

Ambiguity, Monetary Policy and Trend Inflation*

Riccardo M. Masolo

Francesca Monti

Bank of England and CfM

October 8, 2015

Abstract

We develop a model that can explain the evolution of trend inflation in the US in the three decades before the Great Recession as a function of the reduction in uncertainty about the monetary policymaker's behaviour. The model features ambiguity-averse agents and ambiguity regarding the conduct of monetary policy, but is otherwise standard. Trend inflation arises endogenously and has these determinants: the strength with which the central bank responds to inflation, the degree of uncertainty about monetary policy perceived by the private sector, and, if it exists, the inflation target. Given the importance of monetary policy for the determination of trend inflation, we also study optimal monetary policy in the case of lingering ambiguity.

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

Keywords: Ambiguity aversion, monetary policy, trend inflation

1 Introduction

Positive trend inflation is an important feature of the data, but the great majority of macroeconomic models popular in the academic literature are approximated around a zero-inflation steady state. Trend inflation, however, has far reaching effects on the static and dynamic properties of the model and should be accounted for when

*Previously circulated as *Monetary Policy with Ambiguity Averse Agents*. We are grateful to Guido Ascari, Carlos Carvalho, Richard Harrison, Cosmin Ilut, Peter Karadi and Wouter Den Haan. We would also like to thank seminar participants at Oxford University, Bank of Finland, ECB-WGEM, Bank of Korea, Bank of England and Bank of Canada, as well as participants at the 2015 North American Winter Meeting of the Econometric Society, 2015 SNDE, the XVII Inflation Targeting Conference at BCDB, Barcelona GSE summer forum and the 2015 EEA conference for useful comments and suggestions. *Any views expressed are solely those of the authors and so cannot be taken to represent those of the Bank of England or to state Bank of England policy. This paper should therefore not be reported as representing the views of the Bank of England or members of the Monetary Policy Committee or Financial Policy Committee.*

studying the conduct of monetary policy (see Ascari and Sbordone (2014) for an overview of the key results). Even when it is accounted for, trend inflation is not explicitly modeled, and is often considered to be the same as the inflation target. Our work, instead, provides a micro-foundation for trend inflation, linking it to the presence of Knightian uncertainty, or ambiguity, about the conduct of monetary policy. In particular, 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 uncertainty about monetary policy perceived by the private sector.

We illustrate our point in a prototypical new-Keynesian model in which agents are averse to ambiguity, and where the ambiguity regards the monetary policy rule. Ambiguity describes a situation in which there is uncertainty about the probability distribution over states of the world. In our model, ambiguity about the behaviour of the Central Bank has effects also in steady state. In particular, because agents will make their decisions based on a distorted belief about the interest rate rather than on the one actually set by the Central Bank, inflation in steady state will not coincide with the target. In many ways, the dynamics of the model we consider are quite similar to those of the models with exogenous trend inflation studied by Yun (2005) and Ascari and Ropele (2007), yet the policy implications are not necessarily the same. The differences here are that trend inflation 1) arises endogenously because of ambiguity and 2) is also a function the degree of inflation responsiveness of the rule itself. In this sense, our work provides a micro-foundation for trend inflation, an important topic which, as Coibion and Gorodnichenko (2011) point out, has not received great attention in the academic literature. In particular, we find that 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. Most state-of-the-art DSGE models, such as Del Negro and Eusepi (2011) and Del Negro, Schorfheide and Giannoni (2015), resort to a time-varying target inflation in order to capture the rise and fall of inflation and interest rates in that sample. But, as the above authors point out, the assumption that the changes in the target inflation rate are exogenous is somewhat of a convenient short-cut.

The data seems to support the idea that changes in the private sector's confidence and understanding of the policymaker's behavior contribute to determining the level of trend inflation. For example, Crowe and Meade (2008) find evidence that higher transparency (and thus presumably lower ambiguity) is related to lower inflation, in line with the prediction from our model. Another obvious implication of our model is that with less uncertainty about the policymaker's behavior, the beliefs about the interest rate should be less distorted and, thus, closer to the one actually set the Central Bank. This implication seems to have empirical validity as well. For example, Swanson (2006) finds 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 and less diverse in the cross-sectional variety of their interest rate forecasts. Swanson (2006) also shows evidence that strongly suggests increases in Federal Reserve transparency played a role in this improvement. While we do not model explicitly the link between Central Bank communication and the level of Knightian uncertainty about monetary policy, it is quite reasonable to assume that changes in the level of Central Bank transparency

are going to be associated with changes in the level of the public's confidence about their understanding of the monetary policymakers' actions.

In providing a link between trend inflation and Knightian uncertainty, we also contribute to the debate on the so-called *Taylor principle*. In its original formulation it amounted to positing that nominal rates should respond more than one for one to inflation deviations from its target. In simple New-Keynesian DSGEs this intuition carries over: if the central bank responds more than one for one to inflation, the equilibrium is determinate (Galí, 2008, among many others). But this is a by-product of the log-linearization being carried out around the zero-inflation steady-state or by adopting full price indexation¹. However, Ascari and Ropele (2009) and Coibion and Gorodnichenko (2011) clearly show that trend inflation has a dramatic impact on the lowest degree of inflation responsiveness that can ensure equilibrium determinacy. In our analysis we bring this concept one step forward, because trend inflation itself depends on the inflation responsiveness coefficient in the central bank's response function. In other words, in this model by increasing the degree to which it responds to inflation, a Central Bank will not only affect the model dynamics but also the steady state level of trend inflation.

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. In particular, it is optimal for policymakers to target a level of the real interest rate that is higher than the natural rate. Also, the higher the degree of uncertainty, the higher the weight on inflation variability in the policymaker's welfare-based loss function (which extends a result by Coibion, Gorodnichenko and Wieland, 2011, for the case of exogenous trend inflation).

Most of the existing work on ambiguity in macro models, e.g. Hansen and Sargent (2007), has focused on the perspective of a policymaker setting policy under model uncertainty. Here we consider the less studied case in which the private agents, rather than the policymaker, in the model face uncertainty. Bidder and Smith (2012), Adam and Woodford (2012) and Benigno and Paciello (2014) also study models in which the agents are ambiguity averse, but they model ambiguity aversion using the so-called "multiplier" preferences. According to these preferences, agents have a reference belief, but fear misspecification around that belief, so they want to make "robust" decisions. This implies that they evaluate plans using a worst-case belief that minimizes the sum of expected utility and a smooth function that penalizes deviations from a reference belief. We instead use a recursive version of the multiple priors preferences (see Gilboa and Schmeidler (1989) and Epstein and Schneider (2003)), pioneered in business cycle models by Ilut and Schneider (2014). One crucial difference among multiplier preferences and multiple priors utility is that multiple priors utility is not smooth when belief sets differ in mean, as in our model. So, by following the multiple-prior approach we can characterize the effects of ambiguity on the steady state, while the multiplier preferences approach implies that the effect of uncertainty can be identified only by approximating the model at higher orders than a simple linearization. This insight was already highlighted by Ilut and Schneider (2014), who introduce ambiguity

¹Typically to a combination of steady state and past inflation.

about the technology process in their business cycle model. With respect to Ilut and Schneider (2014), we extend the analysis to a situation in which agents have multiple priors on the expectations of an endogenous variable (the interest rate), a fact that amplifies the effects of the under- or over-estimation of the variable at hand.

The rest of the paper is organised as follows. In Section 2 we show that our model matches the main stylized facts about trend inflation and uncertainty about the interest rate. In Section 3 we illustrate the effects of ambiguity in a small linear model to build intuition for our results, while Section 4 provides a description of the model we use for our analysis. Section 5 characterizes the steady state of our economy as a function of the degree of ambiguity and Section 6 presents the log-linear version of the equilibrium conditions alongside the description of sufficient conditions regarding the worst-case definition as well as an analysis of impulse responses. Section 7 studies the optimal monetary policy in the presence of ambiguity and when the policymaker can dispel such ambiguity, while Section 8 is devoted to the conclusions.

2 Empirical evidence

In this Section we show some robust stylized facts about trend inflation and interest rate uncertainty in the United States in the years preceding the Great Recession. We then match them with the predictions of our model, which stack up surprisingly well, given a very standard calibration². We set the scene here, by giving a preview of some of our key theoretical results, further discussed in the subsequent Sections, that are relevant for the analysis of trend inflation.

One key implication of our simple model is that Knightian uncertainty about the behavior of the monetary policymaker drives a wedge between the inflation target and trend inflation (see Result 5.4). The intuition behind this result is that, because agents want to be robust with respect to the uncertainty about the monetary policy rule, they will base their decisions on the worse case scenario, *i.e.* on a distorted belief about the interest rate. The fact that the interest rate used for decision making is different from the actual one causes the divergence of trend inflation from the inflation target. This wedge depends on:

- a. our measure of Knightian uncertainty $\bar{\mu}$, which one can think of as the size of the deviations from the baseline Taylor rule the private sector is entertaining as possible. $\bar{\mu}$ can also be interpreted as measure of the distortion of their beliefs about the interest rate. Its unit of measure is the same as that of the policy rate, be it basis points or percentage points. It is also worth noting that a non-trivial implication of our model is that higher uncertainty leads to *higher* inflation for a given target, in other words our model has a prediction for the sign of the wedge (Result 5.2).
- b. the strength with which policymakers respond to inflation deviations from its target, measured by ϕ which is the coefficient on inflation in our simple Taylor

²The ZLB period cannot be analyzed with this model because a global-solution method would be required: this is the object of ongoing research.




Figure 1: Interdecile range of the forecast for the current quarter 3-month TBill rate from the Survey of Professional Forecasters (4-quarter moving average).

rule. In particular, $\frac{1}{\phi-1}$ measures the elasticity of inflation to our measure of uncertainty. Intuitively, a staunch response by policymakers limits the scope of the impact of uncertainty on inflation.

In sum, the key *accounting* equation for our analysis is the one that describes trend inflation, which, in log-form, reads:

$$\bar{\pi} = \pi^* + \frac{\bar{\mu}}{\phi - 1}, \quad (1)$$

where $\bar{\pi}$ is measured trend inflation, π^* the inflation target and $\bar{\mu} \geq 0$ as it governs the width of an interval around the Taylor rule, so it cannot be negative by definition. Equation (1) implies a tight restriction in that, for a given π^* and ϕ , changes in uncertainty should be reflected directly in changes in the inflation trend.

