A Pitfall of Cautiousness in Monetary Policy*

Stéphane Dupraz, Sophie Guilloux-Nefussi,
and Adrian Penalver
Banque de France

Central banks are often reluctant to take immediate or forceful actions in the face of new information on the economic outlook. To rationalize this cautious approach, Brainard's *attenuation principle* is often invoked: when a policymaker is unsure of the effects of his policies, he should react less than he would under certainty. We show that the Brainard principle, while a wise recommendation for policymaking in general, runs into a pitfall when it is applied to a central bank setting monetary policy. For a central bank, concerns about uncertainty create a cautiousness bias: acting less is justified when taking as given the private sector's expectations of inflation, but acting less shifts these inflation expectations away from the central bank's inflation target. In response to the de-anchoring of expectations, the central bank can easily end up acting as much as it is initially reluctant to do, but without succeeding in putting inflation back on target. This pattern is a feature of policy under discretion: the central bank would often be better off tying its hands and not respond to its concerns about uncertainty.

JEL Codes: E31, E52, E58.

1. Introduction

Central banks must set monetary policy under substantial uncertainty on the economic outlook, as well as the effects of their own policies. Faced with this uncertainty, they often react by attenuating their policy response, or by changing it only gradually. To justify

*We thank Gadi Barlevy, Mario Pietrunti, Ricardo Reis, and Ulf Söderström for useful discussions, and participants at the CEPR-Bank of Finland conference, Sciences Po, CEBRA conference, Banca d'Italia, and Banco Central de Chile for valuable comments. The views expressed herein are those of the authors and do not necessarily reflect those of the Banque de France or the Eurosystem.

this cautious approach to policymaking, they often refer to Brainard (1967), who formally derived what came to be known as the Brainard principle: although a policymaker who is uncertain of the economic outlook should act as if his best expectation were a sure outcome (Theil 1957), a policymaker who is uncertain of the effects of his own policies should act less than he would under certainty.

The logic of Brainard's attenuation principle is not limited to monetary policy, but it became well known by central bankers in the 1990s thanks to Alan Blinder's book on his experience as a central banker (Blinder 1999).[1] Blinder himself declared that the Brainard principle "was never far from [his] mind when [he] occupied the Vice Chairman's office at the Federal Reserve." More recently, and in the context of an increased reliance of central banks on unconventional policies, Powell (2018) nicely summed up the Brainard logic through the following formula: "When unsure of the potency of a medicine, start with a somewhat smaller dose." On the other side of the Atlantic, in March 2019 Mario Draghi (2019) explained the decision of the European Central Bank (ECB) Governing Council in the following terms: "You just do what you think is right and you temper [with] a consideration [that] there is uncertainty. In other words, in a dark room you move with tiny steps." Bernanke (2007) and Carney (2017) make similar references to the Brainard principle.[2] Other influential policymakers, such as Williams (2013) and Praet (2018), provide more extended analysis of the Brainard principle.[3]

The attention central bankers declare paying to the Brainard principle suggests it affects their monetary policy decisions, and can

[1]Blinder was not only instrumental in popularizing Brainard's principle but also in giving it its current interpretation of a rationale for "doing less." Blinder himself refers to the Brainard principle as the *conservatism principle*. We follow Reinhart (2003) in using the more neutral terminology of *attenuation principle*.

[2]See also Villeroy de Galhau (2018): "In the face of uncertainty, one often hears reference to the celebrated Brainard "conservatism principle" [...] But we should go beyond a static view of Brainard's principle (which focuses on one single small step): a dynamic view would include the time dimension and consider how to manage and communicate a sequence of incremental steps."

[3]"In these conditions, the Brainard rationale for gradualism applies with great force: do it as carefully and prudently as possible, at least when you have good reason to believe that the degree of uncertainty as to the direction and size of market reactions is atypically large. Present times, where policy is defined by a multiplicity of instruments that interact in ways that are very imperfectly known, are characterized by an abnormal amount of uncertainty" (Praet 2018).

therefore contribute to explaining the dynamics of inflation. In recent years, however, the issue has received little attention in the academic literature. Motivated by the prolonged undershooting of the inflation target in the euro area during much of the past decade, we make a new contribution to the theory of monetary policy under uncertainty.

We show that the Brainard principle, while a wise recommendation for policymaking in general, runs into a pitfall when it is applied to a central bank setting monetary policy. For concreteness, we focus on interest rate policies. When a disinflationary shock hits, the central bank can push inflation back up by cutting interest rates. The Brainard principle would recommend that, if the central bank is uncertain of the precise effect of an interest rate cut on inflation, it should cut interest rates by less, even if this means letting inflation fall somewhat below target. This recommendation, however, abstracts from the fact that inflation also depends on the private sector's expectations of inflation, a dimension that Brainard's original setup does not incorporate. The central bank takes these inflation expectations as given when it acts under discretion, but if the private sector foresees that the central bank will attenuate its policy response, it forms lower inflation expectations. This pushes inflation further down, and forces the central bank to decrease rates further. The central bank easily ends up decreasing rates by as much as it is initially reluctant to do, but with an inflation rate further below target than if it had not been concerned about uncertainty.

We give the name *cautiousness bias* to this perverse incentive that turns the central bank's concerns over uncertainty against its own interests. The terminology is in direct reference to the inflation bias expounded by Kydland and Prescott (1977) and Barro and Gordon (1983a, 1983b). Like the inflation bias, the cautiousness bias is a feature of policy under discretion: it arises because the central bank fails to internalize the effect of its policy on inflation expectations. Contrary to the inflation bias, however, it does not arise from a desire by the central bank to set output above its natural level. It does not even require the central bank to care about stabilizing output, and applies equally to a central bank that has a single mandate to stabilize inflation only.

Our analysis of the cautiousness bias is motivated by the inflation dynamics of the euro area in recent years. Over much of the past decade, ECB monetary policy decisions oscillated between a cautious and gradual approach in the face of uncertainty (as exemplified by Draghi's recommendation to "move with tiny steps in a dark room," cited above) and bold decisive actions when inflation expectations started to risk disanchoring (such as the decision to start quantitative easing in January 2015). This dual strategy seems in line with the discussion of the Brainard principle given by Peter Praet in 2018. Praet (2018) argues that "a case for gradualism can be made in the context of the uncertainty inherent in economic data, models and parameters, notably in times of unconventional monetary policy," but that "a more aggressive monetary policy response, however, is warranted when there is clear evidence of heightened risks to price stability." Our analysis does not object to this distinction but warns against taking the risks to price stability as exogenous: the disanchoring of inflation expectations that calls for an aggressive policy response can be precisely caused by the earlier desire to attenuate the policy response.

To show the robustness of the cautiousness bias, we study it under various specifications of the Phillips-curve relationship between output and inflation. In Section 2, we start by explaining its logic with the New Classical Phillips curve. We show that in response to shocks foreseen by the private sector, a cautious central bank ends up moving real rates by exactly as much as a central bank that disregards concerns over uncertainty would. However, despite ending up moving real rates by the same amount as a central bank that disregards concerns over uncertainty (which is also the optimal policy under commitment), a cautious central banker suffers greater departures of inflation from its target. In the spirit of Rogoff (1985)'s solution to the inflation bias, we show that society would be better off appointing a central banker who discounts concerns over uncertainty relative to society, even if this means responding to unforeseen shocks too aggressively.

Although the case of the New Classical Phillips curve provides a simple exposition of the cautiousness bias, its absence of dynamics prevents an analysis of the ongoing interplay between interest rate decisions and the response of inflation expectations. In Section 3, we study the dynamics induced by the cautiousness bias under the sticky-information Phillips curve of Mankiw and Reis (2002). With the sticky-information Phillips curve, the private sector only gradually incorporates new information into its inflation expectations. As a result, when a negative shock hits inflation expectations move little at first, and the central bank is able to attenuate the decrease in interest rates. But as the private sector gradually realizes the resulting below-target inflation, inflation is pushed down further, forcing the central back to decrease rates further, ultimately by as much as if it had not been willing to attenuate policy.

In Section 4 we show that the cautiousness bias does not depend on the sluggish adjustment of expectations in the New Classical Phillips curve and sticky-information Phillips curve. It applies equally to the forward-looking New Keynesian Phillips curve (NKPC), which remains the most commonly used Phillips curve in economic modeling. The timing in the manifestation of the bias is however different in this case, due to the front-loaded dynamics the NKPC is known to generate (Ball 1994; Mankiw and Reis 2002). In response to a persistent fall in the natural rate, agents immediately expect that the central bank will let inflation fall below target in the future. As a result, the central bank is forced to decrease interest rates more, as early as on impact. In this, the dynamics of the cautiousness bias under the NKPC resembles the one under the New Classical Phillips curve.

The cautiousness bias we focus on in most of the paper concerns the response of inflation to the underlying shocks. Although the resulting undershooting or overshooting of the inflation target can be very persistent in the face of very persistent shocks, it does not create an incentive for a discretionary central bank to let *average* inflation depart form the inflation target π^*, in contrast to the inflation bias of Kydland and Prescott. In Section 5, we show that this is only due to an implicit assumption of the frameworks used in previous sections. By generalizing the setup, we show that the conflict between the desire to stabilize inflation and the desire to minimize inflation uncertainty can also lead to an average bias, just as the conflict between the desire to stabilize inflation and the desire to stabilize output can lead to both an inflation bias and a stabilization bias (Svensson 1997).

For concreteness we analyze the cautiousness bias in the context of conventional interest rate policies, but its logic applies equally to unconventional policies such as forward guidance and balance sheet policies—at least when these are intended as alternative ways to stimulate aggregate demand. What is key to the cautiousness bias is not the way monetary policy affects aggregate demand, but the way aggregate demand affects inflation through the Phillips curve. Because unconventional policies are precisely the ones whose effects are likely to be the most uncertain (see, e.g., Williams 2013), the importance for a central bank to be aware of a bias toward excessive caution is all the more important when the effective lower bound (ELB) on nominal interest rates only leaves unconventional policies available.

The cautiousness bias has another implication for unconventional policies. Although in our framework *real* interest rates never move more than if the central bank had not tried to attenuate policy, *nominal* interest rates can. In the presence of the ELB, this implies that a central bank can find itself up against the ELB and forced to turn to unconventional policies even though it would not have, had it not tried to attenuate policy.

A number of papers have considered the implications of model uncertainty for the conduct of monetary policy. Clarida, Gali, and Gertler (1999); Estrella and Mishkin (1999); Svensson (1999); Sack (2000); Sack and Wieland (2000); Rudebusch (2001), and more recently Williams (2013), recover Brainard's recommendation for policy attenuation in the context of monetary policy. Subsequent literature has emphasized situations in which Brainard's attenuation principle is overturned and uncertainty calls instead for a more aggressive response.[4] Söderström (2002), Kimura and Kirozumi (2007), and Ferrero, Pietrunti, and Tiseno (2019) consider such situations while still modeling the central bank's uncertainty in a Bayesian way, as in Brainard's original setup (and ours). Söderström (2002) shows policy aggressiveness can be called for when uncertainty bears on the persistence of inflation, in a model with adaptive expectations. Kimura and Kirozumi (2007) show it is appropriate when uncertainty bears on the fraction of firms that form expectations in a rule-of-thumb, adaptive, fashion. Closer to this paper, Ferrero, Pietrunti, and Tiseno (2019) show that uncertainty about the slope of the New Keynesian Phillips curve can lead the central bank to move nominal interest rates by more than under certainty in response to cost-push shocks, if shocks are persistent enough. We

[4]Brainard (1967)'s original paper already contains situations in which uncertainty calls for more aggressive policy, as we discuss in Section 2. See also Chow (1973), Craine (1979), and Walsh (2003).

interpret the result through the lens of the cautiousness bias: the optimal discretionary policy is to attenuate the policy response for given inflation expectations, but the adverse reaction of inflation expectations forces the central bank to ultimately act more. Overall, our contribution is to show that for a central bank, following the Brainard principle ends up generating excessive deviations of inflation from target simply because agents in the economy are forward looking and understand how such cautiousness affects inflation dynamics.

Other papers consider the consequences of modeling the central bank's uncertainty through the minmax approach of robust control instead of the Bayesian approach. They usually find that uncertainty calls for more aggressive policy, in opposition to the Brainard principle. Giannoni (2002) finds that the fear that the worst will happen to output and inflation if the central bank does not track the natural rate provides an incentive to track it more closely—i.e., to move interest rates more aggressively (see also Stock 1999; Tetlow and von zur Muehlen 2001; Onatski and Stock 2002; Söderström and Leitemo 2008). Sargent (1999) finds that uncertainty about the persistence of shocks calls for a more aggressive response of monetary policy. Barlevy (2011) argues that what is conducive to more aggressive policy under robust control is less the minmax approach per se than its application to specific situations.[5] He gives examples where the minmax approach calls for policy attenuation, for the same reason as under the Bayesian approach, and shows that uncertainty on the persistence of shocks calls for aggressive policy in the Bayesian setup as well.[6]

Other arguments for attenuation or gradualism have been put forward that do not rely on the presence of uncertainty.[7] Woodford (2003c) shows that the optimal, history-dependent, monetary policy

[5]See also Onatski and Williams (2003) and Tillmann (2009).

[6]Using as well the minmax approach of robust control, Woodford (2010) and Adam and Woodford (2012) consider the design of monetary policy rules that are robust to the possibility that the private sector forms expectations using a wrong model, even though the central bank is itself sure of the model of the economy.

