

A Simple Model of Monetary Policy under Phillips-Curve Causal Disagreements*

Ran Spiegler[†]

May 19, 2021

Abstract

I study a static textbook model of monetary policy and relax the conventional assumption that the private sector has rational expectations. Instead, the private sector forms inflation forecasts according to a misspecified subjective model that disagrees with the central bank's (true) model over the causal underpinnings of the Phillips Curve. Following the AI/Statistics literature on Bayesian Networks, I represent the private sector's model by a direct acyclic graph (DAG). I show that when the private sector's model reverses the direction of causality between inflation and output, the central bank's optimal policy can exhibit an attenuation effect that is sensitive to the noisiness of the true inflation-output equations.

*Financial support by ERC Advanced Investigator grant no. 692995 is gratefully acknowledged. I thank In-Koo Cho, Stephane Dupraz and Morten Ravn for helpful comments.

[†]Tel Aviv University, University College London and CFM. URL: <http://www.tau.ac.il/~rani>. E-mail: rani@tauex.tau.ac.il.

1 Introduction

Monetary policy is a classic example of an economic interaction in which policy makers' ability to achieve their objectives depends on the accuracy of agents' forecasts of certain economic variables. Consider a textbook static model of the type originated by Kydland and Prescott (1977) and Barro and Gordon (1983). In such a model, the central bank controls a policy variable that affects inflation. The private sector forms an inflation forecast, possibly after observing some signal regarding the central bank's information or its decision. Private-sector expectations are relevant because they affect the joint realization of inflation and output, via some kind of an "expectations-augmented" Phillips curve. It follows that monetary policy involves "*expectations management*". To quote Woodford (2003, p. 15):

“...successful monetary policy is not so much a matter of effective control of overnight interest rates as it is of shaping market expectations of the way in which interest rates, inflation and income are likely to evolve...”

Conventional models constrain the central bank's ability to manage expectations by assuming that the private sector has "rational expectations" - i.e., it fully understands the statistical regularities in its environment, and thus forms an unbiased inflation forecast conditional on its information. To put it differently, the private sector shares the central bank's (true) model of the macroeconomy and uses it to form its beliefs. In Evans and Honkapohja (2005), Thomas Sargent refers to this feature as "communism of models". In this paper, I revisit the textbook model and relax "model communism", by assuming that the private sector forms its forecast according to a different model than the central bank's (true) model. The key difference between the two models is in their *causal treatment of the Phillips Curve*.

The modern history of macroeconomic thought has included a number of forms of Phillips Curves, which differ in their dependent and R.H.S variables. Popular descriptions of the Phillips Curve often impose a causal interpretation on these forms. For instance, when the dependent variable is

inflation, the causal story is about how changes in unemployment affect inflation. When inflation expectations enter the R.H.S, they are presented as another cause of inflation. When the dependent variable is unemployment or output, the story is about how inflation (possibly in addition to inflationary expectations) affects real variables. Cogley and Sargent (2005), following King and Watson (1994), describe this as a distinction between two econometric identification procedures, which they term “Keynesian direction of fit” and “classical direction of fit”. Taking a broader perspective, Hoover (2001) interprets debates in the history of monetary theory (and other areas of macroeconomics) in terms of disagreements over causality.

Therefore, if we wish to model how the private sector forms inflation forecasts using some kind of a Phillips Curve, it seems reasonable to assume that it will involve a causal interpretation. Even if the true model is not fundamentally causal (in the sense that the variables are determined in equilibrium by some system of simultaneous equations with influences in both directions), popular conceptions of the Phillips Curve and how it is integrated with a broader macroeconomic model are likely to be causal in nature.

Following Spiegler (2016,2020a), I adopt the formalism of Bayesian networks (Cowell et al. (1999), Pearl (2009), Koller and Friedman (2009)) and represent the private sector’s causal model with a directed acyclic graph (DAG). Nodes in the graph represent economic variables and directed links represent postulated causal influences. The DAG can be interpreted in at least two ways. First, it may capture intuitive narratives about causal relations between the relevant variables. Second, it may represent an explicit formal model of the kind that professional forecasters sometimes employ, consisting of a recursive system of equations. The private sector fits its subjective model to the steady-state distribution, and then relies on this “estimated model” to produce an inflation forecast. When the model is misspecified, the inflation forecast systematically distorts the true expected inflation conditional on the private sector’s information.