2.1 Stylized facts

The key stylized facts concerning the US economy that we aim to match are the following.

Trend inflation fell from around 5 percent in the late 70s to around 2 percent in the mid 90s and appears not to have moved much since. Coibion and Gorodnichenko (2011) and Ascari and Sbordone (2014), among others, document this clearly.

The equilibrium was indeterminate in the pre-Volcker era and determinate thereafter. This finding goes back to Clarida, Galí and Gertler (2000), but Coibion and Gorodnichenko (2011) convincingly show that the changes in trend




Figure 2: The 4-quarter moving average of the log of the Baker, Bloom and Davis (2015) Monetary Policy Uncertainty Index (solid, LHS axis) and the log of the Interdecile Range of the forecast for the 3-month TBill rate from the Survey of Professional Forecasters (4-quarter moving average) (dashed, RHS axis).

inflation generate a shift from the indeterminacy to the determinacy region in spite of a barely noticeable change in inflation responsiveness³.

Uncertainty about monetary policy followed a downward trend in the 80s and 90s. Following Drechsler (2013) and Ilut and Schneider (2014), for example, we quantify uncertainty by a measure of the dispersion in the set of forecasts from the Philadelphia Feds Survey of Professional Forecasters (SPF), particularly the forecasts of the current quarter's 3-month TBill rates⁴. Figure 1 reports the 4-quarter moving-average⁵ of the interdecile dispersion of this series, which displays an obvious downward trend. The degree of dispersion is an order of magnitude larger in the early 1980s than it is now. While our SPF series does not go back into the 70s, it is interesting to note that the timing of the biggest drop in uncertainty compares favorably with the analysis in Bianchi and Ilut (2015). Bianchi and Ilut (2015) allow for regime changes in monetary and fiscal policy, however, they discuss how inflation started to fall only when agents started to discount the possibility that

³Coibion and Gorodnichenko (2011) estimate a richer specification of the monetary policy rule and some of the parameters (e.g. the interest-rate smoothing coefficient) change across the two regimes, but crucially not the inflation responsiveness coefficient.

⁴We use the nowcasts - which are produced around the middle of the quarter when quite a bit of information about the quarter is already available - to make sure we capture mostly policy uncertainty rather than macroeconomic uncertainty. In this sense we are being slightly more conservative than the assumptions of our model could warrant.

⁵We 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 concerned with, is unaffected.

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

Table 1: Parameter Values

the economy might revert back to the passive monetary-policy regime.

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 largely flattens out. This series is particularly useful because it provides us with a way of calibrating $\bar{\mu}$.

Our claim that uncertainty around the conduct of monetary policy decreased in the 80s and 90s is also supported by a different measure of policy uncertainty, proposed by Baker, Bloom and Davis (2015). Figure 2 shows our SPF-based measure of uncertainty (in logs) clearly co-moves with the moving-average of the log of the Baker, Bloom and Davis (2015) index. 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 limit the sample to the 80s and 90s. Swanson (2006) provides convincing evidence that the improvements in private sector’s forecasts of the policy rate (of which the series in Figure 1 is a manifestation) is unlikely to stem from some kind of improved forecasting technology, since the forecast accuracy for other macro series did not follow a similar pattern. Rather, Swanson suggests it follows from improved central bank communication.

2.2 Model’s predictions

We now turn to showing how our simple, and admittedly stylized, model can reconcile these facts *without resorting to changes in the inflation target or of inflation-responsiveness coefficient*. We think this is important because, as Coibion and Gorodnichenko (2011) showed, it is hard to identify shifts in the monetary policy-maker’s responsiveness to inflation before and after Volcker, even when estimating a much richer model specification as they do. Moreover, as Del Negro, Giannoni and Schorfheide (2015) concede, assuming an exogenous change in the target is somewhat of a shortcut. Models in which the target varies exogenously seem at odds with evidence from the *Blue Book*, a document about monetary policy alternatives produced by Fed staff before each FOMC meeting. These scenarios are based on inflation targets (and earlier on money growth rates) that seem remarkably constant over time. Also, the Federal Reserve officially did not have a target value for inflation until January 2012.

Table 1 reports the calibration of the model we are using in this analysis. For the parameter that governs the inflation responsiveness of the Central Bank, ϕ , we assume the commonly accepted value of 1.5, as in Taylor (1993). Normally this number is assumed to apply only to the post-Volcker era (e.g. Clarida, Gali and Gertler, 2000), while the Taylor principle is not satisfied in the pre-Volcker era. But




Figure 3: Level of annualized trend inflation implied by our measure of forecasters' disagreement, $\phi = 1.5$ and $\pi^* = 1.5\text{pc}$ (lower) or $\pi^* = 2\text{pc}$ (upper).

we can show this need not necessarily be the case when ambiguity is accounted for in the way we propose. The curvature in the preferences for leisure ψ is calibrated to match a unitary Frisch elasticity, while θ and ϵ are calibrated to roughly match the timing at which the equilibrium switched from being indeterminate to determinate. In this sense we take comfort from the fact that a markup of the order of 7 percent and an average price duration of 6 quarters are consistent with our story, since they are both well within the accepted range. We base our assumption that $\pi^* \approx 1.5 - 2$ on the fact that in recent years *Blue Book* scenarios are built on a 1.5 and 2 percent inflation targets.

Given our assumptions on the inflation target and the responsiveness to inflation, Knightian uncertainty of the order of 100bp delivers trend inflation of around 3.5-4 percent, *i.e.* the level of trend inflation Ascari and Sbordone (2014) estimate for the early 1980s. We calibrated the uncertainty using the measure of professional forecasters' disagreement about short rates discussed above. The degree of disagreement of the order of 2 percent in 1981 and 1982 implies a calibration for $\bar{\mu}$ half that, that is 100bp. By the same token, trend inflation of around 2 percent is consistent with our assumptions when uncertainty is of the order of about 20bp, which is what our measure of expectations' disagreement implies in the later part of our sample. Figure 3 shows in a more systematic way, the level of trend inflation implied by our model, given the level of uncertainty about the interest rate implied by the SPF forecasts.

High Knightian uncertainty can easily imply indeterminacy, despite a time-invariant greater-than-one coefficient on inflation in the Taylor rule. Our calibration implies that the equilibrium is determinate for values of trend inflation that exceed the target by at most 2 percent. Based on our maintained assumptions that be consistent with a level of trend inflation of the order of 3-3.5 percent, a value that

the estimate of trend inflation in Ascari and Sbordone (2014) breaches in the early 80s and never reaches again. So, we confirm the finding in Coibion and Gorodnichenko (2011) that **the equilibrium being indeterminate in the 70s and determinate thereafter, despite no change in the inflation responsiveness coefficient**, which always satisfies the Taylor principle.

3 An intuitive example

Moving now to our model setup, we start off by building intuition in the simplest framework we can think of.

Consider a genuinely linear model (e.g. Cochrane (2011)), namely one in which the private sector's behavior can be summarized by a simple Fisher equation, while the Central Bank sets the policy rate according to a linear rule:

$$\begin{aligned}\mathbb{E}_t(i_t - r - \pi_{t+1}) &= 0 \\ i_t &= r + \phi\pi_t \quad \phi > 1,\end{aligned}$$

where r is the constant real rate, π_t is the rate of inflation and i_t is the nominal interest rate. We abstract from any standard exogenous shock to make the exposition as stark as possible.

The problem above can be summarized in the following single dynamic expectation equation:

$$\phi\pi_t = \mathbb{E}_t\pi_{t+1},$$

whose only non-explosive solution is $\pi_t = 0$.

Suppose now the private sector is uncertain about the laws of motion describing the economy. In particular, we consider a situation in which the private sector entertains alternative priors for the law of motion of the nominal rate. It could be because the private sector is uncertain about the Central Bank's estimate of the real rate or because the Central Bank's objective is not perfectly known to the private sector or because agents fear that policymakers might deviate from their proposed rule. We also assume that the agents are not able to give probability weighting of the different alternative priors and, therefore, face ambiguity. The private sector's expectation for the policy rate becomes⁶:

$$i_t = \mathbb{E}_t^\mu [r + \phi\pi_t] = r + \phi\pi_t + \mu,$$

where we use the following definition to capture the effects of ambiguity, or Knightian uncertainty, on the mean:

$$\mathbb{E}_t^\mu x_{t+j} \equiv \mathbb{E}_t x_{t+j} + \mu \quad \forall j \geq 0 \tag{2}$$

\mathbb{E}_t being the mathematical expectation operator. Basically we are considering situations in which the private sector acts as if the mean policy rate might persistently deviate from the level implied by our simple Taylor rule.

⁶Implicit is also a timing assumption we will spell out in greater detail below but basically amounts to saying that expectations for the level of the interest rate in the current quarter need be computed (see Christiano, Eichenbaum and Evans (2005) for a similar approach).

As a result, the economy can now be characterized by the following equation instead:

$$\phi\pi_t = \mathbb{E}_t\pi_{t+1} - \mu \quad (3)$$

which admits the following solution:

$$\pi_t = -\frac{1}{\phi - 1}\mu$$

We pin down the value of μ , by assuming that the agents are ambiguity-averse⁷. Gilboa and Schmedler (1989) and Epstein and Schneider (2003) show, in a static and dynamic context respectively, that, if agents are ambiguity-averse, they will act as if the worst-case scenario, *i.e.* the one that minimizes their welfare over the plausible set, will materialize.

In this simple example, we need to endow the private sector with an *ad hoc* welfare function. A natural candidate is a concave function \mathcal{W} that penalizes deviations of inflation⁸. Then, the agents will form expectations according to the level of μ that solves:

$$\min_{\mu \in [-\bar{\mu}, \bar{\mu}]} \mathcal{W}(\pi_t) \quad (4)$$

Strict concavity implies the only the two boundaries are candidate minima. So, for any non-zero level of ambiguity ($\bar{\mu} > 0$), inflation will persistently deviate from its welfare-optimal level delivering a first-order effect of uncertainty on all the variables in our economy.

This poses an interesting challenge for policy. Suppose that $\mathcal{W}'(0) = 0$, *i.e.* the optimal level of inflation is zero⁹. This corresponds to a situation in which the central bank is following optimal monetary policy, provided the private sector trusts this is indeed the case. In spite of this, inflation will turn out to deviate from this level, generating a strong incentive for policymakers to engage in effective communication to increase the private sector confidence in their policy behavior.

