation and output-gap variability is mitigated, since the coefficient on inflation variability in equation (57) in Appendix C.5 is decreasing in trend inflation while that on output gap variability is constant. So, for a given level of confidence $\overline{\mu}$, we might end up in the paradoxical situation in which increasing ϕ all the way up to $\overline{\phi}$, might result in a reduction of the tradeoff large enough to warrant a smaller ϕ from a dynamic optimality perspective. At the same time, the coefficient on inflation in equation (57) tends to be an order of magnitude larger than that on the output gap for the calibrations one could reasonably try, so it would appear that setting $\phi = \overline{\phi}$ is the optimal thing to do under any circumstance, which is also in line with thourough numerical experiment carried out in Schmitt-Grohé and Uribe (2007).

During the *transition*, our proposed rule would be sub-optimal as Yun (2005) demonstrates. That is because our rule would implement a zero-inflation (in deviation from target) equilibrium which is sub-optimal when starting from $\Delta_{t-1} > 1$, i.e. there is some lingering price dispersion as a legacy from the period before ambiguity was permanently

reduced to zero. As Yun (2005) illustrates, under those circumstances, some deflation (or inflation below target) would be beneficial because it would reduce the price-dispersion inefficiency at a faster pace than a zero-inflation equilibrium. In the limit, however, as price dispersion wanes, our proposed rule would be once more optimal. Which shows how, despite its simplicity, the specification of the policy rule we adopt is extremely robust in this class of economies.

4   Model Validation

We take our model to the data, and show that it can, with a very standard calibration and using the dispersion in the individual SPF nowcasts of the 3-month T-Bill rate as a measure of ambiguity about monetary policy, explain the following stylized facts.

   i. The observed decline in trend inflation over the 80s and 90s.

  ii. The low, and possibly below-target, trend inflation of the last few years.

 iii. The shift from equilibrium indeterminacy to determinacy in the early 1980s

 iv. Paul Volcker's apparent excessive tightening in 1982.

Ambiguity and Trend Inflation.   Our headline measure of ambiguity about the conduct of monetary policy is the interdecile dispersion of nowcasts on the 3-month TBill rate from the Survey of Professional Forecasters, which is available from 1981Q3 onwards. Dispersion is a standard proxy for ambiguity, and we build on the association established above between transparency and forecast dispersion (Swanson, 2006). Finally, an important advantage of the measure we use is that it is measured in basis points, as opposed to being an index, hence it is more readily comparable with measures of inflation.

Figure 2 shows the 4-quarter moving-average of our preferred measure of ambiguity. It is obvious that the degree of dispersion was an order of magnitude larger in the early 1980s than it is now. Another important fact is that, while the degree of dispersion more than halves by the mid 80s, it still clearly trends downward through the mid 90s and then, to a first approximation, flattens out. It is interesting to notice that it flattens somewhere between 25 and 50bps, which means that the usual disagreement among the *hawkish* and *dovish* ends of the professional forecasters pool amounts to situations like the former expecting a 25bp tightening and the latter no change or 25bp loosening, a very reasonable scenario in the late 1990s and early 2000s. In the early 80s, however, that number was some 4-5 times larger. As large as that number can seem from today's vantage point, the swift monetary policy interventions that occurred in the early 80s (see Goodfriend, 2003) can justify these levels.

Our model implies the following relationship between trend inflation and ambiguity, which is simply the log-linear representation of equation (14) when the ambiguity interval is symmetric (as shown in Result 3.1).

$$\overline{\pi}_t \;\; = \;\; \pi_t^* + \frac{\overline{\mu}_t}{\phi - 1} \tag{20}$$

Figure 4: Test for the Null that individual SPF nowcasts are symmetric.
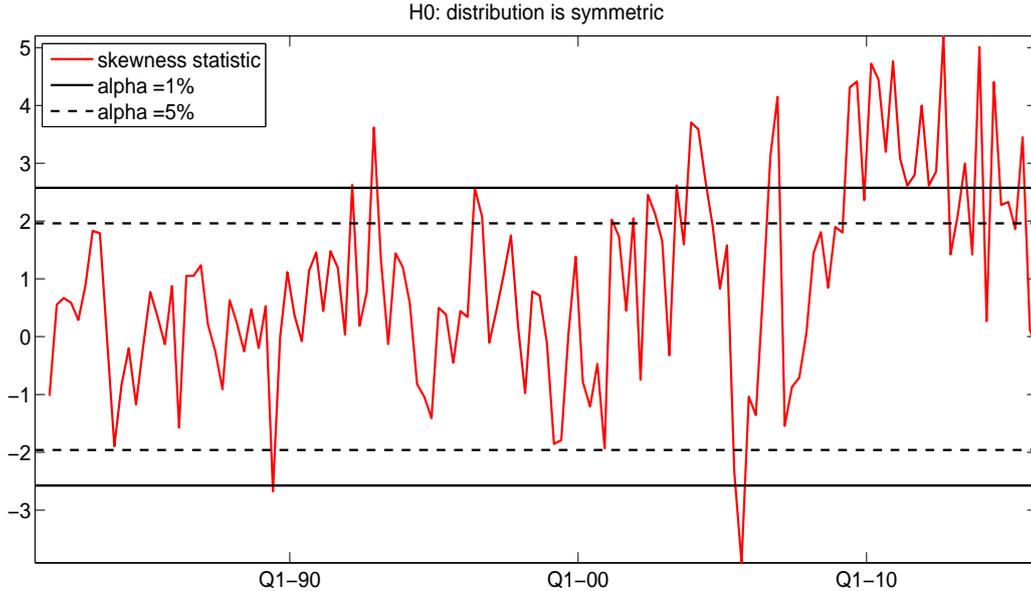
We test for the symmetry of the distribution of the individual SPF nowcasts of the 3-months ahead T-Bill rate. We use D'Agostino (1970) test (in red) and perform it period per period from 1981 to 2014.

This very stark relation implied by our model can be brought to the data in a very simple way. We only need to calibrate ϕ, which we set to 1.5, and the target which we set to 2 percent, as indicated in Table 2. We then use the SPF interdecile dispersion measure described above to calibrate $\overline{\mu}_t$. When the dispersion is roughly symmetric, it makes sense to set $\overline{\mu}_t$ to half the measured dispersion. To verify this assumption, we tested for the symmetry in the dispersion of short-term rate nowcasts using a test developed by D'Agostino (1970). Figure 4 shows that it is remarkably hard to reject symmetry up to the point when rates approached the lower bound. At that point, as one would expect, dispersion started to display a noticeable upward skew, as the zero lower bound provides a natural limit to the extent that agents can forecast low rates.

We can draw a first implication from our model, simply based on this observed change in the features of the dispersion series. That is because, under symmetry, our model predicts trend inflation to exceed target, while in the presence of sufficiently high positive skew the worst-case scenario switches to being one in which inflation is below target. This is consistent with the idea that, in the past, high inflation was the policymakers' and the public's main concern, while in the last few years the main concern has been low inflation instead. And it is also consistent with the decline in the inflation trend. A more

Table 2: Parameter Values		
β	Subjective Discount	.995
ψ	Inverse Frish Elasticity	1
ϕ	Inflation Responsiveness	1.5
ϵ	Demand elasticity	15
θ	Calvo probability	.83

20

Figure 5: Trend inflation implied by our measure of forecasters' disagreement
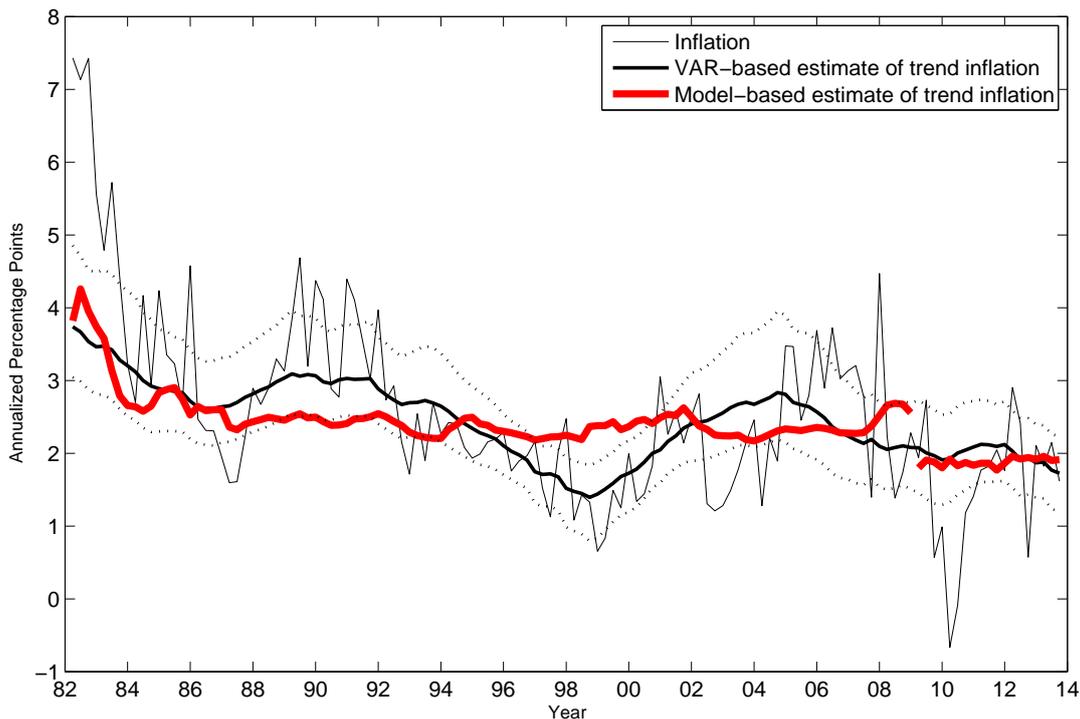
Level of annualized trend inflation implied by our measure of forecasters' disagreement in red, assuming $\phi = 1.5$ and $\pi^* = 2pc$. The thick black line is the measure of trend inflation from the time-varying-parameters VAR, will the dotted lines are is 68% confidence bands.

accurate assessment can be obtained by plugging our measure of ambiguity in equation. (20) and comparing the resulting measure of trend inflation with our estimates where, as the distribution becomes asymmetric, the sign will switch.