[7]Still other arguments have been put forward to explain gradualism positively, without defending it normatively. For instance, Riboni and Ruge-Murcia (2010) argue that gradualism in monetary policy is partly due to the consensus-building approach taken by many monetary policy committees, a decisionmaking procedure that favors the status quo. See also Favaretto and Masciandaro (2016). Spiegler (2021) shows that when the private sector's subjective causal model

under commitment features inertia and can be approximated by a discretionary central bank that puts a cost on abrupt changes in interest rates.[8]

A third argument for gradualism is based on concerns about the stability of the financial system. As argued by Cukierman (1991), interest rate smoothing can be desirable because it mitigates sudden changes in banks' short-term funding costs or long-term asset returns, and therefore in banks' profits and balance sheets. Interestingly, in a recent paper Stein and Sunderam (2018) show that this distinct motive for gradualism can also lead to a time-inconsistency problem: If the central bank dislikes volatile long-term rates for financial stability reasons, it has an incentive to track the natural interest rate only gradually, in order not to reveal information on long-term natural rates that would make long-term rates react too abruptly. But this is taking markets' expectations as given: in equilibrium markets understand that the central bank is moving gradually and adjust their expectations of long-term natural rates accordingly, partly undoing the central bank's efforts. We show that time inconsistency is equally at play when gradualism is driven by uncertainty concerns. The time-inconsistency problem is different between the two models however: in our model the cautiousness bias arises from a failure to internalize inflation expectations, while in Stein and Sunderam's model time inconsistency arises from a failure to internalize expectations of future interest rates—there are no inflation expectations in their model.[9]

---

reverses the direction of causality between inflation and output, the central bank can react by attenuating its policy response relative to when the private sector has rational expectations. Other costs of delaying policy action have also been put forward. Acharya et al. (2019) show that under skill-loss hysteresis, not acting promptly can push the economy into a recession from which it can then only slowly recover, if at all.

[8]Since Woodford's argument for gradualism does not rely on concerns over uncertainty, however, it is not a rationale for attenuating or delaying the policy response more when uncertainty is higher. In particular, there is no rationale for being more reluctant to act when the only instruments available are unconventional instruments with more uncertain effects. As far as Odyssean forward guidance is concerned, it is precisely Woodford's argument in favor of inertia in interest rates that makes committing to keeping rates lower for longer a superior strategy (Eggertsson and Woodford 2003).

[9]As a consequence, the cautiousness bias is at play in our model even though we make the New Keynesian assumption that aggregate demand and inflation

2. A Simple Model of the Cautiousness Bias

In this section, we expose the cautiousness bias in a simple model where the supply side is captured by the New Classical Phillips curve (Lucas 1972). Using the New Classical Phillips curve has two advantages. First, it is the Phillips curve for which the bias appears most transparently. Second, it follows the classic accounts of the inflation bias by Kydland and Prescott (1977) and Barro and Gordon (1983a, 1983b).

2.1 The Problem of the Central Bank

The problem of the central bank is to pick an allocation for inflation π_t, the output gap x_t, and the nominal interest rate i_t that best fits its objective, subject to the constraints imposed by the behavior of the private sector. These constraints are captured by a simple two-equation model. The aggregate-demand side of the economy is represented by the Euler equation:

$$x_t = -\sigma(i_t - E_t(\pi_{t+1})) + E_t(x_{t+1}) + v_t, \tag{1}$$

where σ is the intertemporal elasticity of substitution, and v_t is a possibly autocorrelated exogenous shock with mean zero, observable at period t. The shock v_t captures variations in natural output y_t^n or, equivalently, in the natural rate of interest r_t^n. Specifically, v_t is the function $v_t = -(y_t^n - E_t(y_{t+1}^n))$ of natural output, and connects to the natural rate through $v_t = \sigma r_t^n$. Appendix A derives and discusses the connection between these alternative representations of the fundamental shocks to the Euler equation.

The aggregate supply side of the economy is captured by the New Classical Phillips curve:

$$\pi_t = \kappa x_t + E_{t-1}(\pi_t), \tag{2}$$

where κ is the slope of the Phillips curve. The private sector's past expectations of present inflation, formed at $t-1$, shift the Phillips

depend on the long-term—not short-term—interest rate. In Stein and Sunderam's model, for the time-inconsistency problem to be at play it is necessary that the central bank's interest rate target bears on the short-term rate only.

curve. The New Classical Phillips curve can be derived for instance under the assumption that a fraction of firms set their prices at t with outdated information from $t-1$ (Woodford 2003b; Mankiw and Reis 2010).

As in the literature on the inflation bias which distinguishes between anticipated inflation and surprise inflation (e.g., Barro and Gordon 1983a, 1983b), we allow for the shock v_t to be partially anticipated by the firms that set their prices with outdated $t-1$ information, by letting v_t be partially forecastable with $t-1$ information. Accordingly, we refer to $E_{t-1}(v_t)$ as the *foreseen shocks* and to $v_t - E_{t-1}(v_t)$ as the *unforeseen shocks.*[10]

Crucially, the central bank faces parameter uncertainty. It is uncertain about the value of the structural parameter σ, and entertains several possible values for it. Like Brainard (1967), we follow Savage (1954) in modeling parameter uncertainty in a Bayesian way. The central bank assigns probabilities to every possible value of σ and treats it as a random variable. We note $\bar{\sigma}$ and V_σ the mean and variance of the central bank's subjective beliefs over σ. We assume that V_σ is constant over time. We assume the central bank is certain of the value of κ. We do so because assuming uncertainty bears only on σ is the case most favorable to Brainard's attenuation principle, as will become clear below.

Although the central bank is uncertain of the model of the economy, we assume that the models it entertains are not too far from the actual one, in the spirit of rational expectations. Specifically, we assume that the true value of σ is $\bar{\sigma}$, the mean of the values considered by the central bank. The true dynamics of the economy are therefore given by the Euler equation (1) and Phillips curve (2) with $\sigma = \bar{\sigma}$. Note that we implicitly assume that the private sector is not subject to parameter uncertainty, since we assume that the Euler equation (1) and Phillips curve (2) hold, both of which are derived under the assumption of no parameter uncertainty. As a consequence, the central bank and the private sector have different information sets at t. To avoid any confusion, we denote by $E^*_t(.)$ the expectations of the central bank at t, which are formed without

[10]We take *"anticipated," "foreseen,"* and *"expected"* as synonyms, but reserve *"foreseen"* to the private sector's anticipations of the exogenous shocks and *"expected"* to the private sector's anticipations of inflation.

knowing σ. We reserve the notation $E_t(.)$ for the expectations of the private sector, which knows that $\sigma = \bar{\sigma}$. We assume that the private sector's expectations are part of the central bank's information set. This is meant to capture the fact that central banks have access to—and heavily monitor—measures of the private sector's inflation expectations before taking monetary policy decisions, such as market-based expectations or surveys of professional forecasters.[11]

We assume that the mandate of the central bank is to stabilize inflation only. Its objective is to set inflation π_t to a target π^* at all periods.[12] It has the quadratic loss function:

$$\mathcal{L}_\infty = E_t^* \left(\sum_{k=0}^{\infty} \beta^k (\pi_{t+k} - \pi^*)^2 \right). \tag{3}$$

The assumption of a single inflation mandate is not necessary for our results. Appendix D shows that they hold equally well in the more general case in which the central bank has a dual objective to stabilize both inflation and the output gap. We focus on the case of a single mandate in the body of the paper for two reasons. First, it emphasizes that the cautiousness bias does not arise from a desire to stabilize output at the expense of stabilizing inflation, unlike the

[11]Note that since the private sector's expectations depend on the parameter σ, the central bank could in theory solve for the dependence of expectations on σ and infer the value of σ from expectations. Such an inference is possible in our model because of the simplicity of its stylized two-equation setup and the simplicity of its information structure. The mapping between σ and expectations could be made arbitrary noisy by adding noise to the model, making the inference arbitrarily uninformative. Alternatively, we could assume that the private sector faces the same model uncertainty as the central bank so that the central bank has nothing to learn from the private sector, but at the cost of less standard forms for the Euler equation and Phillips curve. Since the issue is peripheral to our focus, we simply assume away this inference from endogenous signals, as is common in the literature on incomplete perfect knowledge (e.g., Woodford 2003a or Angeletos and La'O 2010).

[12]With the New Classical Phillips curve (2) (and other information-based Phillips curves), the loss function that can be microfounded from the costs of relative price dispersion contains unexpected inflation, not inflation. However, since in practice the mandate of central banks bears on inflation, we assume so here. We consider any arbitrary inflation target (not necessarily zero) for the same reason. Considering other costs of inflation could justify caring about expected inflation (and various values for the inflation target) even when the Phillips curve is based on information frictions.

inflation bias. Therefore, it applies equally to central banks with a single primary mandate, like the ECB. In the appendix, we allow for the central bank to be willing to set output above potential and therefore be subject to the inflation bias, and show that the cautiousness bias and inflation bias arise from distinct perverse incentives. Second, the assumption of a single objective corresponds to the original framework of Brainard (1967).

2.2 Reductio ad Brainard

We now show that this simple monetary model exactly fits into the framework considered by Brainard (1967), up to one key difference: the presence of the expectations of the private sector. In its canonical form, Brainard's model considers a policymaker who seeks to set a single variable on a target through the use of a single instrument. In our case, the single objective is inflation. We pick the single instrument of the central bank to be the real interest rate $r_t \equiv i_t - E_t(\pi_{t+1})$.[13]

By taking expectations at $t-1$ of the New Classical Phillips curve (2), it must be that the expected output gap next period is zero, $E_t(x_{t+1}) = 0$. Plugging in the expression for the output gap from the Euler equation (1) into the Phillips curve (2), we get

$$\pi_t = -\phi r_t + \varepsilon_t + E_{t-1}(\pi_t), \tag{4}$$

where we define $\phi \equiv \sigma\kappa$ and $\varepsilon_t \equiv \kappa v_t$. We denote $\bar{\phi} = \kappa\bar{\sigma}$ the mean of ϕ and $V_\phi = \kappa^2 V_\sigma$ its variance.

Since the relationship (4) only contains period-t variables, the objective of the central bank reduces to setting the interest rate r_t to minimize the present-period loss:

$$\mathcal{L}_t(\varepsilon_t) = E_t^*(\pi_t - \pi^*)^2, \tag{5}$$

[13]To be sure, in practice the central bank sets a path for the nominal interest rate, but the implementation of the optimal policy is an issue distinct from the choice of the optimal policy, which is the one we consider here. The latter is a path for all three variables i_t, π_t, and x_t, subject to the constraints imposed by the Euler equation (1) and Phillips curve (2). Parameterizing the equilibrium through the three variables r_t, π_t, and x_t is simply a convenient change of variables, and one that fits into Brainard's framework.

at all periods t and for all realizations of ε_t, subject to constraint (4). Since there is no ambiguity, we drop the time subscripts. The following lemma takes stock and draws the parallel to Brainard's framework.

LEMMA 1 (Reductio ad Brainard). *The program of the central bank is, for any realization of the shock ε, to pick the interest rate r that minimizes*

$$\mathcal{L}(\varepsilon) = E^*((\pi - \pi^*)^2), \tag{6}$$

subject to

$$\pi = -\phi r + \varepsilon + E_{-1}(\pi), \tag{7}$$

where the central bank observes ε and $E_{-1}(\pi)$, and ϕ is a random variable with mean $\bar{\phi}$ and variance V_ϕ.

Up to the expectations of the private sector $E_{-1}(\pi)$, this is exactly the framework considered by Brainard (1967).

The random variable ϕ captures the policymaker's uncertainty on how its own action r affects its objective π—in our case, the central bank's uncertainty over the interest rate channel. As Brainard emphasized, this is the type of uncertainty that can justify policy attenuation. Uncertainty over ε only—in our case, uncertainty over the natural rate—would result in Theil's certainty equivalence: it would leave the optimal policy unchanged, up to replacing ε by its expected value (Theil 1957). Our assumption that the central bank perfectly observes ε abstracts from this irrelevant form of uncertainty.

We have restricted the uncertainty over the interest rate channel to arise from uncertainty over σ, the elasticity of demand to changes in the real interest rate. It could also arise from uncertainty over κ, the elasticity of inflation to changes in demand. We focus on σ because uncertainty over κ would create correlation between $\varepsilon = \kappa v$ and $\phi = \kappa\sigma$. As Brainard notes, such correlation can turn the recommendation for policy attenuation into a recommendation for policy aggressiveness. Our qualification of the Brainard principle is distinct and does not rely on correlated shocks. Therefore, we restrict uncertainty to σ to focus on the standard case of uncorrelated shocks, which is the one most favorable to policy attenuation.

Our main point is that the presence of the private sector's expectations in the Brainard model (7) makes important changes to its policy recommendations. When the outcome of the policy depends on the expectations of the private sector, we need to distinguish between policy under discretion and policy under commitment. The cautiousness bias is a feature of policy under discretion, when the central bank takes the inflation expectations of the private sector as given. We start with this case.

2.3 Brainard's Attenuation Principle

We first show that if we fix the inflation expectations of the private sector, Brainard's attenuation principle holds unchallenged. To take explicit note of the fact that the central bank does not internalize its impact on expectations, we temporarily denote expectations $e(\pi)$ instead of $E_{-1}(\pi)$. For the moment they do not have to bear any resemblance to equilibrium outcomes.

To understand the trade-off at the heart of Brainard's attenuation principle, it is helpful to decompose the mean squared error in its loss function (6) into the square of the distance of average inflation from its target, and the perceived variance of inflation:

$$\mathcal{L}(\varepsilon) = (E^*(\pi) - \pi^*)^2 + Var^*(\pi), \tag{8}$$

$$\text{where } (E^*(\pi) - \pi^*)^2 = (-\bar{\phi}r + \varepsilon + e(\pi) - \pi^*)^2, \tag{9}$$

$$\text{and } Var^*(\pi) = V_\phi r^2. \tag{10}$$

This expression makes apparent the two—possibly conflicting—objectives of the central bank. It wants to bring its expectation of inflation (conditional on ε) to target, and it wants to minimize the (conditional) variance of inflation. Note that through both objectives—including the one of bringing expected inflation on target given the realization of ε—the goal of the central bank is to minimize the overall variance of inflation.[14]

[14] By setting the conditional expectation of inflation on target for every realization of ε, the central bank minimizes the between variance of inflation. By minimizing the conditional variance of inflation, it minimizes the within variance of inflation. The unconditional loss function is equal to $\mathcal{L} = E(\mathcal{L}(\varepsilon)) = (E(\pi) - \pi^*)^2 + Var(E(\pi|\varepsilon)) + E(Var(\pi|\varepsilon))$, where the second term is the between variance of inflation and the third term is the within variance of inflation. They sum to the total variance of inflation by the law of total variance.