Having set the stage in section, in the next one we build a standard simple New-Keynesian DSGE model that illustrates how this intuition carries over in a general equilibrium model with pricing frictions. Moreover, the micro-founded structure of the economy will allow us to characterize agents' welfare from first principle. In particular, it turns out that welfare is indeed a strictly concave function of inflation around its target. Indeed the loss function is also asymmetric, the worst case corresponding to a situation in which inflation is inefficiently high.

4 The Model

We modify a textbook New-Keynesian model (Galí, 2008) by assuming that the agents face ambiguity about the expected future policy rate. Absent ambiguity,

⁷See for example Dimmock *et al.* (2014) for evidence that the majority of households are ambiguity averse.

⁸We are simply imposing $\mathcal{W}''(\pi_t) < 0$, but note that the widely-used quadratic loss function $-\pi_t^2$ is a special case.

⁹Possibly in deviation from the target. This is not restrictive as we will show later.

the first-best allocation is attained thanks to a sufficiently strong response of the Central Bank to inflation, *i.e.* for $\phi > 1$ 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 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.

4.1 Households

Let $s_t \in S$ be the vector of exogenous states. We use $s^t = (s_1, \dots, s_t)$ to denote the history of the states up to date t . A consumption plan \vec{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:

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

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

$$U_t(\vec{C}; s^t) = \min_{p \in \mathcal{P}_t(s^t)} \mathbb{E}^p u(\vec{C}_t) + \beta U_{t+1}(\vec{C}; s_t, s_{t+1}) \quad (5)$$

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 (5) 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)¹⁰. 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.

As discussed in more detail below, we parametrise the belief set with an interval $[-\bar{\mu}_t, \bar{\mu}_t]$ 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:

$$U_t(\vec{C}; s^t) = \min_{\mu \in [-\bar{\mu}_t, \bar{\mu}_t]} \mathbb{E}^\mu u(\vec{C}_t) + \beta U_{t+1}(\vec{C}; s_t, s_{t+1}) \quad (6)$$

The households' budget constraint is:

$$P_t C_t + B_t = R_{t-1} B_{t-1} + W_t N_t + T_t \quad (7)$$

¹⁰More details and axiomatic foundations for such preferences are in Epstein and Schneider (2003).

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 households in this setup is their perceived uncertainty about the return to their savings R_t . As we describe in more detail in the Subsection 4.3, 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:

$$\frac{1}{C_t} = \mathbb{E}_t^\mu \left[\frac{\beta R_t}{C_{t+1} \Pi_{t+1}} \right] \quad (8)$$

$$N_t^\psi C_t = \frac{W_t}{P_t} \quad (9)$$

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, once more, 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. Hence the intertemporal Euler equation becomes:

$$\frac{1}{C_t} = \mathbb{E}_t \left[\frac{\beta \tilde{R}_t}{C_{t+1} \Pi_{t+1}} \right] \quad (10)$$

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

4.2 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

$$Y_t = \left[\int_0^1 Y_t(i)^{\frac{\varepsilon-1}{\varepsilon}} di \right]^{\frac{\varepsilon}{\varepsilon-1}}. \quad (11)$$

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

$$Y_t(i) = \left(\frac{P_t(i)}{P_t} \right)^{-\varepsilon} Y_t \quad (12)$$

and in the aggregate price index

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

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

$$Y_t(i) = A_t N_t(i), \quad (13)$$

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-optimize their price at any given period, while the others must keep the nominal price unchanged¹¹. Whenever a firm can re-optimize, it sets its price maximising the expected present discounted value of future profits

$$\max_{P_t^*} E_t \left[\sum_{s=0}^{\infty} \theta^s Q_{t+s} \left(\left(\frac{P_t^*(i)}{P_{t+s}} \right)^{1-\varepsilon} Y_{t+s} - \Psi \left(\left(\frac{P_t^*(i)}{P_{t+s}} \right)^{-\varepsilon} Y_{t+s} \right) \right) \right] \quad (14)$$

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 is characterised by the following first-order condition:

$$\frac{P_t^*(i)}{P_t} = \frac{\mathbb{E}_t \sum_{j=0}^{\infty} \theta^j Q_{t+j} \left(\frac{P_{t+j}}{P_t} \right)^{\frac{\varepsilon}{\varepsilon-1}} MC_{t+j}}{\mathbb{E}_t \sum_{j=0}^{\infty} \theta^j Q_{t+j} \left(\frac{P_{t+j}}{P_t} \right)^{\varepsilon-1}},$$

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

$$\frac{P_t^*(i)}{P_t} = \left(\frac{1 - \theta \Pi_t^{\varepsilon-1}}{1 - \theta} \right)^{\frac{1}{1-\varepsilon}}, \quad (15)$$

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

4.3 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

¹¹Or indexed to the inflation when we consider it to be non-zero.

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:

$$R_t = R_t^n (\Pi_t)^\phi, \quad (16)$$

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 rate¹².

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.* $\bar{\mu}_t = 0$:

$$\tilde{R}_t \equiv R_t e^{\mu_t} = R_t = R_t^n (\Pi_t)^\phi$$

which is the nominal rate that implements first-best allocations (together with the subsidy). And, clearly, there is no uncertainty in the standard sense of the word around the expected value, which is then a risk-free rate.

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:

$$\tilde{R}_t = R_t^n (\Pi_t)^\phi e^{\mu_t} \quad (17)$$

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 (16) 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.

¹²While 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 uncertainty about it.

4.4 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 (11), then it follows that

$$Y_t = C_t.$$

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 (13) and then use the demand function (12) to obtain the last equality. Let us define $\Delta_t \equiv \int_0^1 \left(\frac{P_t(i)}{P_t} \right)^{-\varepsilon} di$ 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.

5 The Worst-Case Steady State

Following Ilut and Schneider (2014), we study our model economy log-linearized 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.

5.1 General Properties of the Steady States

5.1.1 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 5.1. *In a steady state with no real growth, inflation depends on the ambiguity distortion parameter as follows:*

$$\Pi(\mu, \cdot) = e^{-\frac{\mu}{\phi-1}}, \quad (18)$$

while the policy rate is:

$$R(\mu, \cdot) = \frac{1}{\beta} e^{-\frac{\phi\mu}{\phi-1}}. \quad (19)$$

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

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

and the opposite for $\mu < 0$.

Proof. Proof in Appendix A. □

Result 5.1 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 ($\mu < 0$)¹³. 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 state¹⁴ $\frac{1}{\beta}$.

5.1.2 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:

$$\Delta(\mu, \cdot) = \frac{(1 - \theta) \left(\frac{1 - \theta \Pi(\mu, \cdot)^{\epsilon-1}}{1 - \theta} \right)^{\frac{\epsilon}{\epsilon-1}}}{1 - \theta \Pi(\mu, \cdot)^{\epsilon}} \quad (20)$$

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 alters reduces labour productivity and ultimately welfare.

¹³In our analysis we do not consider the zero lower bound, because in case it was binding that could not be a steady state. If, however, one wanted to explicitly account for that, a restriction on range of μ would simply take the following form: $\mu < -\frac{\phi-1}{\phi} \log(\beta)$.

¹⁴Here we assume that, absent distortions, inflation would be zero in steady state, but, as discussed previously, all results follow through with non-zero steady state inflation, as long as the firms index to that value of inflation.

5.1.3 Hours, Consumption and Welfare

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

$$N(\mu, \cdot) = \left(\frac{(1 - \theta\Pi(\mu, \cdot)^{\epsilon-1})(1 - \beta\theta\Pi(\mu, \cdot)^\epsilon)}{(1 - \beta\theta\Pi(\mu, \cdot)^{\epsilon-1})(1 - \theta\Pi(\mu, \cdot)^\epsilon)} \right)^{\frac{1}{1+\psi}}, \quad (21)$$

while consumption is:

$$C(\mu, \cdot) = \frac{A}{\Delta(\mu, \cdot)} N(\mu, \cdot) \quad (22)$$

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

$$\mathbb{V}(\mu, \cdot) = \frac{1}{1-\beta} \left(\log(C(\mu, \cdot)) - \frac{N(\mu, \cdot)^{1+\psi}}{1+\psi} \right). \quad (23)$$

Finally note that equation (21) 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 finite¹⁵.

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:

$$\mu > \frac{\phi - 1}{\epsilon} \log(\theta), \quad (24)$$

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 1 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.

5.2 Characterising the Worst-Case Steady State

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. This is a key difference with respect to Ilut and Schneider (2014), who assume that agents are ambiguous about the *exogenous* TFP process. In their paper follows quite naturally that the worst-case steady state is one in which agents under-estimate TFP growth.

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.

¹⁵Indeed, the same condition could be derived from the formula for price dispersion in equation (20).




Figure 4: Steady-state welfare as a function of μ (measured in annualized percentage points).

Result 5.2. 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:

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

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

Proof. Proof in Appendix A. □

Result 5.2 illustrates that the welfare function is locally concave around the first-best (see Figure 4, 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 5.2 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 to see from Figure 4 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.

Result 5.3. For β sufficiently close but below 1 and all the other parameters in the intervals defined in Result 5.2, $\mu = -\bar{\mu}$ minimizes $\mathbb{V}(\mu, \cdot)$ over $[-\bar{\mu}, \bar{\mu}]$, for any sufficiently small $\bar{\mu} > 0$.

Proof. Proof in Appendix A. □

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




Figure 5: Steady state marginal cost as a function of μ (measured in annualized percentage points).

on welfare is not. In particular, positive steady-state inflation - associated with negative levels of μ as shown in Result 5.1 - 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-optimize. These firms will face a very high demand, which in turn will push up their labour demand, and ultimately their marginal costs, as Figure 5 shows. On the other hand, negative inflation will reduce the demand for firms which do not re-optimize and this will reduce their demand for labour and their marginal costs¹⁶. 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 5.1 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, 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

¹⁶Once more, this shows that first best is attained in the absence of ambiguity, when the marginal costs equals the inverse of the markup.

Result summarises Results 5.1 and 5.3 and establishes more formally the effects of ambiguity on the worst-case steady state levels of inflation and the policy rate.