Figure 5 overlays the measure of trend inflation implied by our model to the baseline estimate for the trend (based on Cogley and Sbordone, 2008). The ambiguity-based estimate of trend inflation captures the decline in the secular component of inflation in the data. We find this very encouraging, considering this only depends on the calibration of two parameters, which we set to standard values.

Ambiguity and Equilibrium Determinacy. It is well known that the Taylor principle, which postulates that nominal rates should move more than one for one with inflation, ensures equilibrium determinacy in New-Keynesian DSGE models log-linearized around the zero-inflation steady-state (Gali' 2008). The well known *policy mistakes view* (exemplified by the work of Clarida, Galí and Gertler, 2000), attributes the high and out-of-control inflation of the 1970s to the policy response to inflation failing to meet the Taylor principle. Ascari and Ropele (2009), however, show that, when the model is approximated around positive trend inflation, the Taylor principle is not a guarantee for equilibrium determinacy. Coibion and Gorodnichenko (2011) further explore this aspect and find that the Fed was, in all likelihood, already satisfying the Taylor principle prior to the appointment of Chairman Volcker. They attribute the switch from indeterminate to determinate equilibria to a fall in trend inflation. We push this argument one step

Figure 6: Indeterminacy region (in gray) as a function of the degree of ambiguity expressed in annualized basis points (on the horizontal axis) and the responsiveness of the policy rate to inflation (on the vertical axis). The solid black line corresponds to $\phi = 1.5$ while the black dashed line represents the $(\overline{\mu}, \phi)$ pairs consistent with annualized inflation 1.5 percent above target.
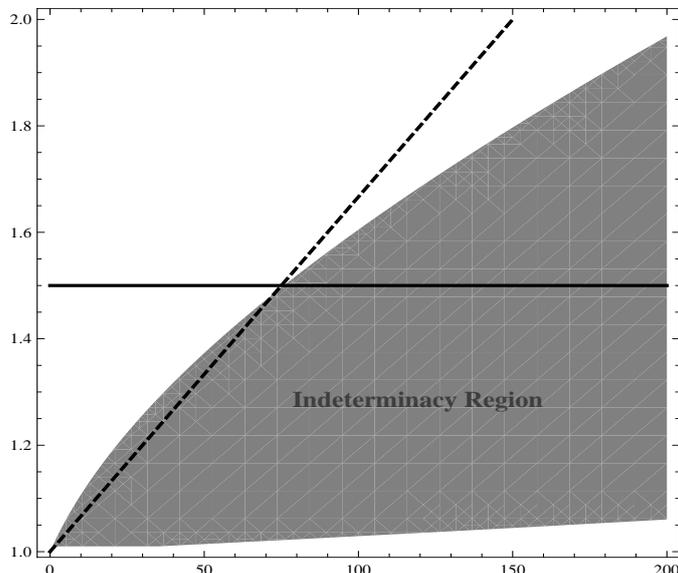
further. That is because trend inflation, in our model, depends, in turn, on ϕ, the inflation responsiveness coefficient. In our setup, higher ϕ's – typically associated with satisfying the Taylor principle – also cause trend inflation to fall. Mechanically, a higher ϕ will have the usual, direct, effect on the eigenvalues of the dynamic system, associated with the idea that if $\phi > 1$ real rates move in the same direction as nominal rates, thus stabilizing the economy via their effect on consumption. Moreover, ϕ will affect the so called *transmission mechanism* because, as shown in Appendix C the parameters of the log-linearized model depend the level of the inflation trend, and therefore on ϕ. This supplementary channel generates a non-trivial relationship between the value of ϕ and equilibrium determinacy and provides another way to validate our model.

Figure 6 illustrates the determinacy region (white area) as a function of the degree of ambiguity and ϕ. The black solid line in Figure 6 marks the $\phi = 1.5$ level. Our reading of the early 1980s is one in which ϕ is constant but $\overline{\mu}$ falls, *i.e.* we are moving leftward on the chart. We calibrate the remaining parameters of our model so that the equilibrium switches to being determinate when the degree of ambiguity falls below 150bps, a value that prevailed in the early 80s. Interestingly, our model can deliver this scenario under a very reasonable calibration. The two key parameters in this respect are demand elasticity ϵ, set at a value which implies a markup of about 7 percent, and the probability of not having a chance to re-optimize prices under the Calvo-pricing scheme θ. θ is set equal to .83, implying an average price duration of 6 quarters. So, the prediction of our model, where we endogenize the trend but we maintain a constant response to inflation, is consistent with the Coibion and Gorodnichenko (2011) description of the transition to a determinate equilibrium.

Optimal Policy Prediction. One of the main challenges to our maintained assumption that the target never changed over the last 35 or so years is the overly tight policy stance over the course of 1982, Goodfriend (2005, p. 248). In our model a similar course of action is perfectly sensible, for a given inflation target, if one accounts for our optimal policy prediction. The immediate consequence of Result 3.1 is, in fact, that in the presence of ambiguity *extra tightening* is called for.

So, while a full quantitative assessment of this effect is left for future research, Chairman Volckers policy is in line with our simple optimal policy prediction. If anything, the limit of our model, in its current form, is that it does not have an explicit mechanism linking that course of action to the increase in confidence regarding the conduct of monetary policy it ensued. Rather, that effect is captured via the subsequent fall in forecasters' forecast dispersion.

Low Inflation since 2009-2010 Over recent years, inflation trend fell below the 2% mark. This is hardly surprising given the proximity of interest rates to their lower bound. Indeed, it is natural to expect that in a similar situation, the worst-case scenario is one in which rates are too high, resulting in a low level of inflation.

This is true in our model as well. We have shown above that welfare falls both when rates are expected to be too high or too low. Our main results depend on the symmetry of the interval over which agents are ambiguous. The lower bound on interest rates, however, poses an inherent constraint to the possibility of excessive rate cuts, thus making it natural to think that forecast dispersion will display a degree of upward skewness. When that is sufficiently high, it results in the worst case switching to one in which inflation is below the target. Figure 4 illustrates the results of statistical testing of the hypothesis that the dispersion of forecasts is symmetric. The null hypothesis that the dispersion is symmetric is hard to reject before the Great Recession, while is starkly rejected from 2009 onwards, when the series displays a significant upward skew. Hence the fall in trend inflation below 2 percent is consistent with the prediction of our model in the aftermath of the Great Recession.

5 Conclusions

We develop a model that features ambiguity-averse agents and ambiguity regarding the conduct of monetary policy, but is otherwise standard. We show that the presence of ambiguity has far-reaching effects, also in steady state. In particular, the model can generate trend inflation endogenously. Trend inflation has three determinants in our model: the inflation target, the strength with which the central bank responds to deviation from the target and the degree of private sector confidence about the monetary policy rule.

Based on a calibration of ambiguity that matches the interdecile dispersion of the SPF nowcasts of the current quarter's 3-monthTBill rates, our model can explain the disinflation of the 80s and 90s as resulting from an increase in the private sector's confidence in their understanding of monetary policy, rather than from changes in target inflation. We can also match the fall in trend inflation since 2009-2010 as an effect of the zero lower

bound. The bound sets a floor to the distortion of the interest rate and makes the worst-case scenario shift from one in which agents fear that the interest rate will be set too low, to one in which they fear it will be set to high. Agents base their consumptions-saving decisions on a distorted rate that is too high and will therefore end up facing low inflation. We also confirm the finding in Coibion and Gorodnichenko (2011) that the equilibrium in the pre-Volcker period might have been indeterminate even though the Taylor principle was satisfied throughout, because of the presence of trend inflation. However in our model the trend inflation itself depends on the inflation responsiveness coefficient in the central bank's response function. In other words, by increasing the degree to which it responds to inflation, a central bank will not only affect the dynamics but also the steady state level of trend inflation.

Finally, given the importance of monetary policy for the determination of trend inflation, we complete the paper studying optimal monetary policy. We can prove analytically that, irrespective of the specifics of the parametrization, the higher the degree of ambiguity, the *more hawkish* a central banker needs to be in order to achieve a comparable degree of welfare. Also, the higher the degree of uncertainty, the higher needs to be the weight on inflation variability in the policymaker's welfare-based loss function. Our results also imply that if a policymaker wanted to be *less hawkish*, he or she should ensure a lower level of ambiguity about monetary policy in order to achieve a comparable degree of welfare.

References

[1] Ascari, G. and T. Ropele, 2007. "Optimal monetary policy under low trend inflation," Journal of Monetary Economics, Elsevier, vol. 54(8), 2568-2583.

[2] Ascari, G. and T. Ropele, 2009. "Trend Inflation, Taylor Principle, and Indeterminacy," Journal of Money, Credit and Banking, vol. 41(8), 1557-1584.

[3] Ascari, G. and A. Sbordone, 2014. "The Macroeconomics of Trend Inflation," Journal of Economic Literature, vol. 52(3), 679-739.

[4] Baker, S. R., N. Bloom, and S. J. Davis, 2015. "Measuring economic policy uncertainty," http://www.policyuncertainty.com/media/BakerBloomDavis.pdf .

[5] Benigno, P. and L. Paciello, 2014. "Monetary policy, doubts and asset prices," Journal of Monetary Economics, 64(2014), 85-98.

[6] Boyarchenko, N., V. Haddad, and M.C. Plosser, 2016. "The Federal Reserve and Market Confidence." Federal Reserve Bank of New York Staff Reports, no. 773

[7] Calvo, G. A., 1983. "Staggered prices in a utility-maximizing framework," Journal of Monetary Economics, Elsevier, vol. 12(3), pages 383-398, September.

[8] Clarida, R., Gali, J. and Mark Gertler, 2000. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," The Quaterly Journal of Economics, Vol. 115, No. 1, pp.147-180.

[9] Clark, T.E. and C. Garciga, 2016. "Trecent Inflation Trends," Economic Trends, Federal Reserve Bank of Cleveland.

[10] Clark, T.E. and T. Nakata, 2008. "Has the behavior of inflation and long-term inflation expectations changed?," Economic Review, Federal Reserve Bank of Kansas City, issue Q I, pages 17-50.

[11] Cogley, T. and T.J. Sargent, 2002. "Evolving Post World War II U.S. Inflation Dynamics." In *NBER Macroeconomics annual 2001*, edited by B.S. Bernanke and K. Rogoff, 331-338, MIT press.

[12] Cogley, T. and A.Sbordone, 2008. "Trend Inflation, Indexation, and Inflation Persistence in the New Keynesian Phillips Curve." *American Economic Review*, 98(5):2101-26.