Denote r^s the interest rate that the central bank sets when it faces no parameter uncertainty, $V_\phi = 0$. In this case the central bank can focus on minimizing the first term (9) in its loss function, and can fully stabilize inflation on target by setting[15]

$$r^s = \bar{r}^n + \frac{e(\pi) - \pi^*}{\bar{\phi}}, \tag{11}$$

where $\bar{r}^n$ denotes the natural rate in the average model:

$$\bar{r}^n \equiv \frac{\varepsilon}{\bar{\phi}} = \frac{v}{\bar{\sigma}}. \tag{12}$$

According to Equation (11), without concerns for parameter uncertainty, the optimal discretionary policy is to track the natural rate, plus a corrective term if inflation expectations are not on target. In this case, it is by responding fully to variations in the natural rate that the central bank reduces inflation volatility and stabilizes the economy.

When the central bank is uncertain of the impact of its rate decision on inflation, $V_\phi > 0$, the policy rate r affects not only the expected value of inflation (9) but also its variance (10). This new dependence captures the fact that, if the central bank is unsure of the consequences of departing from the steady-state rate $r = 0$, its uncertainty is all the greater the larger the departure away from the steady-state rate. Because the policy rate now affects both terms, there is now a trade-off between reaching the inflation target on average (and minimizing the between variance of inflation) and minimizing the variance of inflation. The Brainard principle answers the question of how the central bank solves this trade-off. It can be obtained by taking the first-order condition of the loss function (8).

[15]The result for r^s can equivalently be written $r^s = \kappa v + e(\pi) - \pi^*$. The expression in the text uses the fact that $\bar{\phi} = \kappa\bar{\sigma}$. Note that this decision-theoretic result is valid regardless of whether the central's bank subjective (average) beliefs $\bar{\phi}$ and $\bar{\sigma}$ correspond to the equilibrium ones. The assumption that the central bank's average model corresponds to the equilibrium one—the standard rational expectations assumption—will only intervene later on when we solve for inflation expectations. The same remark applies to the expression of α in Equation (14) below.

LEMMA 2 (Brainard's Attenuation Principle). *Under discretion, the central bank sets the real interest rate as*

$$r = \alpha r^s, \tag{13}$$

$$\text{where } \alpha \equiv \frac{\bar{\phi}^2}{\bar{\phi}^2 + V_\phi}. \tag{14}$$

The central bank solves the trade-off by choosing a midpoint r between the optimal interest rate policy without parameter uncertainty r_s which minimizes the first term, and the steady-state interest rate $r = 0$ which minimizes the second term. Policy becomes biased toward the steady-state interest rate $r = 0$ because the central bank understands the effects of this policy better. Crucially, because α is less than one, the central bank no longer fully reacts to shocks to the natural rate. Uncertainty over the effects of the policy response calls for attenuating the policy response.

Note that under its optimal discretionary policy the central bank does not expect inflation to be on target. Plugging the chosen policy rate (13) into (7), the central bank expects inflation to be

$$E^*(\pi) = \pi^* + (1 - \alpha)(\bar{\phi}\bar{r}^n + e(\pi) - \pi^*) \neq \pi^*. \tag{15}$$

But the central bank is fine with this. It sees it as a cost worth paying to avoid the risks of uncertain policy outcomes.

2.4 The Reaction of Inflation Expectations

The conclusion that the central bank reacts less to shocks is premature, however. The optimal discretionary policy (13) depends on the private sector's expectations of inflation, which are still to be solved for. Crucially, inflation expectations depend on what policy the private sector expects the central bank to implement. If a central bank concerned with parameter uncertainty fights inflation less aggressively, private agents are likely to take it into account in forming inflation expectations.

We solve for the rational expectations of the private sector. Injecting policy (13) into Equation (7), taking expectations E_{-1}, and imposing rational expectations $e(\pi) = E_{-1}(\pi)$ yields

$$E_{-1}(\pi) = \pi^* + \left(\frac{1}{\alpha} - 1\right)\bar{\phi}E_{-1}(\bar{r}^n). \tag{16}$$

When the central bank faces no parameter uncertainty $\alpha = 1$, inflation expectations are on target $E_{-1}(\pi) = \pi^*$, since the private sector rightly anticipates that the central bank will set the policy rate so that inflation is on target under all circumstances. With parameter uncertainty $\alpha < 1$, however, expectations of a natural rate below average leads the private sector to expect below-target inflation. The private sector rightly expects that in this case the cautious central bank will set the real interest rate above the natural rate, creating a negative output gap, and thus below-target inflation.

By plugging the private sector's expectations of inflation in equilibrium (16) into the expression for the policy rate chosen by the central bank (13), we obtain the value of the real interest rate in equilibrium.

PROPOSITION 1 (Brainard Principle Unraveled). *Under the optimal discretionary policy, the real interest rate is in equilibrium:*

$$r = E_{-1}(\bar{r}^n) + \alpha(\bar{r}^n - E_{-1}(\bar{r}^n)). \tag{17}$$

In equilibrium, a central bank with concerns over uncertainty ($\alpha < 1$) attenuates its response only to changes in the natural rate unforeseen by the private sector. To foreseen changes it reacts exactly as much as if it had no concerns about uncertainty.

A cautious central bank ends up reacting just as much to shocks foreseen by the private sector because its reluctance to act pushes inflation to the point at which it is forced to act to the same extent anyway. Assume a shock hits that pushes the natural rate below its average level. Worried by the uncertainty induced if it decreases its policy rate to $\bar{r}^n$, the central bank decides not to fully track the decrease in the natural rate, even if it means letting inflation fall somewhat below target. But if the shock is foreseen by the private sector, this willingness of the central bank to tolerate below-target inflation is, too. Accordingly, the private sector expects lower inflation. Lower inflation expectations put further downward pressure on inflation. In response, the central bank decides to decrease its rate a little more but still in an attenuated manner, which justifies even lower inflation expectations, and so on. Ultimately, inflation expectations fall to the point where they are low enough to convince the central bank to fully match the decrease in the natural rate, as it would have chosen in the absence of concerns for uncertainty.

2.5 The Cautiousness Bias

In reaction to foreseen shocks, the central bank ends up moving its policy rate by as much as if it had no concerns about uncertainty, but inflation ends up further away from target. Formally, the overall departure from target, in response to both unforeseen shocks $\bar{r}^n - E_{-1}(\bar{r}^n)$ and foreseen shocks $E_{-1}(\bar{r}^n)$, is

$$E^*(\pi) - \pi^* = (1-\alpha)\bar{\phi}\left((\bar{r}^n - E_{-1}(\bar{r}^n)) + \frac{1}{\alpha}E_{-1}(\bar{r}^n)\right). \qquad (18)$$

In response to unforeseen shocks $\bar{r}^n - E_{-1}(\bar{r}^n)$, inflation ends up away from target, by $(1-\alpha)\bar{\phi}$ percentage points for every percentage-point change in the natural rate. This is the amount of inflation the central bank was willing to tolerate. But in response to foreseen shocks $E_{-1}(\bar{r}^n)$ inflation ends up further away from target, by an additional factor $1/\alpha$. The outcome in terms of stabilizing inflation is worse than if the central bank ignored policy uncertainty. Since the policy rate is forced into territory the central bank was seeking to avoid, the outcome is as bad in terms of avoiding uncertain outcomes. Overall, in response to foreseen shocks the central bank reaches a worse outcome than if it had not sought to act cautiously. Figure 1 provides a graphical illustration similar to the diagrammatic exposition of the inflation bias by Kydland and Prescott (1977).

The result that the central bank is behaving against its own interest is a feature of policy under discretion. It would not be if policy were chosen under commitment, because this would allow the central bank to internalize the effect of its policy on expectations. Indeed, Appendix B shows the following result.

Proposition 2 (The Cautiousness Bias). *In the optimal allocation under commitment, the policy rate takes the exact same value as under discretion (17), but inflation departs from target by only*

$$E^*(\pi) - \pi^* = (1-\alpha)\bar{\phi}\bigg(\bar{r}^n - E_{-1}(\bar{r}^n)\bigg). \qquad (19)$$

Concerns over uncertainty make the discretionary policy depart from the optimal commitment policy.

Figure 1. Consequences of Attenuating Policy when Expectations React

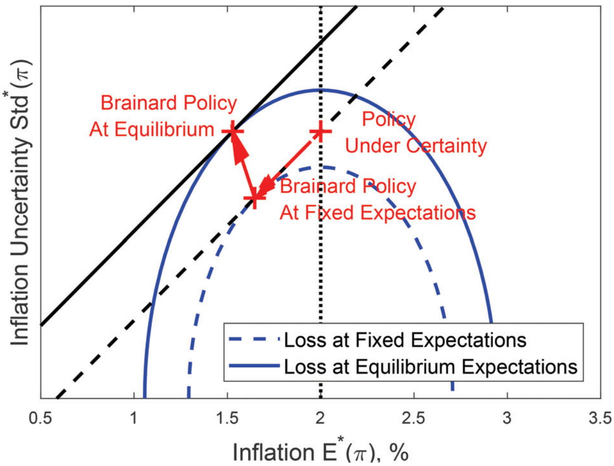

Inflation takes the same value as under discretion in response to unforeseen shocks, but it remains on target in response to foreseen shocks. Under commitment the central bank understands that when shocks are fully foreseen by the private sector, the policy rate can only be equal to the natural rate in equilibrium. It understands that a desire to vary the policy rate by less than the natural rate will only increase inflation expectations up to the point where inflation is enough off target to convince the central bank to vary the policy rate by the full extent of the change in the natural rate. As a consequence, it does not attempt to attenuate the policy response to foreseen shocks and inflation remains on target. It does attenuate the response to unforeseen shocks, however, because these do not risk de-anchoring inflation expectations.

We give the name cautiousness bias to the perverse incentive that turns the central bank's cautiousness into a policy that is no less aggressive, but yields worse stabilization outcomes. Like the inflation bias expounded by Kydland and Prescott (1977) and Barro and Gordon (1983a, 1983b), the cautiousness bias arises because policy chosen under discretion abstracts from the effect of policy on expectations. It differs from the inflation bias, however, in that it does not rely on the desire of the central bank to set output above its natural level. As our assumption of a single inflation mandate highlights, it does not even require the central bank to care about stabilizing output. It arises instead because of the distorted perception of the trade-off between stabilization inflation, and stabilizing the policy rate at values where its effects are better known.[16]

2.6 Guarding Oneself Against Being Cautious

Short of shifting to deciding policy under commitment, what can a central bank—or the society that appoints it—do to guard itself against the cautiousness bias? Rogoff (1985) proposed a solution to realign the incentives of a discretionary policymaker with the preferences of society under commitment: appoint a policymaker whose preferences differ from society's.

We show that society would be better off appointing a central banker who is less cautious than society is. We capture different degrees of cautiousness through different weightings δ of the variance term in the loss function (8):[17]

$$\mathcal{L}(\varepsilon) = (E^*(\pi) - \pi^*)^2 + \delta Var^*(\pi). \tag{20}$$

A discretionary central bank with such preferences still sets policy according to (13), up to a new value for the attenuation coefficient α:

$$\alpha(\delta) = \frac{\bar{\phi}^2}{\bar{\phi}^2 + \delta V_\phi}. \tag{21}$$

[16]Note that the perverse incentive of the cautiousness bias would apply similarly if the central bank's motive for limiting fluctuations in the policy rate—the term in r^2 in its loss function—was driven by something else than concerns over uncertainty—for instance, by concerns over financial stability. Our paper is concerned with Brainard uncertainty, as it is a recurring argument in central bankers' discussions of attenuation and gradualism, but the model can be fruitfully applied to other motives.

[17]The weighting parameter δ allows to encompass several reasons for different degrees of cautiousness. A central banker who perceives less uncertainty on the model parameter ϕ will be akin to one who has a lower δ, since $Var^*(\pi) = V_\phi r^2$. But differences in the degree of cautiousness can reflect pure differences in preferences: the weighting parameter δ can also be seen as capturing different degrees of risk aversion within a class of (squared) mean-variance preferences. Finally, discounting the variance term can be a conscious decision to discount uncertainty concerns in order to counterbalance the cautiousness bias. In this last case, it resembles a form of limited commitment.

We assume that society's true preferences are still captured by the loss function (8), i.e., the loss function (20) with $\delta = 1$, and evaluate the outcome delivered by the various central bankers—different values of δ—according to society's true social preferences.[18] Appendix C shows the following result.

Proposition 3 (Optimal Discounting of Uncertainty Concerns).

- *Unless all shocks are unforeseen by the private sector, it is always desirable to have a central bank that discounts concerns over uncertainty,* $\delta < 1$.
- *The optimal value of* δ *decreases with the proportion of shocks that are foreseen by the private sector.*

A central bank that discounts uncertainty more reacts to shocks more. The benefit is that such a central bank reacts more to foreseen shocks, reducing the cautiousness bias. The cost is that it overreacts to unforeseen shocks. The optimal δ strikes a balance between costs and benefits.

3. The Cautiousness Bias with the Sticky-Information Phillips Curve

We now consider how a more realistic model of the dynamics of inflation affects the workings of the cautiousness bias. While the New Classical Phillips curve is useful to illustrate the logic of the cautiousness bias, its absence of dynamics misses an analysis of the chronology in the policy decisions and their consequences. With the New Classical Phillips curve, the entire dynamics is subsumed into a one-period simultaneous equilibrium: because the private sector anticipates that the central bank will fight deflationary shocks less aggressively at t, inflation expectations are lower at t, and the central

[18]The fact that concerns over uncertainty derive from society's concerns over uncertainty can conflict with our assumption that the private sector faces no parameter uncertainty. As explained in footnote 11, we only make the latter assumption to avoid encumbering the model with a peripheral signal-extraction problem. Alternatively, the social preferences can be interpreted as the preferences of a government which faces the same parameter uncertainty as the central bank.

bank is forced to act at t. To capture the dynamics in a more realistic way, we turn to the sticky-information Phillips curve of Mankiw and Reis (2002). Under the sticky-information Phillips curve, the sequence of events happens sequentially. A negative shock to the natural rate hits; the central bank does not fully track the fall in the natural rate; inflation falls below target; the private sector gradually realizes that inflation is below target and forms lower inflation expectations; lower inflation expectations push inflation down; the central bank is forced to decrease rates further. However, while a cautious central bank can initially attenuate its policy, it still eventually ends up acting as much as if it did not try to attenuate its policy.