Result 5.4. *For β sufficiently close but below one and all the other parameters in the intervals defined in Result 5.2, for any small enough $\bar{\mu} > 0$:*

$$\mathbb{V}_w(\bar{\mu}') > \mathbb{V}_w(\bar{\mu}) \quad \Pi_w(\bar{\mu}') < \Pi_w(\bar{\mu}) \quad R_w(\bar{\mu}') < R_w(\bar{\mu}) \quad \forall 0 \leq \bar{\mu}' < \bar{\mu} \quad (25)$$

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

Proof. Proof in Appendix A. □

5.3 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 5.5. *While parameter values are in the intervals defined in Result 5.2 and $\bar{\mu}$ is a small positive number, given any pair $(\mu, \phi) \in [-\bar{\mu}, 0) \times (1, \infty)$, for any $\mu' \in [-\bar{\mu}, 0)$ 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 $\mu \in (0, \bar{\mu}]$.

Proof. Proof in Appendix A. □

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 6 illustrates Result 5.5 graphically for 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 steeper¹⁷ (as the orange dashed line illustrates)

¹⁷Note 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.




Figure 6: Welfare function for $\phi = 1.5$ (solid red line) and $\phi = 1.4$ (orange dashed line).

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 5.5 shows graphically that the same level of welfare can be attained when $\phi = 1.4$ but only if uncertainty is smaller (of the order of 80bp).

6 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 B presents the log-linear approximation around a generic steady state indexed by μ . By setting $\mu = -\bar{\mu}$, 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 $-\bar{\mu}$ 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 Knightian uncertainty on the determinacy region and we then study the effects of shocks to ambiguity.

6.1 Log-Linear solution

Appendix B.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:

$$\tilde{y}_t = \mathbb{E}_t \tilde{y}_{t+1} - (\phi \pi_t + \hat{\mu}_t - \mathbb{E}_t \pi_{t+1}) \quad (26)$$

$$\pi_t = \kappa_0(-\bar{\mu}, \cdot) \left((1 + \psi) \tilde{y}_t + \psi \hat{\Delta}_t \right) + \kappa_1(-\bar{\mu}, \cdot) \mathbb{E}_t \widehat{F^2}_{t+1} + \kappa_2(-\bar{\mu}, \cdot) \mathbb{E}_t \pi_{t+1} \quad (27)$$

$$\hat{\Delta}_t = \kappa_3(-\bar{\mu}, \cdot) \hat{\Delta}_{t-1} + \kappa_4(-\bar{\mu}, \cdot) \pi_t \quad (28)$$

$$\widehat{F^2}_t = \mathbb{E}_t \left(\kappa_5(-\bar{\mu}, \cdot) \pi_{t+1} + \kappa_6(-\bar{\mu}, \cdot) \widehat{F^2}_{t+1} \right). \quad (29)$$

where $\tilde{y}_t = c_t - a_t$ is the deviation of the output gap from its worst-case steady-state level, $\hat{\Delta}_t$ is the price dispersion, while $\widehat{F^2}_t$ can be interpreted as the present discounted sum of future expected inflation rates in the recursive formulation of the optimal price-setting equation¹⁸. 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 $-\bar{\mu}$. In particular, $\hat{\mu}_t > 0$ implies $\mu_t > -\bar{\mu}$. As a result, $\hat{\mu}_t$ cannot take on negative values since $\hat{\mu}_t = 0$ corresponds to the lower bound of the interval $[-\bar{\mu}, \bar{\mu}]$. 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 (26).

To verify our conjecture that the value function is increasing around $-\bar{\mu}$ 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

$\{a_t, \hat{\Delta}_{t-1}, \hat{\mu}_t\}$ into $\{\tilde{y}_t, \pi_t, \hat{\Delta}_t, \widehat{F^2}_t\}$. 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 $\mu_t = -\bar{\mu}$.
- iv. We will finally verify our conjecture numerically for our preferred calibration, in which we maintain a unitary level of the Frisch elasticity.

6.1.1 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

¹⁸Indeed it corresponds to the log-linearized version of the denominator of the expression on the RHS of equation (15).

are:

$$\pi_t = -\frac{1}{\phi-1}\hat{\mu}_t \quad (30)$$

$$\widehat{F}2_t = -\frac{\kappa_5}{(\phi-1)(1-\kappa_6)}\hat{\mu}_t \quad (31)$$

$$\tilde{y}_t = -\frac{\kappa_1\lambda_F + (\kappa_2-1)\lambda_\pi}{\kappa_0}\hat{\mu}_t \quad (32)$$

$$\hat{\Delta}_t = \kappa_3\hat{\Delta}_{t-1} - \frac{\kappa_4}{\phi-1}\hat{\mu}_t \quad (33)$$

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.

$\widehat{F}2_t$ 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 ($\Delta(-\bar{\mu}, \cdot)$) 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:

$$c_t = \tilde{y}_t + a_t = \lambda_Y\hat{\mu}_t + a_t \quad (34)$$

$$n_t = \tilde{y}_t + \hat{\Delta}_t = \lambda_Y\hat{\mu}_t + \hat{\Delta}_t = \kappa_3\hat{\Delta}_{t-1} + \left(\lambda_Y - \frac{\kappa_4}{\phi-1}\right)\hat{\mu}_t \quad (35)$$

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:

$$\mathbb{V}(a_t, \hat{\Delta}_{t-1}; \hat{\mu}_t) = \log((1 + c_t)C(-\bar{\mu}, \cdot)) - (1 + n_t)N(-\bar{\mu}, \cdot) + \beta \mathbb{E}_t \mathbb{V}(a_{t+1}, \hat{\Delta}_t; \hat{\mu}_t) \quad (36)$$

simply plugging in equations (34) and (35) for c_t and n_t .

To build intuition, consider the linear approximation of the log function around one, $\log(1 + \lambda_Y \hat{\mu}_t + a_t) \simeq \lambda_Y \hat{\mu}_t + a_t$, and the fact that for any reasonable calibration values of hours are indeed very close to unity¹⁹, $N(-\bar{\mu}, \cdot) \simeq 1$. 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 $\hat{\mu}_t = 0$ ²⁰ 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 (36) can be re-written as:

$$\mathbb{V}(a_t, \hat{\Delta}_{t-1}; \hat{\mu}_t) = u(\vec{C}(-\bar{\mu}, \cdot)) + \log(1 + c_t) - n_t N(-\bar{\mu}, \cdot) + \beta \mathbb{E}_t \mathbb{V}(a_{t+1}, \hat{\Delta}_t; \hat{\mu}_t)$$

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

$$d\mathbb{V}(a_t, \hat{\Delta}_{t-1}; \hat{\mu}_t) \equiv \mathbb{V}(a_t, \hat{\Delta}_{t-1}; \hat{\mu}_t) - \mathbb{V}(-\bar{\mu}, \cdot) \quad (37)$$

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:

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

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 (34) and (35) 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 marginal utility of consumption. In

¹⁹For our preferred calibration we obtain 1.0025.

²⁰Remember that $\hat{\mu}_t$ cannot go negative as described above.

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 $\mu_t = -\bar{\mu}$.

Result 6.1. *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 $\hat{\mu}_t = 0 \forall (a_t, \hat{\Delta}_{t-1}) \in [\underline{a}, \bar{a}] \times \mathbb{R}$:*

- i. if $\frac{\lambda_Y}{(1+\bar{a})(1-\beta)} \geq \Xi(-\bar{\mu}, \cdot)$, when $\lambda_Y > 0$
- ii. always, when $\lambda_Y = 0$
- iii. if $\frac{\lambda_Y}{(1+\bar{a})(1-\beta)} \geq \Xi(-\bar{\mu}, \cdot)$, when $\lambda_Y < 0$

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

Proof. See Appendix B.3 □

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, $\Xi(-\bar{\mu}, \cdot) < 0$ for any value of $\bar{\mu}$ so, in reality, the only potentially binding condition is the third. Moreover, if we take the calibration in Table 1 and simply set $\psi = 0$, we find that λ_Y is positive for any degree of uncertainty 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 $\bar{\mu}$ our sufficient conditions are trivially met. Yet, to show the robustness of the result, let us suppose uncertainty 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.

6.1.2 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 Frisch 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:

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

To derive our sufficient conditions (see Appendix B.4 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 negative²¹.
- ii. the derivative is decreasing in the level of technology²²
- 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.

6.2 Uncertainty and Equilibrium Determinacy

It is well known that the Taylor principle (nominal rates moving 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). Ascari and Ropele (2009), however, show that this is not necessarily the case when the model is approximated around a positive level of steady-state inflation.

In our setting, steady-state inflation emerges as the result of model uncertainty and responsiveness to inflation on the side of the Central Bank, which allows us to explore another *margin*. Consistent with the analysis in Ascari and Ropele (2009), a higher ϕ tends to deliver determinacy for any given level of inflation. Here however, a higher ϕ also acts to reduce the level of steady-state inflation for a given degree of uncertainty, which in turns tends to require a lower level of ϕ for determinacy to be attained. In other words, the role of the coefficient governing the response of the interest rate to inflation is twofold thus making it much more powerful.

Figure 7 illustrates this point. A higher degree of ambiguity commands a swifter response to inflation for the equilibrium to be determinate. So we can think of situations where an economy transitions from a determinate to an indeterminate equilibrium, not because the Central Bank changed its response function, but because of a change in confidence on the part of the private sector.

To illustrate our main point, a black solid line in Figure 7 marks the $\phi = 1.5$ level. It is immediate to see that, given our calibration, degrees of ambiguity of the order of about 75bp or higher, which prevailed in the first part of the sample we consider, imply that the equilibrium is indeterminate.

The dashed line, on the other hand, represents the combinations of $\bar{\mu}$ and ϕ such that $\bar{\pi} - \pi^* = 1.5$, i.e. trend inflation is one-and-a-half percentage points above the target (see equation (1)). The crossing point of the two lines, represents the scenario in which uncertainty is of the order of 75bp, inflation is 1.5 percent above target (so around 3-3.5 percent given the assumptions we laid down in Section 2), while

²¹Zero in the limit case in which $\bar{m}u = 0$

²²Similar to the first of the three cases in the previous paragraph.




Figure 7: 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 $(\bar{\mu}, \phi)$ pairs consistent with annualized inflation 1.5 percent above target.

$\phi = 1.5$. In other words, it represents a situation consistent with what happened in the early 80s. And it also corresponds to the point the equilibrium switches from being indeterminate (for higher values of inflation give ϕ) to being determinate, which is one of the points we made in Section 2 and we now properly illustrate.