[13] Coibion, O. and Y. Gorodnichenko, 2011. "Monetary Policy, Trend Inflation, and the Great Moderation: An Alternative Interpretation," American Economic Review, vol. 101(1), pp. 341-70.

[14] Coibion, O., Y. Gorodnichenko and J. Wieland, 2012. "The Optimal Inflation Rate in New Keynesian Models: Should Central Banks Raise Their Inflation Targets in Light of the Zero Lower Bound?," Review of Economic Studies, Oxford University Press, vol. 79(4), pages 1371-1406

[15] D'Agostino, R.B., 1970. "Transformation to Normality of the Null Distribution of *g1*," Biometrika, 57(3), 679-681

[16] Del Negro, M., M.P. Giannoni and F. Schorfheide, 2015. "Inflation in the Great Recession and New Keynesian Models," American Economic Journal: Macroeconomics, American Economic Association, vol. 7(1), pages 168-96, January

[17] Del Negro, M. and S. Eusepi, 2011. "Fitting observed inflation expectations," Journal of Economic Dynamics and Control, Elsevier, vol. 35(12), pages 2105-2131.

[18] Drechsler, I., 2013. "Uncertainty, Time-Varying Fear and Asset Prices," Journal of Finance, vol. 68(5), pp. 1843-1889

[19] Ehrmann, M., S. Eijffinger and M. Fratzscher, 2012. "The Role of Central Bank Transparency for Guiding private Sector Forecasts," Scandinavian Journal of Economics, 114(3),1018-1052.

[20] Epstein, L.G. and M. Schneider, 2003. "Recursive multiple-priors," Journal of Economic Theory, Elsevier, vol. 113(1), pages 1-31.

[21] Galí, J. 2008. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*, Princeton University Press.

[22] Gilboa, I. and D. Schmeidler, 1989. "Maxmin expected utility with non-unique prior," Journal of Mathematical Economics, Elsevier, vol. 18(2), pp. 141-153

[23] Hansen, L.P. and T. J. Sargent, 2007. *Robustness*. Princeton University Press.

[24] Ilut, C. and M. Schneider, 2014. "Ambiguous Business Cycles," American Economic Review, American Economic Association, Vol. 104(8), pp. 2368-99.

[25] Kreps, D.M., 1998. "Anticipated Utility and Dynamics Choice." In *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures 1983-1997*, edited by D.P. Jacobs, E. Kalai and M.I. Kamien, 242-74. Cambridge University Press

[26] Lindsey, D.E., 2003. *A Modern History of FOMC Communication: 1975-2002*, Board of Governors of the Federal Reserve System.

[27] Schmitt-Grohe, S. and M. Uribe, 2007. "Optimal simple and implementable monetary and fiscal rules," Journal of Monetary Economics, Elsevier, vol. 54(6), pages 1702-1725.

[28] Swanson, E.T., 2006. "Have Increases in Federal Reserve Transparency Improved Private Sector Interest Rate Forecasts?," Journal of Money Credit and Banking, vol. 38(3), pp. 791-819

[29] Taylor, J.B., 1993. "Discretion versus Policy Rules in Practice," Carnegie-Rochester Conference Series on Public Policy 39: 195214.

[30] Woodford, M., 2003. Interest and Prices: Foundations of a Theory of Monetary Policy. Princeton University Press.

[31] Yun, T. 2005. "Optimal Monetary Policy with Relative Price Distortions,"American Economic Review, American Economic Association, vol. 95(1), pages 89-109.

A   Some Steady state results

A.0.1   Pricing

In our model firms index their prices based on the first-best inflation, which corresponds to the inflation target and is zero in this case. Because of ambiguity, however, steady-state inflation will not be zero and therefore there will be price dispersion in steady state:

$$\Delta(\mu,\cdot) \;=\; \frac{(1-\theta)\left(\frac{1-\theta\Pi(\mu,\cdot)^{\epsilon-1}}{1-\theta}\right)^{\frac{\epsilon}{\epsilon-1}}}{1-\theta\Pi(\mu,\cdot)^{\epsilon}} \tag{21}$$

Δ is minimised for $\Pi = 1$ - or, equivalently, $\mu = 0$ - and is larger than unity for any other value of μ. As in Yun (2005), the presence of price dispersion reduces labour productivity and ultimately welfare.

A.0.2   Hours, Consumption and Welfare

In a steady state with no real growth, steady-state hours are the following function of μ:

$$N(\mu,\cdot) = \left(\frac{\left(1-\theta\Pi(\mu,\cdot)^{\epsilon-1}\right)\left(1-\beta\theta\Pi(\mu,\cdot)^{\epsilon}\right)}{\left(1-\beta\theta\Pi(\mu,\cdot)^{\epsilon-1}\right)\left(1-\theta\Pi(\mu,\cdot)^{\epsilon}\right)}\right)^{\frac{1}{1+\psi}}, \tag{22}$$

while consumption is:

$$C(\mu,\cdot) = \frac{A}{\Delta(\mu,\cdot)}N(\mu,\cdot) \tag{23}$$

Hence the steady state welfare function takes a very simple form:

$$\mathbb{V}(\mu,\cdot) \;=\; \frac{1}{1-\beta}\left(\log\left(C(\mu,\cdot)\right) - \frac{N(\mu,\cdot)^{1+\psi}}{1+\psi}\right). \tag{24}$$

Finally note that equation (22) delivers the upper bound on steady-state inflation that is commonly found in this class of models (e.g. Ascari and Sbordone (2014)). As inflation grows, the denominator goes to zero faster than the numerator, so it has to be that $\Pi(\mu,\cdot) < \theta^{-\frac{1}{\epsilon}}$ for steady state hours to be finite[19]. Given our formula for steady-state inflation, we can then derive the following restriction on the range of values μ can take on, given our parameters:

$$\mu > \frac{\phi-1}{\epsilon}\log\left(\theta\right), \tag{25}$$

where the right-hand side is negative since $\epsilon > 1$, $\phi > 1$ and $0 < \theta < 1$. To put things in perspective, note that the calibration in Table 2 delivers a bound of the order of -3.1 percent, which is about twice as large as the largest value suggested by our measure of expectations disagreement. So our calibration is unaffected by this bound.

---

[19]Indeed, the same condition could be derived from the formula for price dispersion in equation (21).

Figure 7: Welfare function for $\phi = 1.5$ (solid red line) and $\phi = 1.4$ (orange dashed line).
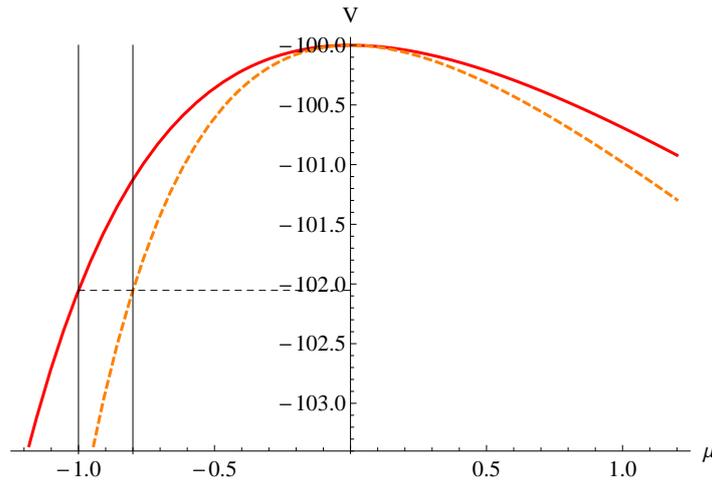
A.1 The role of the inflation response coefficient ϕ

In a model in which the only shock is a technology shock, without ambiguity any inflation response coefficient ϕ larger than one would deliver the first-best allocation (see Galí, 2008), both in steady state and even period by period. As a result, from a welfare perspective any value of $\phi > 1$ would be equivalent. Once ambiguity enters the picture, however, things change and the responsiveness of the Central Bank to inflation interacts with ambiguity in an economically interesting way. In particular, it is possible to view a reduction in ambiguity and the responsiveness to inflation as substitutes in terms of welfare, which can be formalized as follows.

Result A.1. *While parameter values are in the intervals defined in Result 3.2 and $\overline{\mu}$ is a small positive number, given any pair $(\mu, \phi) \in [-\overline{\mu}, 0) \times (1, \infty)$, for any $\mu' \in [-\overline{\mu},\ 0)$ there exists $\phi' \in (1,\ \infty)$ such that:*

$$\mathbb{V}(\mu, \phi') = \mathbb{V}(\mu', \phi)$$

And $\phi' \geq \phi$ iff $\mu' \geq \mu$.
A corresponding equivalence holds for $\mu \in (0,\ \overline{\mu}]$.

Proof. Proof in Appendix B. $\qquad\square$

The intuition behind this relationship between responsiveness to inflation and ambiguity is the following. What ultimately matters for welfare is the steady-state level of inflation: if ambiguity is taken as given, the only way of getting close to first-best inflation is for the Central Bank to respond much more strongly to deviations of inflation from its first-best level. A higher value of ϕ works as an insurance that the response to inflation will be aggressive, which acts against the effect of ambiguity about policy. At the same time, as Schmitt-Grohé and Uribe (2007) suggest, it is practically not very sensible to consider very high values for ϕ, for instance because of the possibility that a modest cost-push shock would cause the policy rate to hit the Zero Lower Bound. So a very high value for ϕ is ultimately not a solution. Figure 7 illustrates Result A.1 graphically for

28

our preferred calibration. Our baseline scenario is presented in the red solid line. If ϕ was lower, say equal to 1.4, the welfare function would become steeper[20] (as the orange dashed line illustrates) because, for a given degree of ambiguity, inflation would be farther away from first best. An application of our result works as follows in this case. Consider a degree of ambiguity of 100bp and $\phi = 1.5$. Figure A.1 shows graphically that the same level of welfare can be attained when $\phi = 1.4$, but only if ambiguity is smaller (of the order of 80bp).

B   Proofs of Steady State Results

Proof of Result 3.1

In steady state, equation 6 becomes:

$$1 = \frac{\beta \tilde{R}(\mu, \cdot)}{\Pi(\mu, \cdot)} \tag{26}$$

From the Taylor rule we get:

$$\tilde{R}(\mu, \cdot) = R^n(\mu, \cdot) \Pi(\mu, \cdot)^\phi e^\mu = \frac{1}{\beta} \Pi(\mu, \cdot)^\phi e^\mu \tag{27}$$

Combining the two, delivers the first part of the result.