3.1 The Problem of the Central Bank

We assume the supply side of the economy is captured by the sticky-information Phillips curve:

$$\pi_t = \kappa x_t + \bar{E}_{t-1}(\pi_t + \zeta \Delta x_t), \tag{22}$$

where ζ is the slope of the short-run aggregate supply (SRAS), $\Delta x_t = x_t - x_{t-1}$ is the growth rate of the output gap, and $\bar{E}_{t-1}$ is notation for the following weighted average of expectations formed at different periods in the infinite past:

$$\bar{E}_{t-1}(\pi_t + \zeta \Delta x_t) = \sum_{j=0}^{\infty} \lambda (1-\lambda)^j E_{t-1-j}(\pi_t + \zeta \Delta x_t). \tag{23}$$

Like the New Classical Phillips curve, the sticky-information Phillips curve models monetary non-neutrality as arising from price setters' imperfect information. Contrary to the New Classical Phillips curve, which assumes all information is incorporated by everyone after one period, the sticky-information Phillips curve assumes that the private sector only gradually learns about the shocks that hit the economy.[19]

[19] Specifically, it can be derived under the assumptions that a fraction λ of firms update their information sets every period, and that the probability of updating information in a given period is independent of how long it has been since the

We only replace the Phillips curve and keep the Euler equation (1) unchanged on the aggregate-demand side. One issue that did not arise in the case of the New Classical Phillips curve is whether to assume that the central bank observes the private sector's current expectation of the output gap tomorrow, $E_t(x_{t+1})$.[20] Assuming one way or the other does not change the results qualitatively. We focus on the case where the central bank does not observe $E_t(x_{t+1})$, as it likely better captures the uncertainty of the central bank on the effect of its policy. Indeed, iterating the Euler equation forward, note that

$$x_t = -\sigma R_t + E_t\left(\sum_{k=0}^{\infty} v_{t+k}\right), \tag{24}$$

where R_t is the long-term real interest rate:

$$R_t \equiv E_t\left(\sum_{k=0}^{\infty} r_{t+k}\right) = r_t + E_t(R_{t+1}). \tag{25}$$

The iterated Euler equation (24) highlights that aggregate demand depends on the effect of the whole sequence of future rates summarized by the long-term real interest rate R_t, i.e., the entire yield curve and not just the short-term real interest rate r_t. In the recursive Euler equation (1), the term $E_t(x_{t+1})$ sums up the effect of the entire yield curve from tomorrow on $E_t(R_{t+1})$ (and of future disturbances) because aggregate demand tomorrow also depends on the entire yield curve from tomorrow on:

$$E_t(x_{t+1}) = -\sigma E_t(R_{t+1}) + E_t\left(\sum_{k=1}^{\infty} v_{t+k}\right). \tag{26}$$

Therefore assuming that the central bank observes $E_t(x_{t+1})$ is akin to assuming that the central bank faces no uncertainty as to the

firm last updated its information. The slope of the sticky-information Phillips curve is a function of the probability λ of renewing one's information set in a given period, $\kappa = \zeta\lambda/(1-\lambda)$.

[20] The issue does not arise with the New Classical Phillips curve because with the New Classical Phillips curve the expected output gap next period is always zero $E_t(x_{t+1}) = 0$.

effect of future short-term real interest rates on current aggregate demand and is only uncertain about the effect of the current short-term real interest rate r_t on current aggregate demand. Since aggregate demand depends on the entire yield curve R_t, assuming that the central bank is equally uncertain about the effect of the entire yield curve R_t is more natural, and we therefore focus on this case. It is handled easily by using the Euler equation in its iterated form (24). Again, the case where the central bank faces uncertainty only on the effect of the short-term interest rate r_t is qualitatively similar.[21]

Plugging in the iterated Euler equation (24) into the sticky-information Phillips curve (22) gives the relationship between the long-term interest rate R_t and inflation π_t at t:

$$\pi_t = -\phi R_t + \kappa E_t \left(\sum_{k=0}^{\infty} v_{t+k} \right) + \bar{E}_{t-1}(\pi_t + \zeta \Delta x_t), \tag{27}$$

where $\phi = \kappa\sigma$. The problem of the central bank under discretion at t is to pick the short-term interest rate $r_t = R_t - E_t(R_{t+1})$ that minimizes the loss $\mathcal{L}_t = E_t^*((\pi_t - \pi^*)^2)$ subject to the constraint (27). Because it acts under discretion, it takes future policies $E_t(R_{t+1})$ as given. As a result, the problem can be phrased equivalently as the one of choosing the long-term interest rate R_t that minimizes the loss $\mathcal{L}_t$.

3.2 The Attenuation Principle in a Dynamic Setup

The central bank still faces the same trade-off as under the New Classical Phillips curve, between bringing its best expectations of

[21] An additional mechanism appears in this case, which tends to make the policy response more front-loaded. In response to a fall in the natural rate, the central bank can expect it will attenuate the decrease in the short-term policy rate tomorrow. As a result, the output gap today becomes more negative and inflation today falls more below target. The central bank is therefore more willing to decrease the short-term interest rate today in order to stabilize inflation today. It is forced into action even earlier on.

inflation on target and minimizing the (conditional) variance of inflation. Its loss function can be written

$$\mathcal{L}_t = (E_t^*(\pi_t) - \pi^*)^2 + Var_t^*(\pi_t), \tag{28}$$

$$\text{where } (E_t^*(\pi_t) - \pi^*)^2 = \left(-\bar{\phi}R_t + \kappa E_t\left(\sum_{k=0}^{\infty} v_{t+k}\right) + \bar{E}_{t-1}(\pi_t - \pi^* + \zeta\Delta x_t)\right)^2, \tag{29}$$

$$Var_t^*(\pi_t) = V_\phi R_t^2. \tag{30}$$

The long-term real interest rate that sets inflation on target—the one the central bank would set absent concerns over uncertainty—is

$$R_t^s = R_t^n + \frac{\bar{E}_{t-1}(\pi_t - \pi^* + \zeta\Delta x_t)}{\bar{\phi}}, \tag{31}$$

where we define again the natural rate in the average model $r_t^n = v_t/\bar{\sigma}$, and the long-term natural rate

$$R_t^n \equiv E_t\left(\sum_{k=0}^{\infty} r_{t+k}^n\right). \tag{32}$$

Without concerns over parameter uncertainty, the optimal discretionary policy is to track the natural rate, plus a corrective term if inflation expectations are not on target.[22] In doing so, it has the long-term rate track the long-term natural rate, plus a corrective term if inflation expectations are not on target.

When the central bank is uncertain about the impact of interest rates on inflation, it faces a trade-off between setting inflation on target on average and minimizing the variance of inflation (30). Taking the first-order condition to the loss function (28) gives that the central bank solves this trade-off by setting the policy rate such

[22] This optimal policy tracks the *long-term* natural rate. As a result, if monetary policy is expected not to track the short-term natural rate tomorrow, optimal discretionary policy requires to make the *short-term* rate today depart from the short-term natural rate today in order to align the long-term rate on the long-term natural rate.

that the long-term rate R_t is at a midpoint between the interest rate R_t^s that puts inflation on target on average, and the interest rate $R_t = 0$ that minimizes the variance of inflation:

$$R_t = \alpha R_t^s, \tag{33}$$

where α is still given by (14). Brainard's attenuation principle materializes once again as a bias toward the policy whose effects the central bank understands best: in our case, keeping long-term interest rates toward their steady-state value.

3.3 Acting Tomorrow Out of Not Acting Today

To assess whether—and for how long—the central bank can indeed attenuate its policy response, we need to solve for the private sector's expectations on which the policy rate (33) depends. In order to do so, Appendix E solves for the dynamics induced by this very policy.

Proposition 4 (Dynamics under the Sticky-Information Phillips Curve). *The dynamics of inflation and the output gap (in the average model) are determined by the system:*

$$x_t = \bar{\sigma} R_t^n - \frac{\alpha}{(1-\alpha)\kappa}(\pi_t - \pi^*), \tag{34}$$

$$\pi_t = \kappa x_t + \bar{E}_{t-1}(\pi_t + \zeta \Delta x_t). \tag{35}$$

We consider the response of the economy (34)–(35) to a persistent fall in the natural rate. We assume the shocks to the natural rate follow an AR(1):

$$r_t^n = \rho r_{t-1}^n + \eta_t, \tag{36}$$

in which case the long-term natural rate $R_t^n = r_t^n/(1-\rho)$ also does. We calibrate the model at a quarterly frequency, as follows. Following Mankiw and Reis (2002), we set the slope of the SRAS to $\zeta = 0.1$, and the frequency of renewing one's information to once a year, $\lambda = 0.25$. This gives a slope of the Phillips curve equal to $\kappa = \zeta\lambda/(1-\lambda) = 0.033$. We set the intertemporal elasticity of substitution to $\bar{\sigma} = 1$. We assume that the uncertainty of the central bank is such that it attenuates its response by a quarter, $\alpha = 0.75$. We set the persistence of the shocks to $\rho = 0.95$.

Figure 2. IRF to a Fall in the Natural Rate under the Sticky-Information Phillips Curve

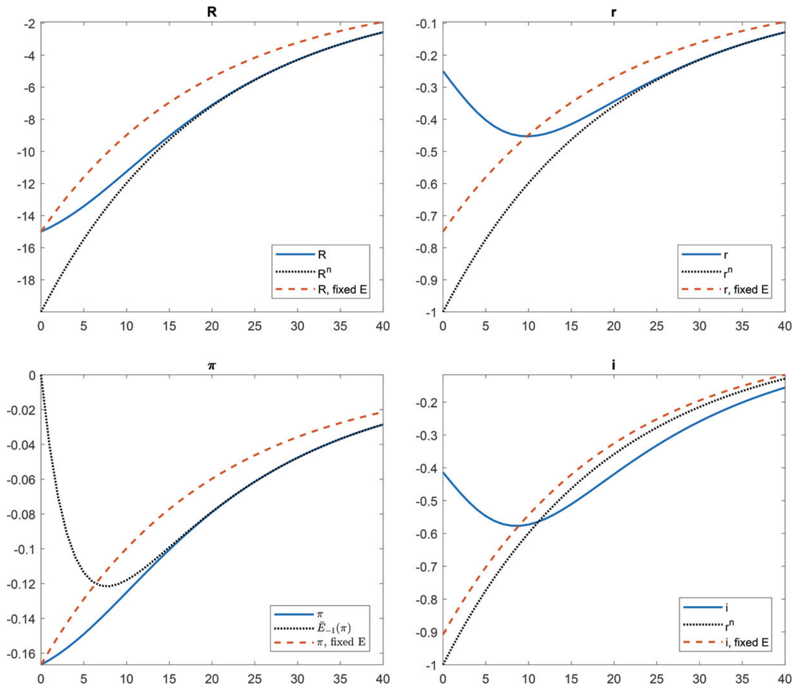

Note: The figure gives the impulse response functions (IRFs) of the long-term real interest rate R, the short-term real interest rate r, inflation π, and the short-term nominal interest rate i to a 1 percentage point decrease in the real natural rate of interest. The dashed lines give the IRFs in the counterfactual case where expectations of inflation and output gap growth $\bar{E}_{t-1}(\pi_t + \zeta \Delta x_t)$ remain constant at π^*. The horizon is expressed in quarters. The IRFs are plotted under the following calibration. The intertemporal elasticity of substitution is $\bar{\sigma} = 1$; the probability of renewing one's information set in the quarter is $\lambda = 0.25$; the slope of the short-run aggregate supply relationship is $\zeta = 0.1$. It implies a slope of the Phillips curve $\kappa = \zeta \times \lambda/(1-\lambda) \backsimeq 0.033$. The uncertainty V_ϕ is such that the central bank attenuates its action by a quarter, $\alpha = 0.75$. The autoregressive root of the AR(1) shock process is $\rho = 0.95$.

We consider a fall of the natural rate by 1 percentage point on impact. We solve for the impulse response function (IRF) using the method of undetermined coefficients, as detailed in Appendix E. Figure 2 gives the responses of interest rates and inflation. As can be seen through the dotted line of the bottom-left panel, on impact inflation expectations stay close to target because most private agents do not notice the shock. As a result, as shown on the upper-left panel, on impact the central bank is able to decrease the long-term real interest rate R (in plain line) by less than the fall in the natural long-term interest rate R^n (in dotted line). Because the long-term interest rate is above its natural level, inflation falls below target, as can be seen through the plain line of the bottom-left panel. As the private sector gradually realizes the fall in inflation, inflation (in plain line) is gradually pushed down below the path it would have taken had expectations stayed fixed (in dashed line). It forces the central back to decrease rates further. Ultimately, the real long-term interest rate ends up tracking the fall in the natural long-term rate as much as if the central bank had not been willing to attenuate policy and had simply tracked the natural rate from the start. Even as the real rate converges to the natural rate, however, inflation remains below target whereas it would have remained on target if the central bank had not been concerned with uncertainty and had tracked the natural rate from the start.[23]

On Figure 2, although the fall in the real rate never exceeds the fall in the natural rate, because of the fall in inflation expectations the *nominal* rate ends up falling by more that it would have absent concerns over uncertainty. Taking into account the effective lower bound on interest rate policies, a central bank subject to the cautiousness bias can therefore find itself up against the ELB even though it would not have if it had not tried to attenuate policy early on.