I incorporate this model of private-sector expectations into a variant on the familiar static textbook model of central-bank policy. The central bank aims to minimize the expectation of a quadratic loss function of the devia-

tions of output and inflation from their exogenous targets. The output target is fixed, whereas the inflation target is random. The central bank chooses a policy instrument in response to its inflation target, trading off the deviation of inflation from its moving target against the deviation of output from its fixed target. A Phillips Curve links these two deviations. Output deviations depend on how the private sector's inflation forecast responds to its information. The key question is therefore how the private sector's narrative regarding the Phillips Curve affects the responsiveness of the private sector's inflation forecast, and in turn the responsiveness of the central bank to the inflation target.

Two qualitative insights emerge from the analysis. First, the private sector's wrong causal model can result in *rigid* inflation forecasts, which in turn leads to monetary-policy attenuation - that is, the central bank's equilibrium response to the inflation target is weaker than in the rational-expectations benchmark. Second, the *variance of the noise terms* in the equations for inflation and output (one of which is a Phillips Curve) can play an important role in the rigid-forecast effect, when the private sector's model reverses the direction of causality between these two variables.

When macroeconomists talk about “flatness” of the Phillips relation, there is ambiguity about whether this means a weak yet precise relation or a noisy relation. In a finite sample, a noisy Phillips relation will appear as a “cloud” of data points, such that the Phillips effect may end up being statistically insignificant - hence, effectively regarded in much the same way as a weak, precise relation. However, when the private sector has rational expectations, the two have very different effects: the variance of the noise terms is irrelevant for the private sector's forecast, hence also for the central bank's optimal policy. This is no longer the case when the private sector's forecasts are based on a misspecified model. Using examples, I show that the variance of the noise terms in the inflation-output equations has a subtle effect on the private sector's inflation forecast and consequently on central-bank policy.

Related literature

The Bayesian-network approach to modeling equilibrium non-rational expectations was presented in Spiegler (2016). In Spiegler (2020a), I presented a

simple monetary-policy example to illustrate the question of whether causal misperceptions can lead to systematically biased expectations. In this paper, I adhere to the standard linear-normal parameterization, such that the private sector’s expectations are correct on average. Therefore, systematically biased inflation forecasts are ruled out; the departure from rational expectations lies in the under-sensitivity of the private sector’s inflation forecast to its information.

The Game Theory literature contains a number of related approaches to equilibrium modeling with non-rational expectations, including analogy-based expectations (Jehiel (2005)), “cursed” beliefs (Eyster and Rabin (2005)) and Berk-Nash equilibrium (Esponda and Pouzo (2016)). Spiegler (2020b) describes the relation between these approaches and the Bayesian-network formalism.

Within the macroeconomic theory literature, there are a few precedents for modeling monetary policy when the rational-expectations assumption is relaxed. Evans and Honkapohja (2001) and Woodford (2013) review dynamic macroeconomic models in which agents form non-rational expectations, and explore implications for monetary policy. Orphanides and Williams (2007) and Garcia-Schmidt and Woodford (2015) are examples of exercises in this tradition. In particular, Orphanides and Williams (2007) study the implications of learning-based private-sector expectations for the structure of central-bank policy.

The most closely related equilibrium concept that is employed in the macroeconomic literature is known as “restricted perceptions equilibrium” (Evans and Honkapohja (2001)), which is based on a notion of coarse beliefs in the same spirit as Piccione and Rubinstein (2003) and Jehiel (2005). Sargent (2001), Cho et al. (2002) and Esponda and Pouzo (2016) study models in which it is the *central bank* that forms non-rational expectations, whereas the private sector is modeled conventionally.

The idea that real-life monetary policy may be insufficiently responsive to shocks has been discussed in the literature since at least Rudenbusch (2001). A number of explanations have been proposed. One of them, originating in Brainard (1967), is that policy attenuation is a response to uncertainty

about relevant parameters or to Knightian uncertainty regarding the central bank's model (see Tetlow and Von zur Muehlen (2001) and Dupraz et al. (2020) for critical discussions). This paper contributes to this literature by proposing a new possible motive for monetary-policy distortion: the private sector's misperception of the causal relation between inflation and output.

2 The Model

The following is a variant on a familiar textbook static model of monetary policy in the tradition of Kydland and Prescott (1977) and Barro and Gordon (1983). It most closely resembles Sargent (1999) and Athey et al. (2005).