The second follows immediately by plugging the resulting expression for inflation into the Taylor rule.

The inequalities result by noting that $\phi > 1$. $\qquad\square$

Proof of Result 3.2

$\mathbb{V}(\mu, \cdot)$, as defined in equation 24, is continuously differentiable around zero. Direct computation, or noting that the first-best allocation is attained in our model when $\mu = 0$, shows that $\frac{\partial \mathbb{V}(\mu, \cdot)}{\partial \mu} = 0$.

Direct computation also delivers:

$$\left. \frac{\partial^2 \mathbb{V}(\mu, \cdot)}{\partial \mu^2} \right|_{\mu=0} = -\frac{\theta\left((\beta-1)^2\theta + \epsilon(\beta\theta-1)^2(1+\psi)\right)}{(1-\beta)(\theta-1)^2(\beta\theta-1)^2(\phi-1)^2(1+\psi)} \tag{28}$$

All the terms are positive given the minimal theoretical restrictions we impose, hence the second derivative is strictly negative and there are no interior minima in a neighbourhood of zero. $\qquad\square$

Proof of Result 3.3

Direct computation shows that the third derivative evaluated at $\mu = 0$ can be expressed

---

[20]Note that, the welfare functions attain the same maximum at $\mu = 0$, which illustrates the fact that the exact value of ϕ is irrelevant in the absence of ambiguity.

as:

$$\left.\frac{\partial^3 \mathbb{V}(\mu,\cdot)}{\partial \mu^3}\right|_{\mu=0} = \frac{\epsilon(2\epsilon-1)\theta(1+\theta)}{(1-\beta)(1-\theta)^3(\phi-1)^3} + \mathcal{R}(\beta) \tag{29}$$

Where, given our parameter restrictions, the first term on the RHS is positive and $\mathcal{R}(\beta)$ is a term in β such that $\lim_{\beta \to 1^-} \mathcal{R}(\beta) = 0$.

Hence, $\lim_{\beta \to 1^-} \left.\frac{\partial^3 \mathbb{V}(\mu,\cdot)}{\partial \mu^3}\right|_{\mu=0} = +\infty$.

Moreover, $\partial \left(\left.\frac{\partial^3 \mathbb{V}(\mu,\cdot)}{\partial \mu^3}\right|_{\mu=0} \right) / \partial\beta$ exists, which ensures continuity of the third derivative in β. Hence the third derivative is positive for any β sufficiently close to but below unity. A third-order Taylor expansion around zero can be used to show that:

$$\mathbb{V}(\mu_0,\cdot) - \mathbb{V}(-\mu_0,\cdot) = \left.\frac{\partial^3 \mathbb{V}(\mu,\cdot)}{\partial \mu^3}\right|_{\mu=0} \frac{2\mu_0^3}{6} + o(\mu_0^4), \tag{30}$$

which is positive for a generic, positive but small value μ_0 thus showing that, the steady state value function attains a lower value at $-\mu_0$ than it does at at μ_0. This, combined with the absence of internal minima (Result 3.2), delivers our result. $\qquad\square$

Proof of Result 3.4

The first inequality follows immediately, as a weak inequality, by considering that $\mathbb{V}_w(\overline{\mu}')$ is the minimum value of welfare on a smaller set than $\mathbb{V}_w(\overline{\mu})$.

The strict inequality follows from the characterization of the worst case in Results 3.2 and 3.3; in particular from the fact that $\mathbb{V}_w(\overline{\mu}) = \mathbb{V}(-\overline{\mu},\cdot)$ and that $\left.\frac{\partial \mathbb{V}(\mu,\cdot)}{\partial \mu}\right|_{\mu<0} > 0$ in the vicinity of $\mu = 0$.

For what concerns inflation, given the formula in Result 3.1, $\phi > 1$ and given that the worst case corresponds to $\mu = -\overline{\mu}$, it is immediate to verify that $\frac{\overline{\mu}}{\phi-1} > \frac{\overline{\mu}'}{\phi-1}$.

$\phi > 1$ also ensures that the Taylor rule is increasing in inflation more than one for one, which delivers the last inequality. $\qquad\square$

Proof of Result A.1

Inspection reveals that μ and ϕ only enter steady-state welfare through the steady-state inflation term $\Pi(\mu,\cdot) = e^{\frac{\mu}{1-\phi}}$. It follows immediately that, for a given μ', $\phi' = 1 + \frac{(\phi-1)\mu}{\mu'}$ implies that (μ,ϕ') is welfare equivalent to (μ',ϕ). $(\mu,\mu') \in [-\overline{\mu},0) \times [-\overline{\mu},0)$ ensures that $\mu' \cdot \mu > 0$ and so $\phi' \in (1,\infty)$ for any $\phi > 1$. The inequalities follow immediately from the definition of ϕ' given above and the fact that both μ and μ' have the same sign. A similar argument would go through for $(\mu,\mu') \in (0,\overline{\mu}] \times (0,\overline{\mu}]$. $\qquad\square$

C   Model Dynamics

To study the dynamic properties of our model, we log-linearize the equilibrium conditions around the worst-case steady state in the usual way. As explained in Ascari and Ropele

(2007), having price dispersion in steady state essentially results in an additional term in the Phillips Curve. Appendix C.1 presents the log-linear approximation around a generic steady state indexed by μ. By setting $\mu = -\overline{\mu}$, we obtain the log-linear approximation to the worst-case steady state. An important caveat is that, in so doing, we are maintaining that the worst case scenario corresponds to $-\overline{\mu}$ in all states of the economy, but will provide sufficient conditions for this to be indeed the case in the next paragraph. Once we have verified our conjecture about the worst-case steady state, we turn our attention to the implications of changes in the agents' confidence in their understanding of the monetary policy rule on the determinacy region and we then study the effects of shocks to ambiguity.

C.1 Log-linearized Equations and Solution

The following equations describe the dynamics of the variables of interest around a generic steady state indexed by μ. Setting $\mu = -\overline{\mu}$ one obtains the log-linear approximation around the worst-case steady state:

$$
\begin{aligned}
c_t &= \mathbb{E}_t c_{t+1} - (\tilde{r}_t - \mathbb{E}_t \pi_{t+1}) \\
w_t &= c_t + \psi n_t \\
\pi_t &= \kappa_0(\mu, \cdot) mc_t + \kappa_1(\mu, \cdot) \mathbb{E}_t \widehat{F2}_{t+1} + \kappa_2(\mu, \cdot) \mathbb{E}_t \pi_{t+1} \\
r_t &= r_t^n + \phi \pi_t \\
\tilde{r}_t &= r_t + \mu_t \\
mc_t &= w_t - a_t \\
y_t &= a_t - \tilde{\Delta}_t + n_t \\
c_t &= y_t \\
r_t^n &= a_{t+1} - a_t \\
y_t^n &= a_t \\
\hat{\Delta}_t &= \Pi(\mu, \cdot)^\epsilon \theta \hat{\Delta}_{t-1} + \epsilon \left(\Pi(\mu, \cdot)^\epsilon \theta - (1 - \Pi(\mu, \cdot)^\epsilon \theta) \frac{\theta}{\left(\frac{1}{\Pi(\mu, \cdot)} \right)^{\epsilon-1} - \theta} \right) \pi_t \\
\widehat{F2}_t &= (\epsilon - 1)\beta \theta \Pi(\mu, \cdot)^{\epsilon-1} \mathbb{E}_t \pi_{t+1} + \beta \theta \Pi(\mu, \cdot)^{\epsilon-1} \mathbb{E}_t \widehat{F2}_{t+1}
\end{aligned}
$$

Where we define:

$$\kappa_0\left(\mu,\cdot\right) \equiv \frac{\left(\left(\frac{1}{\Pi(\mu,\cdot)}\right)^{\epsilon-1}-\theta\right)\left(1-\beta\theta\Pi\left(\mu,\cdot\right)^{\epsilon}\right)}{\theta} \tag{31}$$

$$\kappa_1\left(\mu,\cdot\right) \equiv \beta\left(\left(\frac{1}{\Pi\left(\mu,\cdot\right)}\right)^{\epsilon-1}-\theta\right)\left(\Pi\left(\mu,\cdot\right)-1\right)\Pi\left(\mu,\cdot\right)^{\epsilon-1} \tag{32}$$

$$\kappa_2\left(\mu,\cdot\right) \equiv \beta\Pi\left(\mu,\cdot\right)^{\epsilon-1}\left(\theta(\epsilon-1)(\Pi\left(\mu,\cdot\right)-1)+(1-\epsilon+\epsilon\Pi\left(\mu,\cdot\right))\left(\frac{1}{\Pi\left(\mu,\cdot\right)}\right)^{\epsilon-1}\right) \tag{33}$$

$$\kappa_3\left(\mu,\cdot\right) \equiv \Pi(\mu,\cdot)^{\epsilon}\theta \tag{34}$$

$$\kappa_4\left(\mu,\cdot\right) \equiv \epsilon\left(\Pi(\mu,\cdot)^{\epsilon}\theta-(1-\Pi(\mu,\cdot)^{\epsilon}\theta)\frac{\theta}{\left(\frac{1}{\Pi(\mu,\cdot)}\right)^{\epsilon-1}-\theta}\right) \tag{35}$$

$$\kappa_5\left(\mu,\cdot\right) \equiv (\epsilon-1)\beta\theta\Pi(\mu,\cdot)^{\epsilon-1} \tag{36}$$

$$\kappa_6\left(\mu,\cdot\right) \equiv \beta\theta\Pi(\mu,\cdot)^{\epsilon-1} \tag{37}$$

The equations above (amended for the fact that we allow μ_t to vary) can be summarized in the following system of four equations:

$$\tilde{y}_t = \mathbb{E}_t\tilde{y}_{t+1}-(\phi\pi_t+\hat{\mu}_t-\mathbb{E}_t\pi_{t+1})$$

$$\pi_t = \kappa_0\left(-\overline{\mu},\cdot\right)\left((1+\psi)\tilde{y}_t+\psi\hat{\Delta}_t\right)+\kappa_1\left(-\overline{\mu},\cdot\right)\mathbb{E}_t\widehat{F2}_{t+1}+\kappa_2\left(-\overline{\mu},\cdot\right)\mathbb{E}_t\pi_{t+1}$$

$$\hat{\Delta}_t = \kappa_3\left(-\overline{\mu},\cdot\right)\hat{\Delta}_{t-1}+\kappa_4\left(-\overline{\mu},\cdot\right)\pi_t$$

$$\widehat{F2}_t = \mathbb{E}_t\left(\kappa_5\left(-\overline{\mu},\cdot\right)\pi_{t+1}+\kappa_6\left(-\overline{\mu},\cdot\right)\widehat{F2}_{t+1}\right)$$

Where $\tilde{y}_t \equiv y_t-y_t^n = y_t-a_t$ is the output gap and κ's are defined above. Since they depend on steady state inflation, it is important to note that they are evaluated at the worst-case steady state $\mu=-\overline{\mu}$.