4. The Cautiousness Bias with the New Keynesian Phillips Curve

Both the New Classical Phillips curve and the sticky-information Phillips curve make past expectations of present inflation the

[23]Notice that the decrease in the *short-term* real interest rate r is initially even more attenuated than if inflation expectations stayed anchored, as can be seen on the top-right panel. This is because, as private agents anticipate that the central bank will be forced to decrease rates further in the future, initially the central bank faces a lower yield curve. As anticipations of future short-term rates are low, the central bank has less of an incentive to decrease present short-term rates. Ultimately, however, the short-term real interest rate ends up tracking the natural short-term rate as well.

relevant inflation expectations. Does the cautiousness bias depend on this form of sluggish expectations in the Phillips curve? We show it does not: The cautiousness bias applies equally to the forward-looking New Keynesian Phillips curve (NKPC). The dynamics of events under the NKPC differ from those under the sticky-information Phillips curve. Because the NKPC is not based on the assumption that it takes time for agents to incorporate new information, inflation expectations are not sluggish. Monetary policy is not progressively forced into action as inflation expectations progressively adjust. Instead, inflation expectations respond strongly on impact, and monetary policy is forced into action on impact.

4.1 The Problem of the Central Bank

We assume the supply side of the economy is captured by a standard New Keynesian Phillips curve, where inflation expectations enter as present expectations of future inflation:

$$\pi_t = \kappa x_t + \beta E_t(\pi_{t+1}), \tag{37}$$

where $\beta \in (0,1)$ is the discount factor. We keep the Euler equation unchanged, again in its iterated form (24). Because the New Keynesian Phillips curve is slightly non-vertical in the long-run $\beta < 1$, we focus on the case of a zero-inflation target $\pi^* = 0$ in order to abstract from a desire to exploit a long-run trade-off between inflation and output. Plugging the iterated Euler equation (24) into the New Keynesian Phillips curve (37) gives the relationship between the short-term interest rate r_t chosen at t and inflation π_t at t:

$$\pi_t = -\phi\left(r_t + E_t(R_{t+1})\right) + \kappa E_t\left(\sum_{k=0}^{\infty} v_{t+k}\right) + \beta E_t(\pi_{t+1}), \tag{38}$$

where $\phi = \kappa\sigma$. The problem of the central bank under discretion at t is to pick the short-term interest rate r_t that minimizes the loss $\mathcal{L}_t = E_t^*(\pi_t^2)$ subject to this constraint (38). Because it acts under discretion, it takes future policies $E_t(R_{t+1})$ as given.

4.2 The Cautiousness Bias with Forward-Looking Inflation Expectations

The derivation of the discretionary policy is similar to the case of the sticky-information Phillips curve. Appendix F shows that the Brainard principle still applies: the optimal discretionary policy is to attenuate the response of the long-term interest rate to changes in the long-term natural rate by the factor α,

$$R_t = \alpha \left(R_t^n + \frac{\beta E_t(\pi_{t+1})}{\bar{\phi}} \right). \tag{39}$$

Once again, however, the central bank acts less only for given inflation expectations, and acting less shifts inflation expectations adversely. Appendix F solves for inflation expectations and shows that the cautiousness bias applies equally to the New Keynesian Phillips curve.

PROPOSITION 5 (The Cautiousness Bias under the New Keynesian Phillips Curve). *Assume that the natural rate follows an AR(1) process (36). In equilibrium the long-term rate is*

$$R_t = \alpha \left(\frac{1}{1 - \beta(1-\alpha)\rho} \right) R_t^n. \tag{40}$$

While Brainard's attenuation principle leads the central bank to move rates less by a factor $\alpha < 1$, the reaction of inflation expectations forces the central bank to move rates more by a factor $1/(1 - \beta(1-\alpha)\rho) > 1$.

The timing in the manifestation of the cautiousness bias is specific to the New Keynesian Phillips curve, however. A well-known property of the NKPC is that it produces front-loaded dynamics, where shocks have their maximal effect on impact (Ball 1994; Mankiw and Reis 2002). This applies to the dynamics of the cautiousness bias. When a persistent negative shock to the natural rate hits, agents immediately factor in that the central bank will underreact to the fall in the natural rate, letting inflation fall below target. As a result, their present expectations of future inflation—the ones that enter the New Keynesian Phillips curve—immediately fall. In reaction, the central bank is immediately forced to decrease interest rates further in order to counteract the fall in inflation expectations. The central bank is forced into action as early as on impact, and from then on the change to the real interest rate fades away in proportion to the change in the natural rate.

Because the more persistent a shock is, the more inflation expectations react, persistent shocks force the central bank to react the most vigorously. Inflation expectations do not react at all if the shock is fully transitory $\rho = 0$, and react most when ρ tends to one. In our model, the central bank always act less than if it did not have concerns about uncertainty, if acting more is defined in terms of the *real* interest rate. Even for very persistent shocks, the overall multiplier $\alpha/(1-\beta(1-\alpha)\rho)$ always remains below one.[24] Appendix G shows that the same is true when the relevant Euler equation is the recursive Euler equation (1). Under either assumption on the Euler equation, however, for persistent enough shocks, the central bank moves *nominal* interest rates by more than it would have absent concerns about uncertainty, as shown in Appendices F and G.

That the persistence of shocks mitigates the policy attenuation called for by the Brainard principle is emphasized by Ferrero, Pietrunti, and Tiseno (2019). They consider a model where parameter uncertainty applies to the slope of the New Keynesian Phillips curve κ and show that the response of the central bank's nominal interest rate to cost-push shocks can move from attenuated to accentuated if shocks are persistent enough. We interpret their result through the lens of the cautiousness bias: a central bank concerned with Brainard uncertainty always wants to act less, but under discretion, the adverse reaction of the private sector's expectations can force it to act more.

5. Consequences of Generalizing the Least Uncertain Policy

In all the analysis so far, we have implicitly assumed that the interest rate policy whose effects are least uncertain is the steady-state interest rate. As a consequence, the cautiousness bias has not challenged

[24]Since the short-term rate is $r_t = \alpha\left(\frac{1}{1-\beta(1-\alpha)\rho}\right) r_t^n$, the same is true of the short-term interest rate.

the ability of the central bank to set average inflation on target, in contrast to the inflation bias. By generalizing the setup to allow for any policy to be the least uncertain, we show that the conflict between the desire to stabilize inflation and the desire to minimize inflation uncertainty can lead not only to an overreaction bias but also to an average bias, just as the conflict between the desire to stabilize inflation and the desire to stabilize output can lead to both an inflation bias and a stabilization bias (Svensson 1997). We also show that making the least uncertain policy a function of recently implemented interest rates does not change the equilibrium level of real interest rates, provided the private sector observes past policy rates and adjusts its expectations accordingly.

5.1 Allowing for Any Policy to Be the Least Uncertain

We consider again the framework of Section 2, where the economy's supply side is captured by the New Classical Phillips curve (2). In the setup of Section 2, the unconditional average inflation rate ends up equal to the inflation target, $E(\pi) = \pi^*$. This can be seen by taking the unconditional average of the expression for expected inflation (16), and using the fact that the natural rate is at its steady-state value on average, $E(\bar{r}^n) = 0$. While the cautiousness bias makes inflation depart from target by more, departures from target are symmetric above and below the target. Average inflation remains on target.[25]

The absence of an average bias is however only due to an implicit assumption embedded in Equation (7). The rationale for Brainard's attenuation principle is that, if the central bank is uncertain of the effects of its own action on inflation, uncertainty on inflation is all the greater the more it acts. *Acting more,* however, is only defined relative to a reference point. In Brainard's framework, this reference point is the policy whose effects on inflation are best understood $\underline{r}$, in the sense of minimizing the conditional variance of inflation (10). Equation (7) de facto assumes that the policy whose effects are best understood is keeping the interest rate around the steady-state value

[25]Of course, the fact that average inflation is on target also depends on the fact that we shut down the inflation bias by considering a central bank with a single inflation mandate. See Appendix D for the case where both the cautiousness and inflation biases are potentially at play.

of the natural interest rate, $\bar{r}^n = 0$. This is a justifiable assumption. Because the steady-state natural rate is the interest rate that has been most often implemented, it can be argued it is the interest rate on which most experience has been acquired.

But this is an assumption, and alternative ones are also defensible.[26] For instance, central banks can judge that they are more unsure of the interest rate pass-through when interest rates are low because of potential side effects, and more confident of it when interest rates are high. In a world where the secular, steady-state level of natural interest rates is low, this means that the least uncertain policy $\underline{r}$ is above steady state.

Besides, it can be argued that the policy whose effects are best understood varies across time. For instance, it can be at the level of the policy rates that have been *recently* implemented, which may not correspond to the steady state. Such an assumption is implicit in the reliance on the Brainard principle to justify gradualism, or the terminology of *"conservatism principle"* used by Blinder (1999). For instance, in March 2019, at the time of the quote by Mario Draghi mentioned in the introduction, the EONIA (euro overnight index average) had been at levels below 1.5 percent for more than 10 years, far below its long-term average. In this context, moving cautiously may be better interpreted as tilting nominal rates toward recent low levels, not toward their historical average—which would mean a sharp and sudden increase in rates.

In this section we generalize the setup of Section 2 to allow for the possibility that the policy whose effects are best understood $\underline{r}$ is not the steady-state level of natural rates. We allow for two generalizations. First, we allow $\underline{r}$ to be any arbitrary function of the past. For instance, it can be a weighted average of recent past real interest rates, as would be the case if the central bank is more confident

[26] The importance of the reference point of minimal uncertainty was not lost to Brainard, who cautioned that "some care must be used in interpreting [the attenuation principle]. The gap in this context is not the difference between what policy was 'last period' and what would be required to make the expected value of [the target variable] equal to [its target]. In the example we have used, the gap is the difference between the point where the variance of [the target variable] is least and [the policy instrument] required to give an expected value of [the target variable] equal to [its target]" (Brainard 1967). Sack (1998) and Wieland (2000) embed the Brainard principle in a dynamic learning model to capture the interpretation of the Brainard principle as a recommendation for gradualism.

about the effect of the policies it recently implemented. Second, its unconditional average can differ from the steady-state level of interest rates. We add the time index -1 to $\underline{r}_{-1}$ to emphasize that it is measurable with information available at $t-1$. Equation (7) is now

$$\pi = -\phi(r - \underline{r}_{-1}) - \bar{\phi}\underline{r}_{-1} + \varepsilon + E_{-1}(\pi). \tag{41}$$

The constant term $-\bar{\phi}\underline{r}_{-1}$ is necessary to guarantee that the central bank's average expectation of inflation across all the models it considers is correct, $E^*(\pi) = -\bar{\phi}r + \varepsilon + E_{-1}(\pi)$.

The program of the central bank is still to minimize the loss function (6), which can still be decomposed into a mean term (9) and a variance term. The only difference relative to Section 2 is that, by definition, the variance of inflation

$$Var^*(\pi) = V_\phi(r - \underline{r}_{-1})^2 \tag{42}$$

is now minimized for $r = \underline{r}_{-1}$. The central bank still solves the trade-off between its two objectives by setting the interest rate r to a midpoint between the interest rate (11) that minimizes the mean term (9) and the interest rate $\underline{r}$ that minimizes the variance term (42):

$$r = \alpha r^s + (1 - \alpha)\underline{r}_{-1}, \tag{43}$$

where r^s is still given by (11) and α is still given by (14).

5.2 Missing the Inflation Target on Average

Because the real interest rate does not track the natural rate one-for-one, inflation is not on target, which is anticipated by the private sector. Its rational expectations of inflation are

$$E_{-1}(\pi) = \pi^* + \left(\frac{1}{\alpha} - 1\right)\bar{\phi}(E_{-1}(\bar{r}^n) - \underline{r}_{-1}). \tag{44}$$

It follows in particular that unconditional average inflation is

$$E(\pi) = \pi^* - \left(\frac{1}{\alpha} - 1\right)\bar{\phi}E(\underline{r}_{-1}) \neq \pi^*. \tag{45}$$

Only when $E(\underline{r}_{-1}) = 0$ is average inflation on target π^*. If the central bank understands better how its policy affects inflation around a rate $\underline{r}_{-1} > 0$ which is on average greater that the steady-state natural rate, Brainard's attenuation principle provides an argument for setting the real interest rate above the natural interest rate on average. As a consequence, average inflation is below target π^* on average. Conversely, if the central bank understands better how the economy works around a rate $\underline{r}_{-1} < 0$ lower that the steady-state natural rate, average inflation is above target π^* on average.

Having average inflation not equal to π^* could be desirable, since it could come with the benefit of less uncertain inflation. In order to assess whether the desire to let average inflation depart from π^* constitutes a bias, Appendix B solves for the optimal average inflation rate under commitment to show that it does constitute a bias.[27]

PROPOSITION 6 (A Cautiousness Bias on Average Inflation). *Regardless of the value of $\underline{r}_{-1}$, the optimal average inflation rate is π^*. When $E(\underline{r}_{-1}) \neq 0$, concerns about uncertainty make average inflation depart from this optimal inflation target.*

Therefore, the departure of average inflation from π^* when $E(\underline{r}_{-1}) \neq 0$ is indeed a second manifestation of the cautiousness bias, this time on average inflation. Because the desire of the central bank to systematically tilt the real interest rate toward $E(\underline{r}_{-1})$ is fully anticipated by the private sector, in equilibrium the central bank fails to do so and the real interest rate is still (17). The desire to tilt the interest rate toward $E(\underline{r}_{-1})$ only results in an inflationary bias (if $E(\underline{r}_{-1}) < 0$) or deflationary bias (if $E(\underline{r}_{-1}) > 0$).

As a result, in the generic case when the best-understood policy is not on average the steady-state policy $E(\underline{r}_{-1}) \neq 0$, a discretionary central bank cannot follow a cautious strategy without failing to deliver on its inflation mandate on average inflation. Only in the particular case when the best-understood policy is on average the

[27] The inflation target is usually understood as the average inflation rate desired by the central bank, i.e., the optimal average inflation rate under commitment. Therefore, by referring to π^* as the central bank's inflation target, we have implicitly already assumed that π^* is the average inflation rate under commitment. Appendix B shows it is indeed the case. It is not simply by definition of π^*, however: what the reduced-form preferences (6) assume by construction is only that π^* minimizes the mean term of the loss function. With Brainard uncertainty, another average inflation rate could minimize the variance term if the central bank were able to affect the average real interest rate. But the Phillips-curve constraint (2) imposes that it cannot.

steady-state policy $E(\underline{r}_{-1}) = 0$ can a cautious discretionary central bank deliver an average inflation rate in line with its inflation mandate. In this case, inflation is still off target more often than if the central bank were not cautious, but symmetrically so.