All variables in the model get real values. A central bank observes an exogenous, normally distributed variable θ and chooses an action a . Subsequently, the private sector observes a Gaussian signal t of (θ, a) and forms an inflation forecast e . Following the realization of θ, a, t, e , actual inflation π and actual output y are realized according to the following pair of equations:

$$\begin{aligned}\pi &= a + \varepsilon \\ y &= \pi - \lambda e + \eta\end{aligned}$$

where $\lambda \in (0, 1)$ is a constant that represents the extent to which anticipated inflation offsets the real effect of actual inflation; and ε and η are statistically independent noise terms, $\varepsilon \sim N(0, \sigma_\varepsilon^2)$ and $\eta \sim N(0, \sigma_\eta^2)$.

The central bank's objective is to minimize a standard quadratic loss function:

$$L(\pi, y, \theta) = y^2 + (\pi - \theta)^2$$

Thus, the exogenous state variable θ is interpreted as the central bank's ideal inflation target. The central bank's strategy is a function $(a(\theta))_\theta$ that assigns an action to every realization of θ . The private sector's strategy is the forecast function $(e(t))_t$ that assigns an inflation forecast to every signal t .

Given the central bank's strategy and private sector's forecast function,

we can define a joint distribution p over all six variables θ, a, t, e, y, π . The exogenous components of p are the distribution over θ , the conditional distribution over t (given θ, a) and the conditional distribution of π and y (given θ, a, t, e). The endogenous components are the pair consisting of the central bank's strategy $(a(\theta))_\theta$ and the private sector's forecast $(e(t))_t$.

We now arrive at the key definition of how the private sector forms its inflation forecast. The private sector is endowed with a subjective causal model, which is represented by a *directed acyclic graph* (DAG) $G = (N, R)$, where N is a set of nodes and R is a set of directed links. Each node in N represents a variable. Some variables may be omitted from the causal model, but we assume that t, y and π are always represented. A link represents a perceived direct causal influence.

It will be convenient to enumerate the six variables x_1, \dots, x_6 , such that N is some subset of $\{1, \dots, 6\}$. Denote $x = (x_1, \dots, x_6)$. For every $M \subseteq \{1, \dots, 6\}$, $x_M = (x_i)_{i \in M}$ is the projection of x on the variables represented by M . Let $R(i)$ be the set of nodes j such that G includes the link $j \rightarrow i$. In other words, $R(i)$ represents the set of variables that are considered to be direct causes of the variable x_i . I sometimes use boldface notation of a variable to represent its label, such that for any variable z , its label is \mathbf{z} . For instance, when R postulates that y is a direct cause of π , we can rewrite this as $\mathbf{y} \in R(\boldsymbol{\pi})$.

To arrive at an inflation forecast, the private sector "estimates" its model by performing measurements that quantify the postulated causal links, and combining the results of these measurements in accordance with the causal model. In other words, the private sector *fits* the causal model G to the joint distribution p . Formally, the private sector's belief is defined as follows:

$$p_G(x_N) = \prod_{i \in N} p(x_i \mid x_{R(i)}) \quad (1)$$

This is a *Bayesian-network factorization formula*: it factorizes the objective joint distribution p (defined over x_N) into a product of conditional-probability terms. Each term is extracted from p , but the terms are put together in a way that may result in $p_G \neq p$. As long as p has full support, p_G is a well-defined distribution, which induces well-defined marginal and conditional

distributions. We say that a distribution p is consistent with G if $p_G = p$. Indeed, the objective joint distribution p is consistent with the following DAG, denoted G^* :

$$\begin{array}{ccccc} \theta & \rightarrow & a & \rightarrow & \pi \\ \downarrow & \swarrow & & & \downarrow \\ t & \rightarrow & e & \rightarrow & y \end{array}$$

One interpretation of p_G is that it is the limit belief of a discrete-time process of Bayesian updating, where the prior has full support over all distributions that are consistent with G , and the private sector observes an independent draw from p at every period (see Spiegel (2020b)). Another interpretation, which is appropriate when p is multivariate normal, is that G represents a recursive system of linear regression equations, where every term $p(x_i | x_{R(i)})$ in (1) corresponds to a linear regression equation for x_i , in which the set of regressors is $R(i)$. In this case, (1) defines a so-called *Gaussian Bayesian network* (Koller and Friedman (2009, Ch. 7)). Each equation is estimated via OLS against an arbitrarily large sample, and the estimated equations are put together in accordance with the assumption that the noise terms in the regression equations are independent. This assumption is wrong when G is a misspecified causal model.