It is then possible to verify that the following guesses solve the system above[21]:

$$\pi_t = \lambda_{\pi\Delta}\Delta_{t-1}+\lambda_{\pi\mu}\hat{\mu}_t$$

$$\tilde{y}_t = \lambda_{y\Delta}\Delta_{t-1}+\lambda_{y\mu}\hat{\mu}_t$$

$$\widehat{F2}_t = \lambda_{F\Delta}\Delta_{t-1}+\lambda_{F\mu}\hat{\mu}_t$$

When hours enter the felicity function linearly, however, the solution simplifies further as $\lambda_{\pi\Delta} = \lambda_{y\Delta} = \lambda_{F\Delta} = 0$. As a result, simple analytic expressions for the other undetermined coefficients can be computed, which are reported in the main body of the text.

---

[21]The technology process a_t does not enter the solution of this system because r_t^n is included in the Taylor rule. However it is still part of the state of the economy because it has a role in determining welfare.

C.2   Log-Linear solution

Appendix C.1 reports the log-linear equations that govern the evolution of our economy around the worst-case steady state. They can be summarized into four equations:

$$\tilde{y}_t = \mathbb{E}_t \tilde{y}_{t+1} - (\phi \pi_t + \hat{\mu}_t - \mathbb{E}_t \pi_{t+1}) \tag{38}$$

$$\pi_t = \kappa_0\left(-\overline{\mu},\cdot\right)\left((1+\psi)\tilde{y}_t + \psi\hat{\Delta}_t\right) + \kappa_1\left(-\overline{\mu},\cdot\right)\mathbb{E}_t\widehat{F2}_{t+1} + \kappa_2\left(-\overline{\mu},\cdot\right)\mathbb{E}_t\pi_{t+1} \tag{39}$$

$$\hat{\Delta}_t = \kappa_3\left(-\overline{\mu},\cdot\right)\hat{\Delta}_{t-1} + \kappa_4\left(-\overline{\mu},\cdot\right)\pi_t \tag{40}$$

$$\widehat{F2}_t = \mathbb{E}_t\left(\kappa_5\left(-\overline{\mu},\cdot\right)\pi_{t+1} + \kappa_6\left(-\overline{\mu},\cdot\right)\widehat{F2}_{t+1}\right). \tag{41}$$

where $\tilde{y}_t = c_t - a_t$ is the deviation of the output gap from its worst-case steady-state level, Δ_t is the price dispersion, while $\widehat{F2}_t$ can be interpreted as the present discounted sum of future expected inflation rates in the recursive formulation of the optimal price-setting equation[22]. The κ's are known functions of the underlying deep parameters (including the one governing belief distortion) defined in the Appendix and $\hat{\mu}_t$ measures the deviation of the distortion μ_t from its steady state level $-\overline{\mu}$. In particular, $\hat{\mu}_t > 0$ implies $\mu_t > -\overline{\mu}$. As a result, $\hat{\mu}_t$ cannot take on negative values since $\hat{\mu}_t = 0$ corresponds to the lower bound of the interval $[-\overline{\mu},\ \overline{\mu}]$. In fact, we will demonstrate that welfare is increasing in $\hat{\mu}_t$ in all states of the economy $(a_t, \hat{\Delta}_{t-1})$ under mild conditions, so that we can simply plug $\hat{\mu}_t = 0$ into equation (38).

To verify our conjecture that the value function is increasing around $-\overline{\mu}$ we will proceed in steps:

  i. We will solve for the linear policy functions as a function of the state of the economy. At this stage we will consider μ_t as an exogenous variable, i.e. our policy functions will represent the optimal decision for a given level of $\hat{\mu}_t$. In doing so, we maintain the assumption that the minimizing agent will apply the same distortion to all future expected levels of the policy rate.

  In particular a solution consists of four linear policy functions mapping $\left\{a_t,\ \hat{\Delta}_{t-1},\ \hat{\mu}_t\right\}$ into $\left\{\tilde{y}_t,\ \pi_t,\ \hat{\Delta}_t,\ \widehat{F2}_t\right\}$. In the appendix we report the general solution to this guess-and-verify problem, while in the main body of the text we focus on the special case of linear utility from leisure which allows for an analytic solution and helps build the intuition.

  ii. We will then plug the linear policy functions into the non-linear value function (once more following Ilut and Schneider, 2014).

  iii. For the case in which hours enter the felicity function linearly and the technology process distribution has bounded support, we will provide analytic sufficient conditions under which our value function is increasing around $\hat{\mu}_t = 0$, i.e. around $\mu_t = -\overline{\mu}$.

  iv. We will finally verify our conjecture numerically for our preferred calibration, in which we maintain a unitary level of the Frisch elasticity.

---

[22]Indeed it corresponds to the log-linearized version of the denominator of the expression on the RHS of equation (13).

C.2.1 Special Case: Linear Hours

When hours enter the felicity function linearly ($\psi = 0$), the policy functions simplify and the coefficients can be computed analytically. In particular the policy functions are:

$$\pi_t = -\frac{1}{\phi - 1}\hat{\mu}_t \tag{42}$$

$$\widehat{F2}_t = -\frac{\kappa_5}{(\phi - 1)(1 - \kappa_6)}\hat{\mu}_t \tag{43}$$

$$\tilde{y}_t = -\frac{\kappa_1 \lambda_F + (\kappa_2 - 1)\lambda_\pi}{\kappa_0}\hat{\mu}_t \tag{44}$$

$$\hat{\Delta}_t = \kappa_3 \hat{\Delta}_{t-1} - \frac{\kappa_4}{\phi - 1}\hat{\mu}_t \tag{45}$$

Where $\lambda_F = -\frac{\kappa_5}{(\phi-1)(1-\kappa_6)}$, $\lambda_\pi = -\frac{1}{\phi-1}$ and all the κ's are evaluated at the worst-case steady state though we do not explicitly write this out here for the sake of notation clarity. In this case, neither a_t nor $\hat{\Delta}_{t-1}$ appear in the policy functions. The first is a well-known consequence of the natural rate being included in the Taylor rule, an effect that carries over even to setups in which there is price dispersion in steady state. The second results from linearity in the utility from leisure. The policy function for inflation is particularly interesting, as it directly reflects our discussion of steady state inflation. In the worst-case steady state, inflation is inefficiently high. For lower levels of the beliefs' distortion ($\hat{\mu}_t \geq 0$) inflation will fall, i.e. it will get close to its first-best value. And again, the value of ϕ will be critical. A higher responsiveness to inflation movements will reduce the steady state distortion and will, consequently, reduce the responsiveness of inflation around steady state. Indeed, ϕ shows up at the denominator in all the policy functions, so higher levels of ϕ reduce the effects on all the variables of changes in the distortion of beliefs. $\widehat{F2}_t$ is a discounted sum of future expected inflation rates, hence it moves in the same direction with inflation. The sign of the response of the output gap, on the other hand, varies with the steady-state degree of ambiguity. For low levels of ambiguity, inflation and the output gap move in the same direction as in the standard case, *i.e.* high inflation would correspond to an inefficiently high output gap hence a fall in inflation will correspond to a fall in $\tilde{y}_t$ from a positive value down towards zero. For sufficiently high degrees of ambiguity, however, the steady-state inefficient wedge between hours worked and output ($\Delta(-\overline{\mu}, \cdot)$) grows faster than hours worked and so output will end up below potential even in the face of high steady state inflation (as discussed in Ascari and Sbordone (2014)). As a consequence, a reduction in the distortion around the worst-case steady state will induce a reduction in the output gap in the former case and an increase in the latter, a fact that will play a role when defining sufficient conditions. Finally, it is important to bear in mind that, while the the solution of the model in the four variables described above does not depend on the level of the technology process or of price dispersion, the variables that enter the agents' utility do:

$$c_t = \tilde{y}_t + a_t = \lambda_Y \hat{\mu}_t + a_t \tag{46}$$

$$n_t = \tilde{y}_t + \hat{\Delta}_t = \lambda_Y \hat{\mu}_t + \hat{\Delta}_t = \kappa_3 \hat{\Delta}_{t-1} + \left(\lambda_Y - \frac{\kappa_4}{\phi - 1}\right)\hat{\mu}_t \tag{47}$$

So ultimately the agents's welfare will vary with technology and price dispersion. Indeed, the value function for the problem, using the linear policy functions, can be expressed as:

$$\mathbb{V}\left(a_t, \hat{\Delta}_{t-1}; \hat{\mu}_t\right) = \log\left((1 + c_t)C(-\overline{\mu}, \cdot)\right) - (1 + n_t)N(-\overline{\mu}, \cdot) + \beta\mathbb{E}_t\mathbb{V}\left(a_{t+1}, \hat{\Delta}_t; \hat{\mu}_t\right) \qquad (48)$$

simply plugging in equations (46) and (47) for c_t and n_t. To build intuition, consider the linear approximation of the log function around one, $\log\left(1 + \lambda_Y\hat{\mu}_t + a_t\right) \simeq \lambda_Y\hat{\mu}_t + a_t$, and the fact that, for any reasonable calibration, values of hours are indeed very close to unity[23], $N(-\overline{\mu}, \cdot) \simeq 1$. This results in the value function becoming linear. If we refer to this approximation of the welfare function with v_t and substitute forward, we get:

$$v_t \simeq -\frac{\kappa_3}{1 - \beta\kappa_3}\hat{\Delta}_{t-1} + \mathbb{E}_t\left[\frac{a_t}{1 - \beta L^{-1}}\right] + \frac{\kappa_4}{(1 - \kappa_3)(\phi - 1)}\left(\frac{1}{1 - \beta} - \frac{\kappa_3}{1 - \beta\kappa_3}\right)\hat{\mu}_t$$