The bottom line is that, given a level of ϕ we normally we associate with determinate equilibria and a level of the inflation target of their order of 1.5 – 2 percent, our measure of expectations’ dispersion maps nicely into the idea that the equilibrium switched from indeterminate to determinate in the early 1980s.

6.3 Responses to Changes in Uncertainty

So far we have maintained the assumption that the set of probability models that the agents entertain as possible does not change over time. It is plausible, however, that the agents’ understanding and trust in the monetary policymakers behavior is affected by changes in the agents’ level of confidence, the policymakers’ communication strategy, or the general level of uncertainty. Such shocks to confidence will affect the range over which agents are uncertain. In light of our analysis in the previous section we can maintain:

$$\hat{\mu}_t = -\hat{\bar{\mu}}_t. \quad (39)$$




Figure 8: Responses of (left to right) Output, Annualized Inflation, Hours and Price Dispersion to a 25bp reduction in Knightian uncertainty

In other words, agents will continue to choose the lowest value in the interval to twist all future expectations, but now this value will fluctuate over time²³. For concreteness, let us assume an AR(1) process for $\widehat{\mu}_t$:

$$\widehat{\mu}_t = \rho \widehat{\mu}_{t-1} + \eta_t.$$

and we will maintain that $\rho = .95$ to capture the persistence of these shifts in confidence.

As the previous analysis demonstrated, the impulse responses depend on the steady-state level $\bar{\mu}$. In particular, we discussed above how the sign of the response of output/consumption is potentially sensitive to the size of $\bar{\mu}$ ²⁴. Given our baseline calibration, however, any $\bar{\mu} \geq 7bp$ will deliver a positive output response. 7bp is much smaller than even the smaller value suggested by our measure of expectations dispersions so we focus on numbers larger than that.

For concreteness we calibrate $\bar{\mu}$ to 50bp, which is between the high values around 100bp in the early part of the sample and lower values of the order of 25bp in the latter part.

Looking at Figure 8 it is immediate to notice that inflation falls by just shy of half a percent. If utility were linear in hours, the fall would be exactly $\frac{1}{\phi-1}\widehat{\mu}_t$ as equation (30) demonstrates. In our baseline calibration we are not far from that benchmark, although the marginal cost is affected by the curvature in the disutility from working. Output, as mentioned above, rises. The increase is less than .1 percent, which is small in absolute value but is still interesting for a couple of reasons. For one thing, in our model there is only a simple pricing friction. If, for instance, wages were stickier, the effect would be much larger. We can verify this by simply setting $\psi = 0$, which in and of itself makes wages respond less and increases the real effects of the shock. In our case, that would make the response of output double in size. More importantly, it is interesting to see that a disinflationary shock can drive output up. This is something we would associate with supply shocks.

²³In so doing we implicitly assume that the shocks to ambiguity are small enough so not to violate the sufficient condition derived in the previous section which relies, in the case of non-linear hours at least, on $\widehat{\Delta}_t = 0$. The fact that in the limiting case of linear hours this is not an issue suggests that this is a safe assumption to make.

²⁴See equation (32) for the linear-hours case.

Here the reason for this is primarily the expectations formation process and the fact that higher confidence leads to a smaller wedge in the expectation of *all* future nominal wages. In the next paragraph, we will explore this aspect more in depth, comparing what would happen if the same disinflation was to be achieved via an inflation target shock. Here we turn to analyzing the response of hours. A confidence shock makes agents better off so they tend to want to consume more leisure. That pushes down on hours worked, which is a well known effect in this class of models in response to technology shocks. Here, however, the increased productivity comes from a reduction in price dispersion which builds up over time (as the impulse response in the right-most pane of Figure 8 illustrates). Hence, in a demand-driven economy, hours initially increase to meet the increased demand for consumption and only later fall when the economy becomes more efficient and can support the increased consumption demand with a smaller amount of hours worked.

The bottom line is that a reduction, albeit temporary, in uncertainty will cause inflation to fall from its inefficiently high level. As a consequence, the economy will experience an increase in consumption and an increase in welfare.

In this respect, it is important notice that the welfare effect of confidence shocks is different in different phases of the business cycle. In the linear-hours case this is particularly stark and goes back to the discussion we presented above. The marginal utility of leisure is constant so the phase of the business cycle does not matter. However, the marginal utility of consumption is decreasing in the level of a_t (see equation (38)) so the boost in utility from the increase in consumption is greater during a downturn, i.e. when a_t and hence consumption are relatively low.

6.4 Inflation Target Shock

There is a literature trying to reconcile low-frequency movements in inflation and policy rates with inflation target shocks (e.g. Del Negro and Eusepi (2011)). Typically shocks to a target level of inflation are assumed to be extremely persistent and considered a stand-in for all kinds of changes in the conduct of monetary policy.

Our discussion in Section 2 shows how changes in Knightian uncertainty can explain the evolution in trend inflation in the 80s and 90s without necessarily resorting to changes in the target rate of inflation, which was announced by the Fed only in January 2012. In this section we illustrate how the dynamic responses to inflation target shocks differ from those to a confidence shock. The very fact that they might differ could be surprising in and of itself, because the shocks enter the dynamic system in a very similar fashion. In particular, they both enter the Euler equation which, in their presence becomes:

$$\tilde{y}_t = \mathbb{E}_t \tilde{y}_{t+1} - \left(\phi (\pi_t - \pi_t^*) - \widehat{\mu}_t - \mathbb{E}_t \pi_{t+1} \right) \quad (40)$$

It would appear as if a reduction in the target was isomorphic to a reduction in ambiguity (appropriately scaled by ϕ). It would indeed be so if it wasn't because of the peculiar properties of the shock to confidence.

The confidence shock is a shock to expectations, *i.e.* it changes expectations computed at time t for all future periods:

$$\mathbb{E}_t^\mu x_{t+j} = \mathbb{E}_t x_{t+j} + \widehat{\mu}_t \quad \forall j \geq 0$$




Figure 9: Responses of (left to right) Output, Annualized Inflation, Hours and Price Dispersion to a 25bp reduction in Knightian uncertainty (solid red lines) or to a reduction in the target producing the same effect on inflation (dashed blue lines).

which breaks down the law of iterated expectations. To see this consider the twisted expectation for the generic variable x_{t+2} :

$$\mathbb{E}_t^\mu x_{t+2} = \mathbb{E}_t x_{t+2} + \widehat{\mu}_t \quad (41)$$

$$\mathbb{E}_t (\mathbb{E}_{t+1}^\mu x_{t+2}) = \mathbb{E}_t (\mathbb{E}_{t+1} x_{t+2} + \widehat{\mu}_{t+1}) = \mathbb{E}_t x_{t+2} + \rho \widehat{\mu}_t \quad (42)$$

For example, let us consider the same experiment behind the impulse responses in Figure 8, *i.e.* a reduction in uncertainty to 25bp, down from a steady-state level of 50bp. As an agent, I act as if I knew that policymakers could only deviate by at most 25bp from their rule. Because of worst-case considerations I will then compute the expected policy rate implied by the *rational* expectations for future levels of inflation and the natural rate and then subtract 25bp.

The fact that my confidence level tomorrow will be different (because uncertainty will have started to revert back to steady state at a rate governed by ρ) is quite irrelevant because I need to compute the expectations of *all* future rates *today* to determine my consumption. In other words, the confidence level that matters is that which prevails when expectations are formed, not when the variable is realized²⁵. On the other hand, if model agents are asked what they expect their assessment of the situation in time $t+2$ will be next period (equation (42)), they will correctly recognize that, come period $t+1$, their confidence level will have started to revert back to steady state, hence they will twist the expectations by $\rho \widehat{\mu}_t$.

In light of this discussion, it is not surprising, that the responses to inflation-target and confidence shocks are equivalent when they are both random walk. In that case, in expectations, the current degree of uncertainty would be the same as that prevailing at any future date, so the timing distinction would be irrelevant.

When shocks are stationary, however, the difference emerges first and foremost in consumption, which is the variable for which expectations of future interest rates

²⁵Nothing would prevent us from assuming that in computing expectations for time $t+j$ variables agents use the distortion prevailing at time $t+j$. In fact, if one wants to verify what the effects of that alternative assumption would be, this section would provide the answer, because the responses would be identical to those produced by an inflation-target shock (up to scaling by ϕ). Yet, we find our assumption more compelling because confidence affects expectations directly rather than realizations, so the time at which expectations are computed is critical.

matter. The effect on inflation, on the other hand, is the same up to scaling. In other words, it is always possible to find an inflation-target shock that delivers the same profile for inflation as a given confidence shock.

The interesting aspect is the different response of consumption (and consequently of hours). Figure 9 gives us a clear representation of the difference by reporting the responses to the same confidence shock as Figure 8 alongside the responses induced by the "equivalent" inflation-target shock. The inflation-target shock generates the standard pattern of a contractionary monetary-policy shock²⁶: a fall in inflation, consumption and hours.

The difference is how the gain in efficiency (the fall in price dispersion is the same) is divided up into consumption and leisure. The "expectations" effect induces people to bring consumption forward, which, in turn, generates an increase in hours worked. The ability to substitute in and out of leisure is crucial for this difference to emerge. Indeed as $\psi \rightarrow \infty$ the difference in the responses of consumption and hours to the shocks vanes.

At the end of this analysis we can then derive a simple but sharp policy implication. If a fall in confidence should drive inflation up, neutralizing the impact on inflation by means of policy tightening will create a relatively big recession because of the combined effects on output of the fall in confidence and of the monetary tightening. Restoring confidence, to the extent possible, is definitely the option to be preferred.

7 Optimal Monetary Policy

So far, we have assumed that policymakers follow a rule that *would be* optimal in the absence of Knightian uncertainty (equation (16)). 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 $\phi \rightarrow \infty$. 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 $\bar{\phi}$.

²⁶A contractionary monetary policy shock is again equivalent up to scaling and degree of autocorrelation.

7.1 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 (2011) who derive an approximation suitable for an environment featuring trend inflation and obtain the following:

$$\mathcal{L}_t = \sum_{j=0}^{\infty} \beta^j (\Theta_0 + \Theta_1 \text{var}(\tilde{y}_t) + \Theta_2 \text{var}(\pi_t)) \quad (43)$$

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 uncertainty. This not really surprising since Coibion, Gorodnichenko and Wieland (2011) 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 higher the degree of uncertainty in the economy 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.