5.3 No Impact on Equilibrium Interest Rates

While allowing the least uncertain policy $\underline{r}_{-1}$ to differ from the steady-state level of natural rates affects the equilibrium level of inflation, it makes no change to the equilibrium level of real interest rates. Indeed, plugging the private sector's expectations of inflation in equilibrium (44) into the expression for the policy rate chosen by the central bank (13), we find that the equilibrium real interest rate is still given by Equation (17).

Because under the New Classical Phillips curve the private sector observes $\underline{r}_{-1}$, any change that $\underline{r}_{-1}$ makes to the central bank's policy is anticipated by the private sector and fails to affect equilibrium real interest rates—it only affects inflation. In particular, time variations in $\underline{r}_{-1}$, as would occur if the central bank gradually adjusts its $\underline{r}_{-1}$ to recent levels of the real interest rate, do not generate any persistence in real interest rates, as Equation (43) could at first suggest.

6. Conclusion

Since Alan Blinder's book (Blinder 1999) made Brainard's attenuation principle widely known to central bankers, the economic literature has found several instances in which the Brainard principle proved not robust, with uncertainty calling instead for a more aggressive policy response. Preempting the literature to come—and because Brainard's original paper already emphasized cases in which uncertainty called for policy aggressiveness—Blinder commented: "My intuition tells me that [Brainard's principle] is more general—or at least more wise—in the real world than the mathematics will support."

In this paper, we made a distinct qualification to Brainard's attenuation principle. Focusing on situations in which uncertainty does rationalize policy attenuation, we showed that, when policy outcomes depend on the expectations of the private sector as in monetary policy, the desire to attenuate policy can backfire. It adversely shifts the private sector's inflation expectations, forcing the central bank to ultimately act by as much, but for worse outcomes. Our analysis does not conclude that uncertainty does not justify moving cautiously. But it emphasizes that central banks face a bias toward being *overly* cautious.

Appendix A. Microfoundations of the Shocks to the Euler Equation

In this appendix, we discuss the connection between the alternative representations of the shocks to the Euler equation: through shocks to the underlying fundamentals such as productivity a_t, through shocks to natural output y_t^n, through shocks to the variable $v_t = y_t^n - E_t(y_{t+1}^n) = -\sigma r_t^n$, or through shocks to the natural rate r_t^n. We show that when the central bank faces uncertainty about σ, only the first three are equivalent.

To do so, we first rederive the Euler equation (1) through a standard microfounded model with no capital and technology shocks as the only fundamental disturbance. A representative household consumes C_t, works L_t hours, and invests in B_t nominal riskless bonds in order to maximize intertemporal utility:

$$E_0 \sum_{t=0}^{\infty} \beta^t \left(\frac{C_t^{1-\frac{1}{\sigma}}}{1-\frac{1}{\sigma}} - \frac{L_t^{1+\psi}}{1+\psi} \right), \tag{A.1}$$

where σ is the intertemporal elasticity of substitution, ψ is the inverse of the Frisch elasticity of labor, and $\beta \in (0,1)$ is the household's discount factor.

A unit of consumption costs the price P_t. A unit of labor is paid the nominal wage W_t. The household chooses to invest B_t in nominal riskless bonds yielding the nominal interest rate I_t. The household receives nominal profits Ω_t from firms. It faces the flow budget constraint:

$$P_t C_t + B_t = W_t L_t + \Omega_t + I_t B_{t-1} \tag{A.2}$$

and an additional borrowing constraint that prevents it from entering Ponzi schemes. The household takes all prices as given.

Its optimal labor supply decision is to equate its marginal rate of substitution to the real wage $w_t = W_t/P_t$:

$$L_t^{\psi} C_t^{\frac{1}{\sigma}} = w_t. \tag{A.3}$$

The household's total consumption C_t results from the consumption C_t^i of a continuum $i \in [0, 1]$ of individual goods. We assume standard CES preferences with an elasticity of substitution across goods θ. The household's demand for good i is therefore

$$C_t^i = \left(\frac{P_t^i}{P_t}\right)^{-\theta} C_t. \tag{A.4}$$

Firm i produces good i using the production function:

$$Y_t^i = A_t (N_t^i)^{\alpha}, \tag{A.5}$$

where A_t is an aggregate productivity shock. Under flexible prices, firm i sets its price P_t^i to maximize present-period profits, internalizing the demand curve (A.4) it faces. It charges a markup over marginal costs:

$$P_t^i = \frac{\theta}{\theta - 1} P_t \frac{w_t}{A_t \alpha (N_t^i)^{\alpha-1}}. \tag{A.6}$$

Define natural output as the value of output in a flexible-price equilibrium. In a flexible-price equilibrium all firms set the same price and $A_t \alpha N_t^{\alpha-1} = \frac{\theta}{\theta-1} w_t$, where N_t is total labor demanded by firms. Combining the first-order conditions of the household and the firms and assuming that the goods and labor markets clear, $C_t = Y_t$ and $L_t = N_t$, gives natural output as a function of technology. Using lowercase variables to denote log-deviations from a steady state with $A = 1$, it is given by

$$y_t^n = \frac{\psi + 1}{1 + \psi + \left(\frac{1}{\sigma} - 1\right)\alpha} a_t. \tag{A.7}$$

Natural output is a function of the exogenous shock a_t. Note that natural output depends on the parameter σ, but only because the standard preferences we have assumed make σ parameterize both the intertemporal elasticity of substitution and the income effect on labor supply. Natural output depends on the strength of the income effect but not on the intertemporal elasticity of substitution. What we assume to be uncertain for the central bank is the intertemporal elasticity of substitution, not the strength of the income effect. Therefore, we assume that model uncertainty does not affect the central bank's expectations of natural output.[28]

The household's investment decision results in the Euler equation. Taking into account goods market clearing $C_t = Y_t$, it becomes in log-linear form:

$$y_t = -\sigma(i_t - E_t(\pi_{t+1})) + E_t(y_{t+1}). \tag{A.8}$$

The Euler equation applies in particular under flexible prices, in which case it residually gives the real interest rate in the flexible-price equilibrium, that is, the natural rate:

$$r_t^n = -\frac{1}{\sigma}(y_t^n - E_t(y_{t+1}^n)). \tag{A.9}$$

To rewrite the Euler equation in terms of difference with respect to natural output, define the output gap $x_t \equiv y_t - y_t^n$. The Euler equation is

$$x_t = -\sigma(i_t - E_t(\pi_{t+1})) + E_t(x_{t+1}) + v_t, \tag{A.10}$$

$$v_t \equiv -(y_t^n - E_t(y_{t+1}^n)). \tag{A.11}$$

The disturbance v_t is a function of natural output, therefore of the exogenous shocks. Since $r_t^n = \frac{1}{\sigma} v_t$, in models where agents face no model uncertainty it is customary to express the shock v_t in the Euler equation as exogenous variations in the natural rate:

$$x_t = -\sigma(i_t - E_t(\pi_{t+1}) - r_t^n) + E_t(x_{t+1}). \tag{A.12}$$

[28]There are several ways to make the parameter σ play the role of the elasticity of substitution only, and therefore to explicitly eliminate the dependence of natural output on σ. For instance, we could assume GHH preferences (Greenwood, Hercowitz, and Huffman 1988) to eliminate the income effect on labor supply. Alternatively, we could disentangle the intertemporal elasticity of substitution and the income effect on labor supply through Epstein-Zin preferences (Epstein and Zin 1989). However, in both cases the Euler equation would slightly differ from its canonical form. We thus stick to the standard preferences.

However, the two representations (A.10) and (A.12) are not equivalent when the central bank faces uncertainty over σ. The variables a_t, y_t^n, and v_t are independent of the intertemporal elasticity of substitution, while σ enters the definition (A.9) of r_t^n. Parameterizing the shocks to the Euler equation through an exogenous distribution for r_t^n in Equation (A.12) would spuriously make the effect of disturbances appear dependent on the value of σ, whereas σ multiplies r_t^n in Equation (A.12) only because r_t^n is divided by σ in definition (A.9). While the issue is irrelevant in models without parameter uncertainty, it matters when the central bank faces uncertainty on σ, because it changes the value of i_t for which the variance of x_t is minimal.

Appendix B. Proofs of Propositions 2 and 6: Optimal Policy Under Commitment

When the central bank sets policy under commitment, it understands the effect of its policy on the inflation expectations of the private sector. Because it understands that the private sector forms rational expectations in accordance with (7), it understands that in equilibrium its policy r must satisfy

$$E_{-1}(r) = E_{-1}(r^n). \tag{B.1}$$

Because the constraint (B.1) on policy rates spreads across realizations of ε, the program of the central bank no longer reduces to independent programs for each realization of ε. Instead, the central bank's faces one program at each information node of the private sector. At each node, it chooses the policy rates $r(\varepsilon)$ in each final realization of the shock, and the unique expectation of the private sector e to minimize:

$$\min_{(r(\varepsilon))_\varepsilon, e} E_{-1}(\mathcal{L}(\varepsilon))$$

$$= E_{-1}\left(\left(-\bar{\phi}(r(\varepsilon) - r^n(\varepsilon)) + e - \pi^*\right)^2 + V_\phi r(\varepsilon)^2\right), \tag{B.2}$$

$$\text{s.t. } E_{-1}(r(\varepsilon)) = E_{-1}(r^n(\varepsilon)). \tag{B.3}$$

Noting 2λ the Lagrange multiplier on the constraint, the first-order conditions (FOCs) are

$$\forall \varepsilon, /r(\varepsilon) : \bar{\phi}^2 \left(r(\varepsilon) - r^n(\varepsilon) - \frac{e - \pi^*}{\bar{\phi}} \right) + V_\phi r(\varepsilon) + \lambda = 0, \quad \text{(B.4)}$$

$$/e : e = E_{-1}\left(\bar{\phi}(r(\varepsilon) - r^n(\varepsilon)) + \pi^* \right). \quad \text{(B.5)}$$

Using the constraint (B.3), the FOC (B.5) gives $e = \pi^*$. Inflation expectations are always on target. Taking expectations E_{-1} of the FOC (B.4) and using the constraint (B.3) solves for λ. Substituting the expression for λ in the FOC solves for $r(\varepsilon)$:

$$r(\varepsilon) = E_{-1}(r^n(\varepsilon)) + \alpha(r^n(\varepsilon) - E_{-1}(r^n(\varepsilon))). \quad \text{(B.6)}$$

The policy rate takes the exact same value as under discretion (17). Substituting the policy rate (B.6) into Equation (7) gives the departure of inflation from target under commitment (19) in Proposition 2.

In Section 5, we generalize the setup of Section 2 by replacing Equation (7) with Equation (41). The optimal policy under commitment keeps setting expected inflation on target $E_{-1}(\pi) = \pi^*$ in this case, and therefore unconditional average inflation on target, $E(\pi) = \pi^*$. Indeed, the only difference with respect to the case $\underline{r}_{-1} = 0$ is to replace the first-order condition (B.4) with

$$\forall \varepsilon, /r(\varepsilon) : \bar{\phi}^2 \left(r(\varepsilon) - r^n(\varepsilon) - \frac{e - \pi^*}{\bar{\phi}} \right) + V_\phi(r(\varepsilon) - \underline{r}_{-1}) + \lambda = 0. \quad \text{(B.7)}$$

Equation (B.5) is unchanged. Using the constraint (B.3), it still gives $e = \pi^*$. Following the same steps as in the case $\underline{r} = 0$, one can also check that the policy rate still takes the value (B.6) in this generalized case, as it does under discretion. This proves Proposition 6.

Appendix C. Proof of Proposition 3: Optimal Discounting of Uncertainty Concerns

A central banker that puts a weight $\delta \geq 0$ on the variance term in the loss function (20) sets the interest rate to $r^d = \alpha r^s$, where

$\alpha = \frac{\bar{\phi}^2}{\bar{\phi}^2 + \delta V_\phi}$. As δ increases from zero to infinity, α decreases from 1 to 0: the more concerned he is about uncertainty, the less he reacts to shocks. We can therefore parameterize the central banker's type by how aggressively he reacts to shocks, as captured by α. An α-type central banker acting under discretion delivers an inflation rate of

$$E^*(\pi) - \pi^* = (1-\alpha)\bar{\phi}\left(\frac{1}{\alpha}E_{-1}(r^n) + (r^n - E_{-1}(r^n))\right). \quad \text{(C.1)}$$

Society compares these outcomes using its loss function with $\delta = 1$. The average loss generated by an α-type central banker is

$$E[\mathcal{L}(\varepsilon)] = Var(E^*(\pi) - \pi^*) + Var^*(\pi). \quad \text{(C.2)}$$

The two terms can be written as

$$Var(E^*(\pi) - \pi^*) = (1-\alpha)^2\bar{\phi}^2\left(\frac{1}{\alpha^2}V_E + V_U\right), \quad \text{(C.3)}$$

$$Var^*(\pi) = V_\phi(V_E + \alpha^2 V_U), \quad \text{(C.4)}$$

where $V_E \equiv Var(E_{-1}(r^n))$ is the variance of fluctuations in the natural rate that are expected by the private sector, and $V_U = Var(r^n - E_{-1}(r^n))$ is the variance of fluctuations in the natural rate that are unexpected by the private sector. Therefore, society wants to appoint the central banker whose α minimizes

$$E[\mathcal{L}(\varepsilon)] = \left((1-\alpha)^2\bar{\phi}^2 + \alpha^2 V_\phi\right)\left(\frac{1}{\alpha^2}V_E + V_U\right). \quad \text{(C.5)}$$

Taking the log and differentiating in α, the optimal α satisfies

$$\frac{\alpha V_\phi - (1-\alpha)\bar{\phi}^2}{(1-\alpha)^2\bar{\phi}^2 + \alpha^2 V_\phi} = \frac{1}{\alpha + \alpha^3 \frac{V_U}{V_E}}. \quad \text{(C.6)}$$

The right-hand-side term is decreasing from infinity to $V_E/(V_E+V_U)$ as α increases from 0 to 1. Define the left-hand-side term as the function f:

$$f(\alpha) = \frac{\alpha - \alpha^*}{(1-\alpha)^2\alpha^* + \alpha^2(1-\alpha^*)}, \quad \text{(C.7)}$$

where $\alpha^* = \frac{\bar{\phi}^2}{\bar{\phi}^2 + V_\phi}$ is the value of α of the central banker who has the same preferences as society, $\delta = 1$. The LHS f is negative for $\alpha < \alpha^*$, so it can only cross the RHS over $[\alpha^*, 1]$. The derivative of f has the sign of the quadratic polynomial $P(\alpha) = -\alpha^2 + 2\alpha^*\alpha + \alpha^*(1 - 2\alpha^*)$. The polynomial reaches its maximum at $\alpha = \alpha^*$ and has two real roots. If $\alpha^* > 1/2$, the larger root is greater than 1, so P is positive over $[\alpha^*, 1]$. It follows that f is increasing over $[\alpha^*, 1]$. There is a unique crossing of the RHS and LHS terms in Equation (C.6). If $\alpha^* < 1/2$, then the second root is smaller than 1, so f is increasing then decreasing over $[\alpha^*, 1]$. Yet, since $f(1) = 1 > V_E/(V_E + V_U)$, there is still a unique crossing of the RHS and LHS terms in Equation (C.6). In both cases, the two curves cross at a value greater than α^*, unless $V_E = 0$, in which case the RHS is constantly equal to zero and crosses the LHS at zero. An increase in $V_E/(V_E + V_U)$ causes the RHS to shift up: the optimal α therefore increases with the fraction of shocks that are expected by the private sector.