The expression above confirms our intuition that welfare is increasing in the level of technology and in $\hat{\mu}_t$ (i.e. a lowering of the beliefs' distortion), while it decreases with higher levels of the price dispersion term. This expression is trivially minimized for $\hat{\mu}_t = 0$[24] since $0 \leq \kappa_3 < 1$ and $\kappa_4 > 0$. We can interpret v_t as the average effect of a change in μ_t around the worst-case steady state but since v_t is linear it will not inform us on whether $\hat{\mu}_t > 0$ will improve welfare in all states of the economy. Fortunately, it is immediate to note that the value function (48) can be re-written as:

$$\mathbb{V}\left(a_t, \hat{\Delta}_{t-1}; \hat{\mu}_t\right) = u\left(\overrightarrow{C}(-\overline{\mu}, \cdot)\right) + \log\left(1 + c_t\right) - n_tN(-\overline{\mu}, \cdot) + \beta\mathbb{E}_t\mathbb{V}\left(a_{t+1}, \hat{\Delta}_t; \hat{\mu}_t\right)$$

So we can, without any approximation at this stage, define:

$$d\mathbb{V}\left(a_t, \hat{\Delta}_{t-1}; \hat{\mu}_t\right) \equiv \mathbb{V}\left(a_t, \hat{\Delta}_{t-1}; \hat{\mu}_t\right) - \mathbb{V}(-\overline{\mu}, \cdot) \qquad (49)$$

that is the non-linear difference between the welfare in any state of the economy and welfare in the worst-case steady-state, which can be expressed as:

$$
\begin{aligned}
d\mathbb{V}\left(a_t, \hat{\Delta}_{t-1}; \hat{\mu}_t\right) &= \log\left(1 + c_t\right) - n_tN(-\overline{\mu}, \cdot) + \beta\mathbb{E}_t d\mathbb{V}\left(a_{t+1}, \hat{\Delta}_t; \hat{\mu}_t\right) \\
d\mathbb{V}\left(a_t, \hat{\Delta}_{t-1}; \hat{\mu}_t\right) &= \log\left(1 + \lambda_Y\hat{\mu}_t + a_t\right) - \left(\left(\lambda_Y + \kappa_4\lambda_\pi\right)\hat{\mu}_t + \kappa_3\hat{\Delta}_{t-1}\right)N(-\overline{\mu}, \cdot) \quad (50) \\
&+ \beta\mathbb{E}_t d\mathbb{V}\left(a_{t+1}, \hat{\Delta}_t; \hat{\mu}_t\right)
\end{aligned}
$$

Now the technology process enters the expression nonlinearly. So, while a change in $\hat{\mu}_t$ will affect the utility of leisure in a way that is independent of price dispersion and technology, its effect on the utility from consumption will depend on the level of the technology process. In particular, imagine a situation in which $\lambda_Y < 0$. Equations (46) and (47) show that both consumption and hours will fall. From our analysis above, we know that *on average* the increased amount of leisure will compensate for the fall in consumption. Whether the increased amount of leisure will *actually* compensate for the lower utility from consumption in a given state of the world will, however, depend on the

---

[23]For our preferred calibration we obtain 1.0025.

[24]Remember that $\hat{\mu}_t$ cannot go negative as described above.

marginal utility of consumption. In particular, when a_t is very low, the marginal utility of consumption can be high to the point where $\hat{\mu}_t = 0$ does not minimize welfare. The following result formalizes this intuition, providing sufficient conditions for the welfare function to be increasing in $\hat{\mu}_t$ around $\mu_t = -\overline{\mu}$.

Result C.1. *Consider the economy defined above with $\psi = 0$, $0 \leq \kappa_3 < 1$, $\kappa_4 \geq 0$ and a_t having bounded support. Given linear policy functions, the representative agent's welfare is increasing in $\hat{\mu}_t$ around $\hat{\mu}_t = 0$ $\forall$ $(a_t,\ \hat{\Delta}_{t-1}) \in [\underline{a},\ \overline{a}] \times \mathbb{R}$:*

i. if $\frac{\lambda_Y}{(1+\overline{a})(1-\beta)} \geq \Xi(-\overline{\mu}, \cdot)$, when $\lambda_Y > 0$

ii. always, when $\lambda_Y = 0$

iii. if $\frac{\lambda_Y}{(1+\underline{a})(1-\beta)} \geq \Xi(-\overline{\mu}, \cdot)$, when $\lambda_Y < 0$

where $\Xi(-\overline{\mu}, \cdot) \equiv N(-\overline{\mu}, \cdot)\left(\frac{\lambda_Y}{1-\beta} + \frac{\kappa_4 \lambda_\pi}{(1-\kappa_3)}\left(\frac{1}{1-\beta} - \frac{\kappa_3}{1-\beta\kappa_3}\right)\right)$.

Proof. See Appendix C.2.1 $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

First note that the conditions on κ_3 and κ_4 amount to imposing stationarity in the dynamics of price dispersion around its steady state and the fact that inflation above steady state will result in higher price dispersion. The former is always verified if our model is to have solution, the second is always verified so long as in steady state inflation is above zero (which is the case in our worst case). Assuming a_t has bounded support is only needed to avoid invoking certainty equivalence. More importantly, in our experience, $\Xi(-\overline{\mu}, \cdot) < 0$ for any value of $\overline{\mu}$ so, in reality, the only potentially binding condition is the third. Moreover, if we take the calibration in Table 2 and simply set $\psi = 0$, we find that λ_Y is positive for any degree of ambiguity greater than 7bp, which is less than half the lowest value our series of expectations disagreement would imply. In other words, for any realistic calibration of $\overline{\mu}$ our sufficient conditions are trivially met. Yet, to show the robustness of the result, let us suppose ambiguity was as low as 5bp, even in that case any $\underline{a} > -.73$ would do. So, provided TFP could not fall below steady state by more than about 73 percent, our sufficient condition would be met.

C.2.2 Our preferred calibration: unitary Frisch elasticity

Linearity in the disutility of working is a convenient assumption to illustrate our point analytically, but implies an unrealistic (infinite) value for Frish elasticity. That is why our preferred calibration assumes $\psi = 1$, or a unitary Frisch elasticity. Under this assumption we can still work out the nonlinear value function given linear policy functions, exploiting the fact that hours enter quadratically in this specification. However the term in the past level of the price dispersion does not drop out of the policy functions and the coefficients cannot be computed analytically. In particular, following the same steps laid out in the previous paragraph it is easy to obtain:

$$d\mathbb{V}\left(a_t, \hat{\Delta}_{t-1}; \hat{\mu}_t\right) = \log\left(1 + c_t\right) - N\left(-\overline{\mu}, \cdot\right)^2 \left(n_t + \frac{1}{2}n_t^2\right) + \beta\mathbb{E}_t d\mathbb{V}\left(a_{t+1}, \hat{\Delta}_t; \hat{\mu}_t\right)$$

To derive our sufficient conditions (see Appendix C.3 for a complete derivation) we then compute the derivative of $d\mathbb{V}$ with respect to $\hat{\mu}_t$ and then use the conjecture that $\hat{\Delta}_t = 0$ at

all times. This amounts to saying that our sufficient conditions have been in place for ever. Under this conjecture, it is possible to verify that under our preferred parametrization:

i. the term corresponding to Ξ in the previous derivation is negative[25].

ii. the derivative is decreasing in the level of technology[26]

iii. the sufficient condition that the derivative is increasing around $\hat{\mu}_t = 0$ is then met for any level of the state variables.

iv. this in turn validates our conjecture that both μ_t and Δ_t never deviate from their worst-case steady-state level.

With this, we can proceed with rest of our analysis, taking for granted that the worst-case will correspond to the lower bound of the interval considered by agents in all states of the world.

C.3   Non-Linear Hours: $\psi = 1$

When $\psi = 1$, hours enter our felicity function quadratically, hence we can still use a decomposition similar to the one we used in the linear case to obtain:

$$d\mathbb{V}\left(a_t, \hat{\Delta}_{t-1}; \hat{\mu}_t\right) = \log\left(1 + c_t\right) - N\left(-\overline{\mu}, \cdot\right)^2\left(n_t + \frac{1}{2}n_t^2\right) + \beta\mathbb{E}_t d\mathbb{V}\left(a_{t+1}, \hat{\Delta}_t; \hat{\mu}_t\right) \tag{51}$$

The algebra becomes a bit more cumbersome, however, because the coefficients on lagged price dispersion in the policy functions are no longer zero:

$$\begin{aligned}
d\mathbb{V}\left(a_t, \hat{\Delta}_{t-1}; \hat{\mu}_t\right) &= \log\left(1 + \lambda_{y\Delta}\hat{\Delta}_{t-1} + \lambda_{y\mu}\hat{\mu}_t + a_t\right) \\
&\quad - N\left(-\overline{\mu}, \cdot\right)^2\left(\left(\lambda_{y\Delta}\hat{\Delta}_{t-1} + \lambda_{y\mu}\hat{\mu}_t + \hat{\Delta}_t\right) + \frac{1}{2}\left(\lambda_{y\Delta}\hat{\Delta}_{t-1} + \lambda_{y\mu}\hat{\mu}_t + \hat{\Delta}_t\right)^2\right) \\
&\quad + \beta\mathbb{E}_t d\mathbb{V}\left(a_{t+1}, \hat{\Delta}_t; \hat{\mu}_t\right)
\end{aligned} \tag{52}$$

Substituting forward:

$$\begin{aligned}
d\mathbb{V}\left(a_t, \hat{\Delta}_{t-1}; \hat{\mu}_t\right) &= \mathbb{E}_t\sum_{j=0}^{\infty}\beta^j\log\left(1 + \lambda_{y\Delta}\hat{\Delta}_{t+j-1} + \lambda_{y\mu}\hat{\mu}_t + a_{t+j}\right) \\
&\quad - N\left(-\overline{\mu}, \cdot\right)^2\sum_{j=0}^{\infty}\beta^j\left(\left(\lambda_{y\Delta}\hat{\Delta}_{t+j-1} + \lambda_{y\mu}\hat{\mu}_t + \hat{\Delta}_{t+j}\right) + \frac{1}{2}\left(\lambda_{y\Delta}\hat{\Delta}_{t+j-1} + \lambda_{y\mu}\hat{\mu}_t + \hat{\Delta}_{t+j}\right)^2\right)
\end{aligned} \tag{53}$$

Then note:

$$\hat{\Delta}_{t+j} = \gamma_{\Delta}^{j+1}\hat{\Delta}_{t-1} + \kappa_4\lambda_{\pi\mu}\sum_{l=0}^{j}\gamma_{\Delta}^l\hat{\mu}_t \quad \gamma_{\Delta} \equiv \left(\kappa_3 + \kappa_4\lambda_{\pi\Delta}\right) \tag{54}$$

Note that:

---

[25]Zero in the limit case in which $\overline{m}u = 0$

[26]Similar to the first of the three cases in the previous paragraph.