7.2 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. For our main analytical result we also restrict ourselves to the the economy described in Section 4, i.e. an economy with only TFP shocks. We will address trade-off inducing shocks further down.

Proposition 7.1. *Given the economy described in Section 4, a small $\bar{\mu} > 0$ and restricting $\underline{\phi}(-\bar{\mu}, \cdot) < \phi \leq \bar{\phi}$, the following rule is statically and dynamically optimal in its class:*

$$R_t = R_t^* \Pi_t^{\bar{\phi}} \quad (44)$$

where

$$R_t^* = R_t^n e^{\delta^*(\bar{\mu}, \bar{\phi}, \cdot)} \quad (45)$$

and

$$0 < \delta^*(\bar{\mu}, \bar{\phi}; \cdot) < \bar{\mu} \quad (46)$$

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

Proof. See Appendix B.5 □

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 $\phi = \bar{\phi}$ is optimal should not be surprising at this point²⁷. 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 *naive* policymaker who realizes the private sector is systematically underestimating its policy rate by $\bar{\mu}$. Its response could amount to systematically setting rates higher than its standard Taylor rule would predict by the same amount $\bar{\mu}$. If these were just parameters and there was no minimization involved, this policy action would implement first best²⁸. This would be naive because, so long as there is some uncertainty lingering the first-best outcome will never be attained. Or, in other words, if $\delta = \bar{\mu}$ 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 $\delta = \bar{\mu}$ 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.

²⁷Note that lower bound $\underline{\phi}(-\bar{\mu}, \cdot)$ is simply meant to capture the lowest degree of responsiveness to inflation that ensures determinacy for a given degree of steady-state ambiguity.

²⁸That is because in this economy steady state inflation is $\Pi(-\bar{\mu}, \delta(\cdot), \cdot) = e^{-\frac{-\bar{\mu} + \delta(\cdot)}{\phi - 1}}$ so by setting $\delta = \bar{\mu}$, steady state inflation would be zero, if μ did not respond to that.

Another aspect of Proposition 7.1 is that the rule in equation (44) 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 7.1 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 7.1. *Given any constrained-optimal monetary policy plan, a monetary policy rule with the same functional form as that in Proposition 7.1 can be made welfare equivalent for a suitably high level of $\bar{\phi}$.*

Proof. See appendix B.6 □

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 dynamic optimality would require a numerical exercise, but can definitely draw some general conclusions about the efficiency of the steady state. In particular, as discussed above, 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 (43) is decreasing in trend inflation while that on output gap variability is constant. So, for a given level of uncertainty $\bar{\mu}$, we might end up in the paradoxical situation in which increasing ϕ all the way up to $\bar{\phi}$, 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 (43) 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 $\phi = \bar{\phi}$ is the optimal thing to do under most any circumstance, which is also in line with thorough 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 uncertainty 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 vanes, 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.

8 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 uncertainty 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 uncertainty about monetary policy perceived by the private sector.

Based on a calibration of uncertainty that matches the interdecile dispersion of the SPF forecasts of the current quarter's TBill rates, our model can explain the disinflation of the 80s and 90s as resulting from an increase in the private sector confidence in their understanding of monetary policy, rather than from changes in target inflation. And 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.

To get a better insight in the workings of confidence shocks of this type, we illustrate how macro variables respond differently to a shock to uncertainty than they do to a standard inflation target shock. In particular we show that the response of output tends to be positive, even in the face of a shock that reduces inflation, because of the peculiar way agents form expectations about the future.

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 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] Adam, K. and M. Woodford, 2012. "Robustly optimal monetary policy in a microfounded New Keynesian model," *Journal of Monetary Economics*, Elsevier, vol. 59(5), pp. 468-487.
- [2] Ascari, G. and T. Ropele, 2007. "Optimal monetary policy under low trend inflation," *Journal of Monetary Economics*, Elsevier, vol. 54(8), 2568-2583.
- [3] Ascari, G. and T. Ropele, 2009. "Trend Inflation, Taylor Principle, and Indeterminacy," *Journal of Money, Credit and Banking*, vol. 41(8), 1557-1584.
- [4] Ascari, G. and A. Sbordone, 2014. "The Macroeconomics of Trend Inflation," *Journal of Economic Literature*, vol. 52(3), 679-739.
- [5] Baker, S. R., N. Bloom, and S. J. Davis, 2015. "Measuring economic policy uncertainty," <http://www.policyuncertainty.com/media/BakerBloomDavis.pdf> .
- [6] Bean, C., 2013. "Global aspects of unconventional monetary policies," Jackson Hole Economic Policy Symposium, (August).
- [7] Benigno, P. and L. Paciello, 2014. "Monetary policy, doubts and asset prices," *Journal of Monetary Economics*, 64(2014), 85-98. Bidder, R.M. and M.E. Smith, 2012. "Robust animal spirits," *Journal of Monetary Economics*, Elsevier, vol. 59(8), pp. 738-750.
- [8] Bianchi, F. and C. Ilut, 2015. "Monetary/Fiscal Policy Mix and Agents' Beliefs," mimeo, Duke University.
- [9] Calvo, G. A., 1983. "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics*, Elsevier, vol. 12(3), pages 383-398, September.
- [10] Christiano, L.J, M. Eichenbaum and C.L. Evans, 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, University of Chicago Press, vol. 113(1), pp. 1-45.
- [11] Clarida, R., Gali, J. and Mark Gertler, 2000. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *The Quarterly Journal of Economics*, Vol. 115, No. 1, pp.147-180.
- [12] Cochrane, J.H., 2011. "Determinacy and Identification with Taylor Rule," *Journal of Political Economy*, vol. 119(3), pp. 565-615
- [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] Crowe, C. and E.E. Meade, 2008. "Central bank independence and transparency: Evolution and effectiveness," *European Journal of Political Economy*, Elsevier, vol. 24(4), pp. 763-777.

- [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] Draghi, M., 2014. "Monetary policy communication in turbulent times," Speech at the Conference De Nederlandsche Bank 200 years: Central banking in the next two decades, (April).
- [19] Drechsler, I., 2013. "Uncertainty, Time-Varying Fear and Asset Prices," *Journal of Finance*, vol. 68(5), pp. 1843-1889
- [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. "Three Types of Ambiguity," *Journal of Monetary Economics*, Volume 59, Issue 5, pp. 422-445.
- [24] Hansen, L.P. and T. J. Sargent, 2012. "Three Types of Ambiguity," *Journal of Monetary Economics*, Volume 59, Issue 5, pp. 422-445.
- [25] Ilut, C. and M. Schneider, 2014. "Ambiguous Business Cycles," *American Economic Review*, American Economic Association, Vol. 104(8), pp. 2368-99.
- [26] Kiley, M.T. 2007. "Is Moderate-to-High Inflation Inherently Unstable?," *International Journal of Central Banking*, vol. 3(2), pp. 173-201.
- [27] Levin, T. and J. Piger, 2004. "Is inflation persistence intrinsic in industrial economies?" *European Central Bank Working Paper* 334.
- [28] 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.
- [29] 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
- [30] Taylor, John B., 1993. "Discretion versus Policy Rules in Practice". *Carnegie-Rochester Conference Series on Public Policy* 39: 195-214.
- [31] Taylor, J.B., 1993. "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy* 39: 195-214.
- [32] Woodford, M., 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton University Press.
- [33] Yellen, J. 2013. "Communication in Monetary Policy," speech at the Society of American Business Editors and Writers 50th Anniversary Conference, Washington, D.C., (April).

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

A Proofs of Steady State Results

Proof of Result 5.1

In steady state, equation 10 becomes:

$$1 = \frac{\beta \tilde{R}(\mu, \cdot)}{\Pi(\mu, \cdot)} \quad (47)$$

From the Taylor rule we get:

$$\tilde{R}(\mu, \cdot) = R^n(\mu, \cdot) \Pi(\mu, \cdot)^\phi e^\mu = \frac{1}{\beta} \Pi(\mu, \cdot)^\phi e^\mu \quad (48)$$

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 $\phi > 1$. \square

Proof of Result 5.2

$\mathbb{V}(\mu, \cdot)$, as defined in equation 23, 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:

$$\left. \frac{\partial^2 \mathbb{V}(\mu, \cdot)}{\partial \mu^2} \right|_{\mu=0} = - \frac{\theta ((\beta - 1)^2 \theta + \epsilon (\beta \theta - 1)^2 (1 + \psi))}{(1 - \beta)(\theta - 1)^2 (\beta \theta - 1)^2 (\phi - 1)^2 (1 + \psi)} \quad (49)$$

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. \square

Proof of Result 5.3

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

$$\left. \frac{\partial^3 \mathbb{V}(\mu, \cdot)}{\partial \mu^3} \right|_{\mu=0} = \frac{\epsilon(2\epsilon - 1)\theta(1 + \theta)}{(1 - \beta)(1 - \theta)^3(\phi - 1)^3} + \mathcal{R}(\beta) \quad (50)$$

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

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

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:

$$\mathbb{V}(\mu_0, \cdot) - \mathbb{V}(-\mu_0, \cdot) = \left. \frac{\partial^3 \mathbb{V}(\mu, \cdot)}{\partial \mu^3} \right|_{\mu=0} \frac{2\mu_0^3}{6} + o(\mu_0^4), \quad (51)$$

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 μ_0 . This, combined with the absence of internal minima (Result 5.2), delivers our result. \square

Proof of Result 5.4

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

The strict inequality follows from the characterization of the worst case in Results 5.2 and 5.3; in particular from the fact that $\mathbb{V}_w(\bar{\mu}) = \mathbb{V}(-\bar{\mu}, \cdot)$ 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 5.1, $\phi > 1$ and given that the worst case corresponds to $\mu = -\bar{\mu}$, it is immediate to verify that $\frac{\bar{\mu}}{\phi-1} > \frac{\bar{\mu}'}{\phi-1}$. $\phi > 1$ also ensures that the Taylor rule is increasing in inflation more than one for one, which delivers the last inequality. \square

Proof of Result 5.5

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) . $(\mu, \mu') \in [-\bar{\mu}, 0) \times [-\bar{\mu}, 0)$ 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 $(\mu, \mu') \in (0, \bar{\mu}] \times (0, \bar{\mu}]$. \square