Appendix D. Proof in the Case of a Dual Mandate

We generalize the case of a single inflation mandate considered in the main text of Section 2 to allow for a dual objective to stabilize both inflation and the output gap. In doing so, we also allow for the possibility that the central bank seeks to set output above potential $x^* > 0$, which will result in an inflation bias. The present-period loss of the central bank is

$$\mathcal{L}(\varepsilon) = E^*((\pi - \pi^*)^2) + \lambda E^*((x - x^*)^2), \tag{D.1}$$

where λ is the preference weight of the central bank on stabilizing the output gap. The appearance of the real interest rate in the determination of the output gap,

$$x = -\sigma r + v, \tag{D.2}$$

and inflation,

$$\pi = \kappa(-\sigma r + v) + e(\pi), \tag{D.3}$$

are unchanged. The mean squared errors in the loss function (D.1) can still be decomposed into a term for squared distances to targets and a term for variances:

$$\mathcal{L}(\varepsilon) = \Big((E^*(\pi) - \pi^*)^2 + \lambda(E^*(x) - x^*)^2\Big) + \Big(Var^*(\pi) + \lambda Var^*(x)\Big), \tag{D.4}$$

$$\text{where } (E^*(\pi) - \pi^*)^2 + \lambda(E^*(x) - x^*)^2 = (-\kappa\bar{\sigma}r + \varepsilon + e(\pi) - \pi^*)^2 + \lambda(-\bar{\sigma}r + v - x^*)^2, \tag{D.5}$$

$$\text{and } Var^*(\pi) + \lambda Var^*(x) = (\kappa^2 + \lambda)V_\sigma r^2. \tag{D.6}$$

Denote r^s the interest rate that the central bank sets when it faces no model uncertainty, $V_\sigma = 0$. In this case the central bank can focus on minimizing the first term (D.5) in its loss function. It sets

$$r^s = r^n + \frac{\kappa}{\bar{\sigma}(\kappa^2 + \lambda)}(e(\pi) - \pi^*) - \frac{\lambda}{\bar{\sigma}(\kappa^2 + \lambda)}x^*. \tag{D.7}$$

A desire to stabilize the output gap $\lambda > 0$ changes the optimal discretionary policy relative to the case in which the central bank has no concerns about uncertainty in two ways. First, the central bank reacts less to departures of inflation expectations from target. This is regardless of whether the central bank seeks to set output above potential $x^* > 0$. Second, when $x^* > 0$ the central bank seeks to set the interest rate lower in order to set the output gap higher. This last feature of the discretionary policy results in Kydland and Prescott (1977) and Barro and Gordon (1983a, 1983b)'s inflation bias. Rational expectations of inflation are above target:

$$E_{-1}(\pi) = \pi^* + \frac{\lambda}{\kappa}x^* > \pi^*, \tag{D.8}$$

but the output gap is $x = 0$.[29]

When the central bank is uncertain about the impact of its rate decision on inflation and the output gap, $V_\sigma > 0$, the policy rate

[29] Monetary policy could surprise the private sector by responding to unexpected shocks to the natural interest rate, which would make the output gap depart from zero (although it would need to be zero on average). The central bank has no interest in doing so here, however, because there are no cost-push shocks.

r also affects the variance term (10) in the loss function. The central bank solves this problem by choosing a midpoint r between the optimal interest rate policy without model uncertainty r_s which minimizes the first term, and the steady-state interest rate $r = 0$ which minimizes the second term:

$$r = \alpha r^s, \tag{D.9}$$

where α is still given by (14).

We solve for the rational expectations that this policy generates. Injecting policy (D.9) into Equation (7), taking expectations E_{-1}, and imposing rational expectations $e(\pi) = E_{-1}(\pi)$ yields

$$E_{-1}(\pi) = \pi^* + \left(\frac{1}{\alpha} - 1\right) \frac{\bar{\sigma}(\kappa^2 + \lambda)}{\kappa} E_{-1}(r^n) + \frac{\lambda}{\kappa} x^*. \tag{D.10}$$

Inflation expectations can deviate from target for two reasons. First is the inflation bias, as noted in the case in which there are no concerns about uncertainty: a desire to set output above potential $\lambda > 0$ results in higher inflation expectations. Second is the cautiousness bias: a concern about parameter uncertainty $\alpha < 1$ results in lower (higher) inflation expectations when the natural rate is below (above) steady state. The generalization to the case of a dual mandate shows both that the cautiousness bias is robust to a dual mandate, and that the cautiousness bias and inflation bias arise from distinctly different perverse incentives.

Plugging expectations (D.10) into the optimal policy rate (D.9), the expression for the real interest rate is exactly the same as (17) under a single inflation mandate. In equilibrium the central bank attenuates its response only to unforeseen changes in the natural rate. It does so in exactly the same proportions as in the case of a single inflation mandate.

Appendix E. Proof of Proposition 4: Dynamics of the System and IRF in the Case of the Sticky-Information Phillips Curve

Injecting the solution (33) for the long-term interest rate into the Euler equation (24) (of the true model) gives the output gap as

$$x_t = \bar{\sigma}(1 - \alpha) R_t^n - \frac{\alpha}{\kappa} \bar{E}_{t-1}(\pi_t - \pi^* + \zeta \Delta x_t). \tag{E.1}$$

Using the sticky-information Phillips curve (22) to replace the last expectation term gives Equation (34). Equation (35) is simply the sticky-information Phillips curve.

We solve for the IRF for a natural rate shock using the method of undetermined coefficients, following Mankiw and Reis (2007). Under the assumption of an AR(1) process for the natural rate r_t^n, and denoting $\hat{\pi}_t = \pi_t - \pi^*$ the deviation of inflation from its target, the dynamic system is

$$x_t = \frac{\bar{\sigma}}{1-\rho} r_t^n - \frac{\alpha}{(1-\alpha)\kappa}\hat{\pi}_t, \tag{E.2}$$

$$\hat{\pi}_t = \kappa x_t + \bar{E}_{t-1}(\hat{\pi}_t + \zeta \Delta x_t). \tag{E.3}$$

Under the assumption of an AR(1) process for the natural rate r_t^n, its Wold decomposition is

$$r_t^n = \sum_{k=0}^{\infty} \rho^k \eta_{t-k}. \tag{E.4}$$

We look for a solution where $\hat{\pi}_t$ and x_t are functions of the fundamental shock only. We write their Wold decompositions:

$$\hat{\pi}_t = \sum_{k=0}^{\infty} \phi_k^{\pi} \eta_{t-k}, \tag{E.5}$$

$$x_t = \sum_{k=0}^{\infty} \phi_k^{x} \eta_{t-k}, \tag{E.6}$$

with coefficients $(\phi_k^\pi)_k$ and $(\phi_k^x)_k$ to be determined. To translate the dynamic system (E.2)–(E.3) into equations in ϕ_k^π and ϕ_k^x, note that

$$\begin{aligned}&\bar{E}_{t-1}(\pi_t + \zeta \Delta x_t)\\ &\quad = \sum_{k=1}^{\infty} \Big(1 - (1-\lambda)^k\Big)\Big(\phi_k^\pi + \zeta(\phi_k^x - \phi_{k-1}^x)\Big)\eta_{t-k}.\end{aligned} \tag{E.7}$$

The dynamic system (E.2)–(E.3) therefore can be written as

$$\sum_{k=0}^{\infty} \phi_k^x \eta_{t-k} = \frac{\bar{\sigma}}{1-\rho} \sum_{k=0}^{\infty} \rho^k \eta_{t-k} - \frac{\alpha}{(1-\alpha)\kappa} \sum_{k=0}^{\infty} \phi_k^\pi \eta_{t-k}, \tag{E.8}$$

$$\sum_{k=0}^{\infty} \phi_k^\pi \eta_{t-k} = \kappa \sum_{k=0}^{\infty} \phi_k^x \eta_{t-k} + \sum_{k=1}^{\infty} \left(1 - (1-\lambda)^k\right)\left(\phi_k^\pi + \zeta(\phi_k^x - \phi_{k-1}^x)\right)\eta_{t-k}. \tag{E.9}$$

Identifying the coefficients, it implies the following difference equations in $(\phi_k^\pi)_k$ and $(\phi_k^x)_k$:

$$\forall k \geq 0, \phi_k^x = \frac{\bar{\sigma}}{1-\rho}\rho^k - \frac{\alpha}{(1-\alpha)\kappa}\phi_k^\pi, \tag{E.10}$$

$$\forall k \geq 1, \phi_k^\pi = \kappa\phi_k^x + \left(1 - (1-\lambda)^k\right)\left(\phi_k^\pi + \zeta(\phi_k^x - \phi_{k-1}^x)\right), \tag{E.11}$$

$$\text{for } k = 0, \phi_0^\pi = \kappa\phi_0^x. \tag{E.12}$$

Using Equation (E.11) to eliminate ϕ_k^π in (E.10) gives the following first-order difference equation in ϕ_k^x:

$$\forall k \geq 1, \left((1-\lambda)^k + \frac{\alpha}{1-\alpha}\left(1 + \frac{\zeta}{\kappa}\left(1 - (1-\lambda)^k\right)\right)\right)\phi_k^x = \left(\frac{\alpha\zeta}{(1-\alpha)\kappa}\left(1 - (1-\lambda)^k\right)\right)\phi_{k-1}^x + \frac{\bar{\sigma}}{1-\rho}\rho^k(1-\lambda)^k. \tag{E.13}$$

This gives ϕ_k^x as a function of ϕ_{k-1}^x. We obtain the entire sequence of $(\phi_k^x)_k$ from the initial condition $\phi_0^x = \frac{(1-\alpha)\bar{\sigma}}{1-\rho}$. We then recover the solution for inflation from (E.11). The solutions for the interest rates follow.

Appendix F. Proof of Proposition 5: Derivation in the Case of the NKPC

Since the central bank under discretion takes the future policy rates $E_t(R_{t+1})$ as given, the central bank equivalently chooses the long-term real interest rate $R_t = r_t + E_t(R_{t+1})$. The long-term real interest rate that sets inflation on target is

$$R_t^s = R_t^n + \frac{\beta E_t(\pi_{t+1})}{\bar{\phi}}. \tag{F.1}$$

The long-term real interest rate that minimizes the within variance of inflation is $R_t = 0$. Taking the first-order condition of the loss function $\mathcal{L}_t = E_t^*(\pi_t^2)$ where π_t satisfies (38) gives that the central bank sets its policy rate so that the long-term rate is equal to the weighted average (39) of these two rates. Injecting the long-term rate into (38) (for the true model) gives inflation as

$$\pi_t = (1-\alpha)\left(\bar{\phi} R_t^n + \beta E_t(\pi_{t+1})\right) \tag{F.2}$$

or, written with time-series polynomials,

$$\pi_t = \Big(I - \beta(1-\alpha)F\Big)^{-1}(1-\alpha)\bar{\phi} R_t^n, \tag{F.3}$$

where I is the identity polynomial and F is the forward polynomial.

Under the assumption that r_t^n (and therefore R_t^n) follows an AR(1) process (36), the solution to (F.2) is

$$\pi_t = \frac{(1-\alpha)\bar{\phi}}{1-\beta(1-\alpha)\rho} R_t^n. \tag{F.4}$$

This implies that the private sector forms expectations of inflation:

$$E_t(\pi_{t+1}) = \frac{(1-\alpha)\bar{\phi}}{1-\beta(1-\alpha)\rho} \rho R_t^n. \tag{F.5}$$

Injecting these inflation expectations into the solution for the long-run rate (39) gives (40). The short-term rate $r_t = R_t - E_t(R_{t+1})$ is similarly

$$r_t = \alpha \left(\frac{1}{1-\beta(1-\alpha)\rho}\right) r_t^n. \tag{F.6}$$

It follows that the ex ante nominal long-term interest rate $I_t = R_t + E_t\left(\sum_{k=0}^{\infty} \pi_{t+k+1}\right)$ is

$$I_t = \left(\frac{\alpha + (1-\alpha)\bar{\phi}\frac{\rho}{1-\rho}}{1-\beta(1-\alpha)\rho}\right) R_t^n, \tag{F.7}$$

whereas it is $I_t = R_t^n$ in the absence of concerns about uncertainty. The coefficient in front of R_t^n tends toward infinity as ρ tends toward 1. Therefore, for persistent enough shocks, the central bank ends up moving the nominal long-term interest rate by more than it would have in the absence of concerns about uncertainty. Since the nominal short-term interest rate $i_t = r_t + E_t(\pi_{t+1})$ is similarly

$$i_t = \left(\frac{\alpha + (1-\alpha)\bar{\phi}\frac{\rho}{1-\rho}}{1-\beta(1-\alpha)\rho} \right) r_t^n, \tag{F.8}$$

the same conclusion applies to the nominal short-term rate.