- shocks to TFP are non-inflationary in this economy and hours do not respond to a_t at all.

- hours only depend on $\hat{\mu}_t$, which is a choice variable for the purpose of this exercise, hence there is no uncertainty involved. Intuitively, after the minimization step is carried out, the path for hours in the foreseeable future is known without uncertainty, because the only source of randomness (the productivity shock) will not affect the level of hours worked in the future.

In this economy, TFP shocks do not generate inflation. If, on top of that, $\hat{\mu}_t = 0$ in every period, $\hat{\Delta}_{t-1} = 0$ at all times. So we conjecture that $\hat{\Delta}_{t-1} = 0$, which amounts to assuming that the conditions that make $\hat{\mu}_t = 0$ have been holding for ever and then work out the sufficient conditions that guarantee that to be the case. Using our conjecture:

$$\hat{\Delta}_{t+j} = \kappa_4 \lambda_{\pi\mu} \frac{1 - \gamma_\Delta^{j+1}}{1 - \gamma_\Delta} \hat{\mu}_t \tag{55}$$

So now we can take the derivative of $d\mathbb{V}$ w.r.t. $\hat{\mu}_t$:

$$\frac{\partial \left(d\mathbb{V} \left(a_t, \ \hat{\Delta}_{t-1} \right) \right)}{\partial \hat{\mu}_t} = \sum_{j=0}^{\infty} \beta^j \frac{\lambda_{y\Delta} \kappa_4 \lambda_{\pi\mu} \frac{1-\gamma_\Delta^j}{1-\gamma_\Delta} + \lambda_{y\mu}}{\left(1 + \lambda_{y\Delta} \hat{\Delta}_{t+j-1} + \lambda_{y\mu} \hat{\mu}_t + a_{t+j} \right)}$$

$$- N\left(-\overline{\mu}, \cdot\right)^2 \sum_{j=0}^{\infty} \beta^j \left(\left(\lambda_{y\Delta} \kappa_4 \lambda_{\pi\mu} \frac{1-\gamma_\Delta^j}{1-\gamma_\Delta} + \lambda_{y\mu} + \kappa_4 \lambda_{\pi\mu} \frac{1-\gamma_\Delta^{j+1}}{1-\gamma_\Delta} \right) \left(1 + \lambda_{y\Delta} \hat{\Delta}_{t+j-1} + \lambda_{y\mu} \hat{\mu}_t + \hat{\Delta}_{t+j} \right) \right)$$

If we now evaluate the derivative at $\hat{\mu}_t = 0$ and $\hat{\Delta}_t = 0$ at all times we get:

$$\left. \frac{\partial \left(d\mathbb{V} \left(a_t, \ \hat{\Delta}_{t-1} \right) \right)}{\partial \hat{\mu}_t} \right|_{\hat{\mu}_t = 0, \hat{\Delta}_{t-1} = 0} = \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \frac{\lambda_{y\Delta} \kappa_4 \lambda_{\pi\mu} \frac{1-\gamma_\Delta^j}{1-\gamma_\Delta} + \lambda_{y\mu}}{(1 + a_{t+j})}$$

$$- N\left(-\overline{\mu}, \cdot\right)^2 \sum_{j=0}^{\infty} \beta^j \left(\lambda_{y\Delta} \kappa_4 \lambda_{\pi\mu} \frac{1-\gamma_\Delta^j}{1-\gamma_\Delta} + \lambda_{y\mu} + \kappa_4 \lambda_{\pi\mu} \frac{1-\gamma_\Delta^{j+1}}{1-\gamma_\Delta} \right)$$

Or:

$$\left. \frac{\partial \left(d\mathbb{V} \left(a_t, \ \hat{\Delta}_{t-1} \right) \right)}{\partial \hat{\mu}_t} \right|_{\hat{\mu}_t = 0, \hat{\Delta}_{t-1} = 0} = \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \frac{\lambda_{y\Delta} \kappa_4 \lambda_{\pi\mu} \frac{1-\gamma_\Delta^j}{1-\gamma_\Delta} + \lambda_{y\mu}}{(1 + a_{t+j})} - \frac{N\left(-\overline{\mu}, \cdot\right)^2 \kappa_4 \lambda_{\pi\mu}}{(1-\beta)(1-\gamma_\Delta)} \left(1 + \lambda_{y\Delta} + \lambda_{y\mu} \right)$$

$$+ \frac{N\left(-\overline{\mu}, \cdot\right)^2 \kappa_4 \lambda_{\pi\mu}}{(1-\beta\gamma_\Delta)(1-\gamma_\Delta)} \left(\lambda_{y\Delta} + \gamma_\Delta \right)$$

$$= \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \frac{\lambda_{y\Delta} \kappa_4 \lambda_{\pi\mu} \frac{1-\gamma_\Delta^j}{1-\gamma_\Delta} + \lambda_{y\mu}}{(1 + a_{t+j})} - \Xi\left(-\overline{\mu}, \cdot\right) \tag{56}$$

Where $\Xi\left(-\overline{\mu}, \cdot\right) \equiv \frac{N(-\overline{\mu}, \cdot)^2 \kappa_4 \lambda_{\pi\mu}}{(1-\beta)(1-\gamma_\Delta)} \left(1 + \lambda_{y\Delta} + \lambda_{y\mu} \right) - \frac{N(-\overline{\mu}, \cdot)^2 \kappa_4 \lambda_{\pi\mu}}{(1-\beta\gamma_\Delta)(1-\gamma_\Delta)} \left(\lambda_{y\Delta} + \gamma_\Delta \right)$.

As before, the marginal-disutility-from-labor block is just a number, the only source of uncertainty pertaining to the marginal utility of consumption. Again we can exploit the fact that, on a period-by-period basis, the marginal utility is monotonic in the level of technology, to compute sufficient conditions in the form of bounds

on the expected value. Clearly, we now have to rely on the numerical values of λ's. As it turns out, for our preferred calibration ($\overline{\mu} = 50\text{bp}$), $\Xi(-\overline{\mu}, \cdot) < 0$. The term governing the marginal utility of consumption is more complicated now because of the presence of the γ_Δ^j term. Numerically it possible to verify that the individual coefficients in the sum are positive. Hence the sum is minimized at $\overline{a}$[27], but even then it would correspond to positive number. In other words our sufficient condition is met for any value of the states.

C.4 Linear Approximation to the Welfare Function

The case in which utility is linear in hours lends itself to a very convenient approximation of the welfare function so this is the avenue we pursue, but it is easy to verify numerically that these results are robust to different values of ψ. When $\psi = 0$ our one-period felicity can be approximated as:

$$
\begin{aligned}
u_t &= \log\left(\frac{(1 + a_t)A(1 + n_t)N(-\overline{\mu}, \cdot)}{(1 + \hat{\Delta}_t)\Delta(-\overline{\mu}, \cdot)}\right) - (1 + n_t)N(-\overline{\mu}, \cdot) \\
&\simeq Const + (1 - N(-\overline{\mu}, \cdot))n_t - \hat{\Delta}_t + a_t \\
&\simeq Const + (1 - N(-\overline{\mu}, \cdot))\left(\hat{\Delta}_t + c_t - a_t\right) - \hat{\Delta}_t + a_t \\
&\simeq Const - N(-\overline{\mu}, \cdot)\hat{\Delta}_t + N(-\overline{\mu}, \cdot)a_t + (1 - N(-\overline{\mu}, \cdot))c_t
\end{aligned}
$$

We know that $N(-\overline{\mu}, \cdot) \simeq 1$[28] so variations in the effects of changes in $\hat{\mu}_t$ can be safely approximated by $-\hat{\Delta}_t$. Since the steady state value of Δ is also very close to one for the sake of this analysis we will use log differences and level differences interchangeably and obtain:

$$
v_t \simeq -\hat{\Delta}_t + a_t + \beta\mathbb{E}_t v_{t+1}
$$

Using the law of motion $\hat{\Delta}_t$ and π_t delivers the result in the main body of the text.

C.5 A policy-independent loss function

A quadratic approximation to the policymaker's loss function in the tradition of Woodford (2003) will serve us well for this purpose, because it is independent of policy and, as such, allows us not only to pin down optimal policy but also to rank suboptimal alternatives. Specifically, we follow Coibion, Gorodnichenko and Wieland (2012) who derive an approximation suitable for an environment featuring

---

[27]We are obviously ruling out values of $a < -1$ because they would imply negative TFP in levels.

[28]It equals 1 when $\overline{\mu} = 0$ and it takes on values of the order of 1.01 or smaller for reasonable degrees of ambiguity, given our calibration (it equals 1.00054 for our baseline calibration).

trend inflation and obtain the following:

$$\mathcal{L}_t \;=\; \sum_{j=0}^{\infty} \beta^j \left(\Theta_0 + \Theta_1 var(\tilde{y}_t) + \Theta_2 var(\pi_t) \right) \tag{57}$$

One key difference, relative to the standard case is that a constant shows up, measuring systematic loss. In fact, in our results we will discuss policy optimality along two dimension (which mimic the discussion in Woodford (2003, p. 412)):

I. Systematic or *static* (or, again, average in Woodford's words), i.e. the loss that emerges even in steady state

II. *Dynamic*, the inefficiency that emerges when shocks buffet the economy

Another crucial consideration is that, as it turns out, $\Theta_1 = \frac{1+\psi}{2}$ is independent of ambiguity while Θ_2 is a complicated increasing function of ambiguity. This is not really surprising since Coibion, Gorodnichenko and Wieland (2012) find that the corresponding parameter on inflation variability is increasing in trend inflation and we have documented above that trend inflation is a positive function of ambiguity in our setting. It is, however, very important for our analysis, because it shows that the lower the confidence, the more the central bank has to "focus" on inflation, possibly at the cost of not responding to variations in the output gap. In other words, ambiguity exacerbates the effects of tradeoff-inducing shocks. Endowed with a welfare-based loss function, we can now turn to characterizing optimal policy, which we do in the next paragraph.