B Derivations and Proofs of Model Dynamics

B.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 $\mu = -\bar{\mu}$ one obtains the log-linear approximation around the worst-case steady state:

$$\begin{aligned}
c_t &= \mathbb{E}_t c_{t+1} - (\tilde{r}_t - \mathbb{E}_t \pi_{t+1}) \\
w_t &= c_t + \psi n_t \\
\pi_t &= \kappa_0(\mu, \cdot) m c_t + \kappa_1(\mu, \cdot) \mathbb{E}_t \widehat{F}^2_{t+1} + \kappa_2(\mu, \cdot) \mathbb{E}_t \pi_{t+1} \\
r_t &= r_t^n + \phi \pi_t \\
\tilde{r}_t &= r_t + \mu_t \\
m c_t &= w_t - a_t \\
y_t &= a_t - \tilde{\Delta}_t + n_t \\
c_t &= y_t \\
r_t^n &= a_{t+1} - a_t \\
y_t^n &= a_t \\
\hat{\Delta}_t &= \Pi(\mu, \cdot)^\epsilon \theta \hat{\Delta}_{t-1} + \epsilon \left(\Pi(\mu, \cdot)^\epsilon \theta - (1 - \Pi(\mu, \cdot)^\epsilon \theta) \frac{\theta}{\left(\frac{1}{\Pi(\mu, \cdot)}\right)^{\epsilon-1} - \theta} \right) \pi_t \\
\widehat{F}^2_t &= (\epsilon - 1) \beta \theta \Pi(\mu, \cdot)^{\epsilon-1} \mathbb{E}_t \pi_{t+1} + \beta \theta \Pi(\mu, \cdot)^{\epsilon-1} \mathbb{E}_t \widehat{F}^2_{t+1}
\end{aligned}$$

Where we define:

$$\kappa_0(\mu, \cdot) \equiv \frac{\left(\left(\frac{1}{\Pi(\mu, \cdot)} \right)^{\epsilon-1} - \theta \right) (1 - \beta \theta \Pi(\mu, \cdot)^\epsilon)}{\theta} \quad (52)$$

$$\kappa_1(\mu, \cdot) \equiv \beta \left(\left(\frac{1}{\Pi(\mu, \cdot)} \right)^{\epsilon-1} - \theta \right) (\Pi(\mu, \cdot) - 1) \Pi(\mu, \cdot)^{\epsilon-1} \quad (53)$$

$$\kappa_2(\mu, \cdot) \equiv \beta \Pi(\mu, \cdot)^{\epsilon-1} \left(\theta(\epsilon - 1)(\Pi(\mu, \cdot) - 1) + (1 - \epsilon + \epsilon \Pi(\mu, \cdot)) \left(\frac{1}{\Pi(\mu, \cdot)} \right)^{\epsilon-1} \right) \quad (54)$$

$$\kappa_3(\mu, \cdot) \equiv \Pi(\mu, \cdot)^\epsilon \theta \quad (55)$$

$$\kappa_4(\mu, \cdot) \equiv \epsilon \left(\Pi(\mu, \cdot)^\epsilon \theta - (1 - \Pi(\mu, \cdot)^\epsilon \theta) \frac{\theta}{\left(\frac{1}{\Pi(\mu, \cdot)}\right)^{\epsilon-1} - \theta} \right) \quad (56)$$

$$\kappa_5(\mu, \cdot) \equiv (\epsilon - 1) \beta \theta \Pi(\mu, \cdot)^{\epsilon-1} \quad (57)$$

$$\kappa_6(\mu, \cdot) \equiv \beta \theta \Pi(\mu, \cdot)^{\epsilon-1} \quad (58)$$

The equations above (amended for the fact that we allow μ_t to vary) can be

summarized in the following system of four equations:

$$\begin{aligned}
\tilde{y}_t &= \mathbb{E}_t \tilde{y}_{t+1} - (\phi \pi_t + \hat{\mu}_t - \mathbb{E}_t \pi_{t+1}) \\
\pi_t &= \kappa_0(-\bar{\mu}, \cdot) \left((1 + \psi) \tilde{y}_t + \psi \hat{\Delta}_t \right) + \kappa_1(-\bar{\mu}, \cdot) \mathbb{E}_t \widehat{F2}_{t+1} + \kappa_2(-\bar{\mu}, \cdot) \mathbb{E}_t \pi_{t+1} \\
\hat{\Delta}_t &= \kappa_3(-\bar{\mu}, \cdot) \hat{\Delta}_{t-1} + \kappa_4(-\bar{\mu}, \cdot) \pi_t \\
\widehat{F2}_t &= \mathbb{E}_t \left(\kappa_5(-\bar{\mu}, \cdot) \pi_{t+1} + \kappa_6(-\bar{\mu}, \cdot) \widehat{F2}_{t+1} \right)
\end{aligned}$$

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 $\mu = -\bar{\mu}$.

It is then possible to verify that the following guesses solve the system above²⁹:

$$\begin{aligned}
\pi_t &= \lambda_{\pi\Delta} \Delta_{t-1} + \lambda_{\pi\mu} \hat{\mu}_t \\
\tilde{y}_t &= \lambda_{y\Delta} \Delta_{t-1} + \lambda_{y\mu} \hat{\mu}_t \\
\widehat{F2}_t &= \lambda_{F\Delta} \Delta_{t-1} + \lambda_{F\mu} \hat{\mu}_t
\end{aligned}$$

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.

B.2 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:

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

We know that $N(-\bar{\mu}, \cdot) \simeq 1$ ³⁰ so variations 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

²⁹The 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.

³⁰It equals 1 when $\bar{\mu} = 0$ 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). Setting $\psi = 1$ as in our baseline or $\psi = 0$ as is here

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.

B.3 Proof of Result 5.1

$$\begin{aligned} d\mathbb{V}(a_t, \hat{\Delta}_{t-1}; \hat{\mu}_t) &= \log(1 + \lambda_Y \hat{\mu}_t + a_t) - \left((\lambda_Y + \kappa_4 \lambda_\pi) \hat{\mu}_t + \kappa_3 \hat{\Delta}_{t-1} \right) N(-\bar{\mu}, \cdot) \\ &+ \beta \mathbb{E}_t d\mathbb{V}(a_{t+1}, \hat{\Delta}_t; \hat{\mu}_t) \end{aligned} \quad (59)$$

Substituting forward and using the law of motion for price dispersion delivers:

$$\begin{aligned} d\mathbb{V}(a_t, \hat{\Delta}_{t-1}; \hat{\mu}_t) &= \mathbb{E}_t \left(\sum_{j=0}^{\infty} \beta^j \log(1 + \lambda_Y \hat{\mu}_t + a_{t+j}) \right) - N(-\bar{\mu}, \cdot) \frac{1}{1 - \beta \kappa_3} \hat{\Delta}_{t-1} \\ &- N(-\bar{\mu}, \cdot) \left(\frac{\lambda_Y}{1 - \beta} + \frac{\kappa_4 \lambda_\pi}{(1 - \kappa_3)} \left(\frac{1}{1 - \beta} - \frac{\kappa_3}{1 - \beta \kappa_3} \right) \right) \hat{\mu}_t \end{aligned} \quad (60)$$

Then we can compute:

$$\left. \frac{\partial \left(d\mathbb{V}(a_t, \hat{\Delta}_{t-1}) \right)}{\partial \hat{\mu}_t} \right|_{\hat{\mu}_t=0} = \mathbb{E}_t \left(\sum_{j=0}^{\infty} \beta^j \frac{\lambda_Y}{(1 + a_{t+j})} \right) - \Xi(-\bar{\mu}, \cdot) \quad (61)$$

Where $\Xi(-\bar{\mu}, \cdot) \equiv N(-\bar{\mu}, \cdot) \left(\frac{\lambda_Y}{1 - \beta} + \frac{\kappa_4 \lambda_\pi}{(1 - \kappa_3)} \left(\frac{1}{1 - \beta} - \frac{\kappa_3}{1 - \beta \kappa_3} \right) \right)$ does not depend on the state variables.

Suppose now that $\lambda_Y > 0$ then:

$$\left. \frac{\partial \left(d\mathbb{V}(a_t, \hat{\Delta}_{t-1}) \right)}{\partial \hat{\mu}_t} \right|_{\hat{\mu}_t=0} > \frac{\lambda_Y}{(1 + \bar{a})(1 - \beta)} - \Xi(-\bar{\mu}, \cdot) \quad (62)$$

Where the inequality will always be strict unless \bar{a} is an absorbing state, which it cannot be if we think of it as a deviation from steady state unless $\bar{a} = 0$, i.e. the technology process is constant. This equation provides a lower bound on the derivative of the welfare function around steady state, hence it provides our sufficient condition.

If $\lambda_Y < 0$ then:

$$\left. \frac{\partial \left(d\mathbb{V}(a_t, \hat{\Delta}_{t-1}) \right)}{\partial \hat{\mu}_t} \right|_{\hat{\mu}_t=0} > \frac{\lambda_Y}{(1 + \underline{a})(1 - \beta)} - \Xi(-\bar{\mu}, \cdot) \quad (63)$$

Which provides the sufficient condition for this case.