Appendix G. The Case of the NKPC and the Recursive Euler Equation

Plugging the recursive Euler equation (24) into the NKPC (37), the relationship between the short-term interest rate chosen at t and inflation at t is

$$\pi_t = -\phi r_t + \kappa E_t(x_{t+1}) + \kappa v_t + \beta E_t(\pi_{t+1}), \tag{G.1}$$

where $\phi = \kappa\sigma$. Because the central bank acts under discretion, it chooses r_t at t taking $E_t(x_{t+1})$ and $E_t(\pi_{t+1})$ as given. It does so to minimize the loss $\mathcal{L}_t = E_t^*(\pi_t^2)$. The interest rate that sets inflation on target is

$$r_t^s = r_t^n + \frac{E_t(x_{t+1})}{\bar{\sigma}} + \frac{\beta E_t(\pi_{t+1})}{\bar{\phi}}. \tag{G.2}$$

The interest rate that minimizes the within variance of inflation is $r_t = 0$. The central bank sets its policy rate to the weighted average of these two rates:

$$r_t = \alpha r_t^s. \tag{G.3}$$

Injecting this solution for the short-term interest rate into (G.1) and using the NKPC (37) to replace future output gap $E_t(x_{t+1})$ gives inflation (in the true model) as

$$\pi_t = (1-\alpha)\Big(\bar{\phi} r_t^n + (1+\beta)E_t(\pi_{t+1}) - \beta E_t(\pi_{t+2}) \Big). \tag{G.4}$$

This is a second-order difference equation in π_t. Noting F the forward time-series operator, the stationary solution is

$$\pi_t = \left[I - (1-\alpha)(1+\beta)F + (1-\alpha)\beta F^2\right]^{-1}(1-\alpha)\bar{\phi}r_t^n. \quad \text{(G.5)}$$

Under the assumption that r_t^n follows an AR(1) process (36), the solution to (G.5) is

$$\pi_t = \frac{(1-\alpha)\bar{\phi}}{1-(1-\alpha)(1+\beta)\rho + (1-\alpha)\beta\rho^2} r_t^n. \quad \text{(G.6)}$$

Expectations of inflation $E_t(\pi_{t+1})$ and of the output gap $E_t(x_{t+1})$ follow. Plugging them into the expression for the short-term interest-rate (G.3) gives the solution for the short-term interest rate:

$$r_t = \alpha\left(\frac{1}{1-(1-\alpha)\rho(1+\beta(1-\rho))}\right) r_t^n. \quad \text{(G.7)}$$

As in the case of the iterated Euler equation, Brainard's attenuation principle leads the central bank to move rates by less by a factor $\alpha < 1$, but the reaction of inflation expectations forces the central bank to move rates by more, this time by a factor $1/(1-(1-\alpha)\rho(1+\beta(1-\rho))) > 1$. Because the more persistent a shock is, the more inflation expectations react, persistent shocks force the central bank to react the most vigorously. The central bank varies the real interest rate by *more* than the natural rate if and only if the coefficient on r_t^n in (G.7) is greater than 1. This happens if and only if the second-order polynomial,

$$P(\rho) = \beta\rho^2 - (1+\beta)\rho + 1, \quad \text{(G.8)}$$

takes negative values. Because the roots of the polynomial are 1 and $1/\beta$, this never happens for $\rho \in [0,1]$. However, as in the case of the iterated Euler equation, the short-term nominal interest rate can vary more than if the central bank has no concerns over Brainard uncertainty. Indeed, the solution for the short-term nominal interest rate $i_t = r_t + E_t(\pi_{t+1})$ is

$$i_t = \left(\frac{\alpha + (1-\alpha)\bar{\phi}\rho}{1-(1-\alpha)\rho(1+\beta(1-\rho))}\right) r_t^n. \quad \text{(G.9)}$$

The coefficient in front of r_t^n tends toward $1 + \left(\frac{1-\alpha}{\alpha}\right)\bar{\phi} > 1$ as ρ tends toward 1. Therefore, for persistent enough shocks, the central bank ends up moving the nominal short-term interest rate by more than it would have in the absence of concerns about uncertainty.

References

Acharya, S., J. Bengui, K. Dogra, and S. L. Wee. 2019. "Slow Recoveries and Unemployment Traps: Monetary Policy in a Time of Hysteresis." Working Paper.

Adam, K., and M. Woodford. 2012. "Robustly Optimal Monetary Policy in a Microfounded New Keynesian Model." *Journal of Monetary Economics* 59 (5): 468–87.

Angeletos, G.-M., and J. La'O. 2010. "Noisy Business Cycles." *NBER Macroeconomics Annual 2009*, Vol. 24, ed. D. Acemoglu, K. Rogoff, and M. Woodford, 319–78. NBER, University of Chicago Press.

Ball, L. 1994. "Credible Disinflation with Staggered Price-Setting." *American Economic Review* 84 (1): 282–89.

Barlevy, G. 2011. "Robustness and Macroeconomic Policy." *Annual Review of Economics* 3: 1–24.

Barro, R., and D. Gordon. 1983a. "A Positive Theory of Monetary Policy in a Natural Rate Model." *Journal of Political Economy* 91 (4): 589–610.

———. 1983b. "Rules, Discretion and Reputation in a Model of Monetary Policy." *Journal of Monetary Economics* 12 (1): 101–21.

Bernanke, B. 2007. "Monetary Policy under Uncertainty." Speech given at the 32nd Annual Economic Policy Conference, Federal Reserve Bank of St. Louis, October 19.

Blinder, A. 1999. *Central Banking in Theory and in Practice*. MIT Press.

Brainard, W. 1967. "Uncertainty and the Effectiveness of Policy." *American Economic Review* 57 (2): 411–25.

Carney, M. 2017. "Lambda." Speech given at the London School of Economics, January 16.

Chow, G. 1973. "Effect of Uncertainty on Optimal Control Policies." *International Economic Review* 14 (3): 632–45.

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

Craine, R. 1979. "Optimal Monetary Policy with Uncertainty." *Journal of Economic Dynamics and Control* 1 (1): 59–83.

Cukierman, A. 1991. "Why Does the Fed Smooth Interest Rates?" In *Monetary Policy on the 75th Anniversary of the Federal Reserve System*, ed. M. Belongia, 111–47. Kluwer Academic Publishers.

Draghi, M. 2019. "Transcripts of the Press Conference of March 7, 2019." European Central Bank.

Eggertsson, G., and M. Woodford. 2003. "Zero Bound on Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity* 1 (Spring): 139–211.

Epstein, L., and S. Zin. 1989. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework." *Econometrica* 57 (4): 937–69.

Estrella, A., and F. Mishkin. 1999. "Rethinking the Role of NAIRU in Monetary Policy: Implications of Model Formulation and Uncertainty." In *Monetary Policy Rules,* ed. J. Taylor, 405–36. NBER, University of Chicago Press.

Favaretto, F., and D. Masciandaro. 2016. "Doves, Hawks and Pigeons: Behavioral Monetary Policy and Interest Rate Inertia." *Journal of Financial Stability* 27 (December): 50–58.

Ferrero, G., M. Pietrunti, and A. Tiseno. 2019. "Benefits of Gradualism or Costs of Inaction? Monetary Policy in Times of Uncertainty." Working Paper, Bank of Italy.

Giannoni, M. 2002. "Does Model Uncertainty Justify Caution? Robust Optimal Monetary Policy in a Forward-Looking Model." *Macroeconomic Dynamics* 6 (1): 111–44.

Greenwood, J., Z. Hercowitz, and G. Huffman. 1988. "Investment, Capacity Utilization and the Real Business Cycle." *American Economic Review* 78 (3): 402–17.

Kimura, T., and T. Kirozumi. 2007. "Optimal Monetary Policy in a Micro-founded Model with Parameter Uncertainty." *Journal of Economic Dynamics and Control* 31 (2): 399–431.

Kydland, F., and E. Prescott. 1977. "Rules Rather than Discretion: The Inconsistency of Optimal Plans." *Journal of Political Economy* 85 (3): 473–92.

Lucas, R. 1972. "Expectations and the Neutrality of Money." *Journal of Economic Theory* 4 (2): 103–24.

Mankiw, G., and R. Reis. 2002. "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics* 117 (4): 1295–1328.

———. 2007. "Sticky Information in General Equilibrium." *Journal of the European Economic Association* 5 (2–3): 603–13.

———. 2010. "Imperfect Information and Aggregate Supply." In *Handbook of Monetary Economics*, Vol. 3A, ed. B. Friedman and M. Woodford. Elsevier.

Onatski, A., and J. Stock. 2002. "Robust Monetary Policy under Model Uncertainty in a Small Model of the U.S. Economy." *Macroeconomic Dynamics* 6 (1): 85–110.

Onatski, A., and N. Williams. 2003. "Modeling Model Uncertainty." *European Economic Association* 1 (5): 1087–1122.

Powell, J. 2018. "Monetary Policy in a Changing Economy." Speech given at the Changing Market Structure and Implications for Monetary Policy Economic Symposium sponsored by the Federal Reserve Bank of Kansas City, Jackson Hole, Wyoming, August 24.

Praet, P. 2018. "Assessment of Quantitative Easing and Challenges of Policy Normalisation." Speech given at the ECB and Its Watchers XIX Conference, Frankfurt am Main, March 14.

Reinhart, V. 2003. "Making Monetary Policy in an Uncertain World." In *Monetary Policy and Uncertainty: Adapting to a Changing Economy,* Proceedings of an Economic Policy Symposium sponsored by the Federal Reserve Bank of Kansas City, Jackson Hole, Wyoming, August 28–30.

Riboni, A., and F. Ruge-Murcia. 2010. "Monetary Policy by Committee: Consensus, Chairman Dominance, or Simple Majority." *Quarterly Journal of Economics* 125 (1): 363–416.

Rogoff, K. 1985. "The Optimal Degree of Commitment to an Intermediate Monetary Target." *Quarterly Journal of Economics* 100 (4): 1169–89.

Rudebusch, G. 2001. "Is the Fed Too Timid? Monetary Policy in an Uncertain World." *Review of Economics and Statistics* 83 (2): 203–17.

Sack, B. 1998. "Uncertainty, Learning, and Gradual Monetary Policy." FEDS Working Paper No. 1998-34, Board of Governors of the Federal Reserve System.

———. 2000. "Does the Fed Act Gradually? A VAR Analysis." *Journal of Monetary Economics* 46 (1): 229–56.

Sack, B., and V. Wieland. 2000. "Interest-Rate Smoothing and Optimal Monetary Policy: A Review of Recent Empirical Evidence." *Journal of Economics and Business* 52 (1–2): 205–28.

Sargent, T. 1999. "Comment on "Policy Rules for Open Economies" by Laurence Ball." In *Monetary Policy Rules*, ed. J. Taylor, 144–54. NBER, University of Chicago Press.

Savage, L. 1954. *The Foundations of Statistics.* New York: Dover Publication Inc.

Söderström, U. 2002. "Monetary Policy with Uncertain Parameters." *Scandinavian Journal of Economics* 104 (1): 125–45.

Söderström, U., and K. Leitemo. 2008. "Robust Monetary Policy in the New-Keynesian Framework." *Macroeconomic Dynamics* 12 (S1): 126–35.

Spiegler, R. 2021. "A Simple Model of Monetary Policy under Phillips-Curve Causal Disagreements." Working Paper.

Stein, J., and A. Sunderam. 2018. "The Fed, the Bond Market, and Gradualism in Monetary Policy." *Journal of Finance* 73 (3): 1015–60.

Stock, J. 1999. "Comment on 'Policy Rules for Inflation Targeting' by Glenn Rudebusch and Lars Svensson." In *Monetary Policy Rules,* ed. J. Taylor, 253–59. NBER, University of Chicago Press.

Svensson, L. 1997. "Optimal Inflation Targets, 'Conservative' Central Banks, and Linear Inflation Contracts." *American Economic Review* 87 (1): 98–114.

———. 1999. "Inflation Targeting: Some Extensions." *Scandinavian Journal of Economics* 101 (3): 337–61.

Tetlow, R., and P. von zur Muehlen. 2001. "Robust Monetary Policy with Misspecified Models: Does Model Uncertainty Always Call for Attenuated Policy?" *Journal of Economic Dynamics and Control* 25 (6–7): 911–49.

Theil, H. 1957. "A Note on Certainty Equivalence in Dynamic Planning." *Econometrica* 25 (2): 346–49.

Tillmann, P. 2009. "Optimal Monetary Policy with an Uncertain Cost Channel." *Journal of Money, Credit and Banking* 41 (5): 885–906.

Villeroy de Galhau, F. 2018. "Monetary Policy in the Present Environment: Reconciling Clarity and Flexibility, as Well as Independence and Cooperation." Speech given at the Group of Thirty's annual International Banking Seminar, Bali, Indonesia, October 14.

Walsh, C. 2003. "Implications of a Changing Economic Structure for the Strategy of Monetary Policy." In *Monetary Policy and Uncertainty: Adapting to a Changing Economy,* Proceedings of an Economic Policy Symposium sponsored by the Federal Reserve Bank of Kansas City, Jackson Hole, Wyoming, August 28–30.

Wieland, V. 2000. "Monetary Policy, Parameter Uncertainty and Optimal Learning." *Journal of Monetary Economics* 46 (1): 199–228.

Williams, J. 2013. "A Defense of Moderation in Monetary Policy." *Journal of Macroeconomics* 38 (Part B): 137–50.

Woodford, M. 2003a. "Imperfect Common Knowledge and the Effects of Monetary Policy." In *Knowledge, Information and Expectations in Modern Macroeconomics*, ed. J. S. P. Aghion, R. Frydman, and M. Woodford. Princeton University Press.

———. 2003b. *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton University Press.

———. 2003c. "Optimal Interest-Rate Smoothing." *Review of Economic Studies* 70 (4): 861–86.

———. 2010. "Robustly Optimal Monetary Policy with Near-Rational Expectations." *American Economic Review* 100 (1): 274–303.