Because the private sector perceives statistical regularities through the prism of an incorrect model, the subjective belief p_G may systematically distort the correlation structure of the actual distribution p .

The private sector uses its belief p_G to obtain the inflation forecast:

$$e = E_G(\pi | t) = \int_{\pi} \pi dp_G(\pi | t) \quad (2)$$

For instance, if G is

$$\begin{array}{ccccc} \theta & \rightarrow & a & \rightarrow & y \\ & \searrow & \downarrow & & \downarrow \\ & & t & & \pi \end{array}$$

then

$$E_G(\pi | t) = \int_{\theta, a, y, \pi} p(\theta, a | t) p(y | a) p(\pi | a) \pi$$

As we shall see, because G is a misspecified model, e may be sensitive to all features of the joint distribution of p , including those that would not matter under rational expectations.

Definition 1 (Equilibrium) *The pair $(a(\theta))_\theta$ and $(e(t))_t$ constitutes an equilibrium if: (i) for every θ , $a(\theta)$ maximizes the central bank's expected payoff given $(e(t))_t$; and (ii) $(e(t))_t$ is given by (2).*

Definition 2 (Linear equilibrium) *An equilibrium $(a(\theta))_\theta$ and $(e(t))_t$ is linear if $a(\theta)$ and $e(t)$ are linear functions.*

Throughout the paper, I restrict attention to linear equilibria. This ensures that the joint distribution p is multivariate normal.

Rational-expectations benchmark

As a benchmark, suppose that the private sector has rational expectations. That is, its subjective DAG is G^* itself. Furthermore, suppose it is fully informed of θ, a . Then, its inflation forecast conditional on θ, a is

$$e = E(\pi \mid \theta, a) = E(\pi \mid a) = a$$

Therefore,

$$y = a - \lambda a + \varepsilon + \eta$$

The central bank's expected cost from an action a given the realization of θ is:

$$(1 - \lambda)^2 a^2 + (a - \theta)^2 + 2\sigma_\varepsilon^2 + \sigma_\eta^2$$

Therefore, the central bank's optimal action given θ is

$$a^*(\theta) = \frac{1}{(1 - \lambda)^2 + 1} \theta$$

The solution $a^*(\theta)$ is not fully responsive to changes in θ . However, as λ tends to one - such that only unanticipated inflation matters for output - the central bank's policy approaches fully flexible targeting. Another key observation is that $a^*(\theta)$ is invariant to σ_ε^2 and σ_η^2 .

Discussion: Model-based forecasts

The private sector’s forecasting method consists of two steps: if first estimates a causal model and then uses the estimated model to generate an inflation forecast. The exact interpretation of this process depends on whether we think of private-sector agents as lay persons who reason about macroeconomic processes informally and intuitively, or as professional forecasters. Under the former interpretation, the idea is that agents have some qualitative causal perceptions of how macroeconomic variables relate to one another, and they interpret observational data in light of these perceptions.

As to the latter interpretation, it is clear that not all professional forecasting is model-based. However, models do play a role in forecasting. One virtue of models is that they ground the forecast in a “story” that can be communicated to other parties (e.g. see Edge et al. (2008)). Another advantage of a model is that it is a simple vehicle for making a multitude of conditional predictions. See, for example, the following quote from a recent speech by Stanley Fisher: “The economy is an extremely complicated mechanism, and every macroeconomic model is a vast simplification of reality. . . the large scale of FRB/US is an advantage in that it can perform a wide variety of computational “what if” experiments.”¹ For a critical discussion of theory-based forecasting, see Giacomini (2015).

Thus, we should not think of the private sector as being exclusively interested in predicting inflation - otherwise, there would be no need to rely on a model; the private sector could simply measure the steady-state conditional expectation $E(\pi | t)$. Instead, the private sector employs a model as a multi-purpose vehicle for making sense of statistical regularities in the macroeconomic environment, and for making conditional predictions as the occasion arises. The model is a knowing simplification of a complex environment, and therefore the private sector expects it to get some things wrong, and would be unfazed if some of the model’s assumptions and predictions turned out to be incorrect (defending this stance with familiar statements like “every model is wrong” or “it takes a model to beat a model”).

¹See <https://www.federalreserve.gov/newsevents/speech/fischer20170211a.htm>.