C.6 Proof of Proposition 3.1

The only purpose of the lower bound $\underline{\phi}\left(-\overline{\mu}, \cdot\right) \geq 1$ is to ensure equilibrium determinacy. It is not otherwise relevant as it will always be optimal to have as high a ϕ as possible. Computing the steady state of the model, it is easy to verify that:

$$\Pi\left(\mu, \delta\left(\cdot\right), \cdot\right) = e^{-\frac{\mu+\delta(\cdot)}{\phi-1}} \tag{58}$$

while all the other steady-state expressions, as a function of inflation, remain unchanged. Hence, if the denote with $\mathbb{V}_\delta$ the value function of the economy in which δ enters the Taylor rule we get that:

$$\mathbb{V}_\delta(\mu, \cdot) = \mathbb{V}(\mu + \delta(\cdot), \cdot) \tag{59}$$

Graphically, this amounts to shifting the function leftward by δ. So $\mathbb{V}_\delta$ inherits all the properties of $\mathbb{V}$ established in Results 3.2 and 3.3, except $\mathbb{V}_\delta$ is maximized at $-\delta$, the value of μ delivering zero steady state inflation. Having established this, the proof of static optimality proceeds in three steps by first assuming a range for δ and verifying the optimal value of ϕ over that range, then verifying that for the

optimal value of ϕ the optimal value of δ is pinned down by the equality in our proposition and, finally, by establishing that the optimal value of δ indeed falls in the range we assumed in the first part of our proof.

   i. $\overline{\phi}$ *is the welfare-maximizing value of* $\phi \in [\underline{\phi}, \overline{\phi}]$ $\forall$ $\overline{\mu} > \delta > 0$.

Following the same logic as in Result A.1, it is easy to verify that for $1 < \phi' < \overline{\phi}$, there exists a μ' s.t.:

$$\mathbb{V}\left(\mu', \phi', \cdot\right) = \mathbb{V}\left(-\overline{\mu}, \overline{\phi}, \cdot\right) \tag{60}$$

In particular:

$$\mu' \;\; = \;\; -\overline{\mu}\frac{\phi' - 1}{\overline{\phi} - 1} - \delta\frac{\overline{\phi} - \phi'}{\overline{\phi} - 1} \tag{61}$$

Our restriction on δ implies that $\frac{\partial \mu'}{\partial \phi'} = \frac{-\overline{\mu}+\delta}{\overline{\phi}-1} < 0$. Since $\phi' \in (1, \overline{\phi})^{29}$ we know that $0 > -\delta > \mu' > -\overline{\mu}$. Strict concavity (Result 3.2) and the fact that the maximum is attained at $-\delta$ then imply:

$$\mathbb{V}\left(-\overline{\mu}, \overline{\phi}, \cdot\right) = \mathbb{V}\left(\mu', \phi', \cdot\right) > \mathbb{V}\left(-\overline{\mu}, \phi', \cdot\right) \tag{62}$$

   ii. $\delta^*(\overline{\mu}, \overline{\phi}; \cdot)$ *defined by* $\mathbb{V}\left(-\overline{\mu} + \delta^*(\overline{\mu}, \overline{\phi}; \cdot), \cdot\right) = \mathbb{V}\left(\overline{\mu} + \delta^*(\overline{\mu}, \overline{\phi}; \cdot), \cdot\right)$ *is welfare maximizing for* $\phi = \overline{\phi}$.

The following lemma characterizes the optimal level of δ under very general conditions.

Lemma C.1. *Assuming that* $\mathbb{V}(\mu, \cdot)$ *takes only real values over some interval* $(-\overline{m}, \overline{m})$, *is continuous, strictly concave and attains a finite maximum at* $\mu = \mu_0 \in (-\overline{m}, \overline{m})$; *if* ϕ *is fixed and* $\overline{\mu} > 0$, *then the optimal level of* δ *is pinned down by the following condition.*

$$\delta^*(\overline{\mu}): \quad \mathbb{V}\left(-\overline{\mu} + \delta^*(\overline{\mu}), \cdot\right) = \mathbb{V}\left(\overline{\mu} + \delta^*(\overline{\mu}), \cdot\right) \tag{63}$$

Proof. First we define μ_0 to be the value that maximizes $\mathbb{V}\left(\mu, \cdot\right)$. Strict concavity ensures it is unique.

A number of different cases then arise:

   1. $\mu_0 \in (-\overline{\mu}, \overline{\mu})$: then $\mathbb{V}'(-\overline{\mu}, \cdot) > 0 > \mathbb{V}'(\overline{\mu}, \cdot)^{30}$

      a. $\mathbb{V}(-\overline{\mu}, \cdot) < \mathbb{V}(\overline{\mu}, \cdot)$. Together with strict concavity this implies that

---

[29] Where we consider the lowest possible value for $\underline{\phi}$, i.e. unity

[30] With an abuse of notation we use derivatives here but we do not need differentiability. We just need the function to be strictly increasing and strictly decreasing for values of μ respectively smaller and larger than μ_0, which is ensured by strict concavity.

$\mu_{ws} = -\overline{\mu}$. Then there exists a small enough $\delta > 0$ such that

$$\mathbb{V}(-\overline{\mu}, \cdot) < \mathbb{V}(-\overline{\mu} + \delta, \cdot) < \mathbb{V}(\overline{\mu} + \delta, \cdot) < \mathbb{V}(\overline{\mu}, \cdot).$$

So now the worst case $\mu'_{ws} = -\overline{\mu} + \delta$ generates a higher level of welfare. The worst-case welfare can be improved until the second inequality above holds with equality. Continuity ensures such a level of δ^* exists. Any value of $\delta > \delta^*$ will, however, make welfare in the worst case decrease, and the second inequality above would reverse the sign.

    b. $\mathbb{V}(-\overline{\mu}, \cdot) > \mathbb{V}(\overline{\mu}, \cdot)$. Together with strict concavity, this implies that $\mu_{ws} = \overline{\mu}$. Then there exists a small enough $\delta < 0$ delivering the same as above

    c. $\mathbb{V}(-\overline{\mu}, \cdot) = \mathbb{V}(\overline{\mu}, \cdot)$. There is no room for improvement. Any $\delta \neq 0$ would lower the worst-case welfare.

2. $\mu_0 \geq \overline{\mu}$. Strict concavity implies that $\mathbb{V}(-\overline{\mu}, \cdot) < \mathbb{V}(\overline{\mu}, \cdot)$. Hence $\mu_{ws} = -\overline{\mu}$. For all $0 \leq \delta \leq \mu_0 - \overline{\mu}$

$$\mathbb{V}(-\overline{\mu}, \cdot) < \mathbb{V}(-\overline{\mu} + \delta, \cdot) < \mathbb{V}(\overline{\mu} + \delta, \cdot) \leq \mathbb{V}(\mu_0, \cdot)$$

For δ just above $\mu_0 - \overline{\mu}$ we fall in case 1a above.

3. $\mu_0 \leq -\overline{\mu}$. Strict concavity implies that $\mathbb{V}(-\overline{\mu}, \cdot) > \mathbb{V}(\overline{\mu}, \cdot)$. Hence $\mu_{ws} = \overline{\mu}$. For all $\mu_0 - \overline{\mu} \leq \delta \leq 0$

$$\mathbb{V}(\mu_0, \cdot) \geq \mathbb{V}(-\overline{\mu} + \delta, \cdot) > \mathbb{V}(\overline{\mu} + \delta, \cdot) > \mathbb{V}(\overline{\mu}, \cdot)$$

For δ just above $\mu_0 - \overline{\mu}$ we fall in case 1b above.

$\square$

The conditions of the Lemma apply to our case for sufficiently small degrees of ambiguty (as maintained in Results 3.2 and 3.3), which also ensure that $\overline{m}$ meets the condition in equation (25), which we assume throughout. The Lemma holds for a generic fixed ϕ, so it obviously applies to $\phi = \overline{\phi}$.

iii. *in the economy described in Section 3.1, $0 < \delta^*(\overline{\mu}, \overline{\phi}; \cdot) < \overline{\mu}$.*

Results 3.2 and 3.3 ensure that our economy falls under case 1a of Lemma C.1. This proves that $\delta^*(\overline{\mu}) > 0$. Suppose now that $\delta^*(\overline{\mu}) \geq \overline{\mu}$. That would push the argmax of the welfare function outside (or on the boundary) of $[-\overline{\mu}, \overline{\mu}]$, which can never be optimal given strict concavity (similar arguments to cases 2 and 3 in Lemma C.1).

These three points complete the proof of the static optimality of our proposed rule. Dynamic optimality follows immediately by noting that the first-order solution to our model (Appendix B.1) implies that both inflation and the output gap do not vary with TFP, hence the variance of both the output gap and inflation - equation (57) - is minimized. $\qquad\square$

C.7 Proof of Corollary 3.1

This sufficient condition can be derived even for $\delta = 0$, so we will assume that for expositional simplicity. Setting δ optimally will simply make the *suitably high level of* $\overline{\phi}$ somewhat lower.

Consider any policy plan delivering utility v_0 in steady state. Suppose that is welfare-superior to the policy currently in place $\mathbb{V}(-\overline{\mu}, \phi, \cdot) < v_0 \leq \mathbb{V}(0, \phi, \cdot)$, where the latter is the first-best allocation so it cannot be improved upon.

Results 3.2 and 3.3 ensure the value function is strictly increasing for $\mu < 0$ so there exists a μ', $-\overline{\mu} < \mu' \leq 0$, s.t. $\mathbb{V}(\mu', \phi, \cdot) \geq v_0$. Result A.1 then ensures there also exists ϕ' s.t. $\mathbb{V}(-\overline{\mu}, \phi', \cdot) = \mathbb{V}(\mu', \phi, \cdot)$. Since our proposed monetary policy rule is also dynamically optimal for any value of ϕ that guarantees determinacy, any $\overline{\phi} \geq \phi'$ allows our monetary policy to deliver at least as high a welfare level as the alternative delivering v_0. $\qquad\square$