Finally if $\lambda_Y = 0$, the derivative does not depend on a_t at all and it is positive

because $0 \leq \kappa_3 < 1$ and $\kappa_4 \geq 0$ ³¹:

$$\left. \frac{\partial \left(d\mathbb{V} \left(a_t, \hat{\Delta}_{t-1} \right) \right)}{\partial \hat{\mu}_t} \right|_{\hat{\mu}_t=0} = N(-\bar{\mu}, \cdot) \left(\frac{\kappa_4}{(1-\kappa_3)(\phi-1)} \left(\frac{1}{1-\beta} - \frac{\kappa_3}{1-\beta\kappa_3} \right) \right) > 0 \quad (64)$$

Note that the expression above would equal zero if $\bar{\mu} = 0$ since $\kappa_4(\bar{\mu} = 0) = 0$. However as $\bar{\mu} \rightarrow 0$, $\lambda_Y \rightarrow -\frac{(1-\beta)\theta}{(\phi-1)(1-\theta)(1-\beta\theta)} < 0$ which contradicts $\lambda_Y = 0$. \square

B.4 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:

$$d\mathbb{V} \left(a_t, \hat{\Delta}_{t-1}; \hat{\mu}_t \right) = \log(1 + c_t) - N(-\bar{\mu}, \cdot)^2 \left(n_t + \frac{1}{2} n_t^2 \right) + \beta \mathbb{E}_t d\mathbb{V} \left(a_{t+1}, \hat{\Delta}_t; \hat{\mu}_t \right) \quad (65)$$

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

$$\begin{aligned} d\mathbb{V} \left(a_t, \hat{\Delta}_{t-1}; \hat{\mu}_t \right) &= \log \left(1 + \lambda_{y\Delta} \hat{\Delta}_{t-1} + \lambda_{y\mu} \hat{\mu}_t + a_t \right) \\ &- N(-\bar{\mu}, \cdot)^2 \left(\left(\lambda_{y\Delta} \hat{\Delta}_{t-1} + \lambda_{y\mu} \hat{\mu}_t + \hat{\Delta}_t \right) + \frac{1}{2} \left(\lambda_{y\Delta} \hat{\Delta}_{t-1} + \lambda_{y\mu} \hat{\mu}_t + \hat{\Delta}_t \right)^2 \right) \\ &+ \beta \mathbb{E}_t d\mathbb{V} \left(a_{t+1}, \hat{\Delta}_t; \hat{\mu}_t \right) \end{aligned} \quad (66)$$

Substituting forward:

$$\begin{aligned} d\mathbb{V} \left(a_t, \hat{\Delta}_{t-1}; \hat{\mu}_t \right) &= \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \log \left(1 + \lambda_{y\Delta} \hat{\Delta}_{t+j-1} + \lambda_{y\mu} \hat{\mu}_t + a_{t+j} \right) \\ &- N(-\bar{\mu}, \cdot)^2 \sum_{j=0}^{\infty} \beta^j \left(\left(\lambda_{y\Delta} \hat{\Delta}_{t+j-1} + \lambda_{y\mu} \hat{\mu}_t + \hat{\Delta}_{t+j} \right) + \frac{1}{2} \left(\lambda_{y\Delta} \hat{\Delta}_{t+j-1} + \lambda_{y\mu} \hat{\mu}_t + \hat{\Delta}_{t+j} \right)^2 \right) \end{aligned} \quad (67)$$

Then note:

$$\hat{\Delta}_{t+j} = \gamma_{\Delta}^{j+1} \hat{\Delta}_{t-1} + \kappa_4 \lambda_{\pi\mu} \sum_{l=0}^j \gamma_{\Delta}^l \hat{\mu}_t \quad \gamma_{\Delta} \equiv (\kappa_3 + \kappa_4 \lambda_{\pi\Delta}) \quad (68)$$

Note that:

- 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.

³¹probably need to specify this is true in the worst-case because inflation is too high

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:

$$\hat{\Delta}_{t+j} = \kappa_4 \lambda_{\pi\mu} \frac{1 - \gamma_{\Delta}^{j+1}}{1 - \gamma_{\Delta}} \hat{\mu}_t \quad (69)$$

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

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

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

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

Or:

$$\begin{aligned} \left. \frac{\partial (d\mathbb{V}(a_t, \hat{\Delta}_{t-1}))}{\partial \hat{\mu}_t} \right|_{\hat{\mu}_t=0, \hat{\Delta}_{t-1}=0} &= \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \frac{\lambda_{y\Delta} \kappa_4 \lambda_{\pi\mu} \frac{1 - \gamma_{\Delta}^j}{1 - \gamma_{\Delta}} + \lambda_{y\mu}}{(1 + a_{t+j})} - \frac{N(-\bar{\mu}, \cdot)^2 \kappa_4 \lambda_{\pi\mu}}{(1 - \beta)(1 - \gamma_{\Delta})} (1 + \lambda_{y\Delta} + \lambda_{y\mu}) \\ &\quad + \frac{N(-\bar{\mu}, \cdot)^2 \kappa_4 \lambda_{\pi\mu}}{(1 - \beta\gamma_{\Delta})(1 - \gamma_{\Delta})} (\lambda_{y\Delta} + \gamma_{\Delta}) \\ &= \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \frac{\lambda_{y\Delta} \kappa_4 \lambda_{\pi\mu} \frac{1 - \gamma_{\Delta}^j}{1 - \gamma_{\Delta}} + \lambda_{y\mu}}{(1 + a_{t+j})} - \Xi(-\bar{\mu}, \cdot) \end{aligned} \quad (70)$$

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

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 ($\bar{\mu} = 50\text{bp}$), $\Xi(-\bar{\mu}, \cdot) < 0$. The term governing the marginal utility of consumption is more complicated now because of the presence of the γ_{Δ}^j term. Numerically it is possible to verify that the individual coefficients in the sum are positive. Hence the sum is minimized at \bar{a}^{32} , but even then it would correspond to positive number. In other words our sufficient condition is met for any value of the states.

³²We are obviously ruling out values of $a < -1$ because they would imply negative TFP in levels.

B.5 Proof of Proposition 7.1

The only purpose of the lower bound $\underline{\phi}(-\bar{\mu}, \cdot) \geq 1$ 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:

$$\Pi(\mu, \delta(\cdot), \cdot) = e^{-\frac{\mu + \delta(\cdot)}{\phi - 1}} \quad (71)$$

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

$$\mathbb{V}_\delta(\mu, \cdot) = \mathbb{V}(\mu + \delta(\cdot), \cdot) \quad (72)$$

Graphically, this amounts to shifting the function leftward by δ . So \mathbb{V}_δ inherits all the properties of \mathbb{V} established in Results 5.2 and 5.3, except \mathbb{V}_δ 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. $\bar{\phi}$ is the welfare-maximizing value of $\phi \in [\underline{\phi}, \bar{\phi}] \forall \bar{\mu} > \delta > 0$.

Following the same logic as in Result 5.5, it is easy to verify that for $1 < \phi' < \bar{\phi}$, there exists a μ' s.t.:

$$\mathbb{V}(\mu', \phi', \cdot) = \mathbb{V}(-\bar{\mu}, \bar{\phi}, \cdot) \quad (73)$$

In particular:

$$\mu' = -\bar{\mu} \frac{\phi' - 1}{\phi - 1} - \delta \frac{\bar{\phi} - \phi'}{\phi - 1} \quad (74)$$

Our restriction on δ implies that $\frac{\partial \mu'}{\partial \phi'} = \frac{-\bar{\mu} + \delta}{\phi - 1} < 0$. Since $\phi' \in (1, \bar{\phi})$ ³³ we know that $0 > -\delta > \mu' > -\bar{\mu}$.

Strict concavity (Result 5.2) and the fact that the maximum is attained at $-\delta$ then imply:

$$\mathbb{V}(-\bar{\mu}, \bar{\phi}, \cdot) = \mathbb{V}(\mu', \phi', \cdot) > \mathbb{V}(-\bar{\mu}, \phi', \cdot) \quad (75)$$

³³Where we consider the lowest possible value for $\underline{\phi}$, i.e. unity

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

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

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

$$\delta^*(\bar{\mu}) : \quad \mathbb{V}(-\bar{\mu} + \delta^*(\bar{\mu}), \cdot) = \mathbb{V}(\bar{\mu} + \delta^*(\bar{\mu}), \cdot) \quad (76)$$

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

A number of different cases then arise:

1. $\mu_0 \in (-\bar{\mu}, \bar{\mu})$: then $\mathbb{V}'(-\bar{\mu}, \cdot) > 0 > \mathbb{V}'(\bar{\mu}, \cdot)$ ³⁴
 - a. $\mathbb{V}(-\bar{\mu}, \cdot) < \mathbb{V}(\bar{\mu}, \cdot)$. Together with strict concavity this implies that $\mu_{ws} = -\bar{\mu}$. Then there exists a small enough $\delta > 0$ such that

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

So now the worst case $\mu'_{ws} = -\bar{\mu} + \delta$ 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. $\mathbb{V}(-\bar{\mu}, \cdot) > \mathbb{V}(\bar{\mu}, \cdot)$. Together with strict concavity, this implies that $\mu_{ws} = \bar{\mu}$. Then there exists a small enough $\delta < 0$ delivering the same as above
 - c. $\mathbb{V}(-\bar{\mu}, \cdot) = \mathbb{V}(\bar{\mu}, \cdot)$. There is no room for improvement. Any $\delta \neq 0$ would lower the worst-case welfare.
2. $\mu_0 \geq \bar{\mu}$. Strict concavity implies that $\mathbb{V}(-\bar{\mu}, \cdot) < \mathbb{V}(\bar{\mu}, \cdot)$. Hence $\mu_{ws} = -\bar{\mu}$. For all $0 \leq \delta \leq \mu_0 - \bar{\mu}$

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

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

³⁴With 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.

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

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

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

The conditions of the Lemma apply to our case if \bar{m} meets the condition in equation (24), which we assume throughout. The Lemma holds for a generic fixed ϕ , so it obviously applies to $\phi = \bar{\phi}$.

- iii. *in the economy described in Section 4, $0 < \delta^*(\bar{\mu}, \bar{\phi}; \cdot) < \bar{\mu}$.*

Results 5.2 and 5.3 ensure that our economy falls under case 1a of Lemma B.1. This proves that $\delta^*(\bar{\mu}) > 0$. Suppose now that $\delta^*(\bar{\mu}) \geq \bar{\mu}$. That would push the argmax of the welfare function outside (or on the boundary) of $[-\bar{\mu}, \bar{\mu}]$, which can never be optimal given strict concavity (similar arguments to cases 2 and 3 in Lemma B.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 (43) - is minimized. □

B.6 Proof of Corollary 7.1

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 $\bar{\phi}$* somewhat lower.

Consider any policy plan delivering utility v_0 in steady state. Suppose that is welfare-superior to the policy currently in place $\mathbb{V}(-\bar{\mu}, \phi, \cdot) < v_0 \leq \mathbb{V}(0, \phi, \cdot)$, where the latter is the first-best allocation so it cannot be improved upon. Results 5.2 and 5.3 ensure the value function is strictly increasing for $\mu < 0$ so there exists a μ' , $-\bar{\mu} < \mu' \leq 0$, s.t. $\mathbb{V}(\mu', \phi, \cdot) \geq v_0$. Result 5.5 then ensures there also exists ϕ' s.t. $\mathbb{V}(-\bar{\mu}, \phi', \cdot) = \mathbb{V}(\mu', \phi, \cdot)$.

Since our proposed monetary policy rule is also dynamically optimal for any value of ϕ that guarantees determinacy, any $\bar{\phi} \geq \phi'$ allows our monetary policy to deliver at least as high a welfare level as the alternative delivering v_0 .