3 The Basic Result

The model assumes that the private sector has a correct causal model of the joint behavior of the variables θ, a, t, e . The private sector's error is that it gets the direct causes of π or y wrong. In the true model $G^* = (N^*, R^*)$, $\mathbf{y} \notin R^*(\boldsymbol{\pi})$ - that is, output is not a direct cause of inflation, because causality runs in the opposite direction. Our first result establishes that if the private sector's causal model G shares this property, its inflation forecast in a linear equilibrium is unaffected by σ_ε^2 and σ_η^2 . Moreover, the equilibrium can be "rationalized": it can be obtained by assuming that $G = G^*$ while modifying the private sector's signal function, given by $(p(t | \theta, a))$.

Proposition 1 *Suppose $G = (N, R)$ satisfies $\mathbf{y} \notin R(\boldsymbol{\pi})$. Then:*

- (i) *Any linear equilibrium is invariant to $(\sigma_\varepsilon^2, \sigma_\eta^2)$.*
- (ii) *Any linear equilibrium can be obtained from a specification of the primitives in which the private sector's causal model is G^* , and the private sector is either fully informed or entirely uninformed of (θ, a) .*

Proof. (i) By assumption, G coincides with G^* over the nodes θ, a, t, e . Therefore,

$$\begin{aligned} E_G(\pi | t) &= \sum_{\theta, a, e} p_G(\theta, a, e | t) \sum_{\pi} p_G(\pi | \theta, a, t, e) \pi \\ &= \sum_{\theta, a, e} p_G(\theta, a, e | t) \sum_{\pi} p_G(\pi | x_{R(\boldsymbol{\pi})}) \pi \\ &= \sum_{\theta, a, e} p(\theta, a, e | t) \sum_{\pi} p_G(\pi | x_{R(\boldsymbol{\pi})}) \pi \end{aligned}$$

By assumption, $R(\boldsymbol{\pi}) \subseteq \{\boldsymbol{\theta}, \mathbf{a}, \mathbf{t}, \mathbf{e}\}$. Therefore,

$$p_G(\pi | x_{R(\boldsymbol{\pi})}) = \sum_{a'} p(a' | x_{R(\boldsymbol{\pi})}) p(\pi | a')$$

such that

$$E_G(\pi | t) = \sum_{x_{R(\boldsymbol{\pi})}} p(x_{R(\boldsymbol{\pi})} | t) \sum_{a'} p(a' | x_{R(\boldsymbol{\pi})}) \sum_{\pi} p(\pi | a') \pi \quad (3)$$

Since

$$\sum_{\pi} p(\pi | a') \pi = E(a' + \varepsilon) = a'$$

we obtain

$$E_G(\pi | t) = \sum_{x_{R(\pi)}} p(x_{R(\pi)} | t) \sum_{a'} p(a' | x_{R(\pi)}) a'$$

The terms $p(x_{R(\pi)} | t)$ and $p(a' | x_{R(\pi)})$ only derive from $(p(\theta, a, t, e))$, and are therefore invariant to $(\sigma_{\varepsilon}^2, \sigma_{\eta}^2)$.

(ii) Once again, recall that by assumption, $R(\pi) \subseteq \{\boldsymbol{\theta}, \mathbf{a}, \mathbf{t}, \mathbf{e}\}$. In a linear equilibrium, a is a deterministic function of θ , and e is a deterministic function of t . Therefore, if $R(\pi)$ includes \mathbf{a} ($\boldsymbol{\theta}$), it is immaterial whether it also includes $\boldsymbol{\theta}$ (\mathbf{a}). Likewise, if $R(\pi)$ includes \mathbf{t} (\mathbf{e}), it is immaterial whether it also includes \mathbf{e} (\mathbf{t}).

It follows that we only need to examine three cases. First, suppose that $\mathbf{a} \in R(\pi)$ (or, equivalently, $\boldsymbol{\theta} \in R(\pi)$). Then,

$$\begin{aligned} E_G(\pi | t) &= \sum_{\theta, a} p(\theta, a, e | t) E(\pi | x_{R(\pi)}) \\ &= \sum_{\theta, a} p(\theta, a, e | t) E(\pi | a) \\ &= \sum_a p(a | t) E(\pi | a) \\ &= E(\pi | t) \end{aligned}$$

where the second equality follows from the fact that under the objective distribution, $\pi \perp (\theta, t, e) | a$. It follows that if $\mathbf{a} \in R(\pi)$, the private sector's inflation forecast is consistent with rational expectations, even if we do not change its signal function.

Second, suppose that $\{\mathbf{a}, \boldsymbol{\theta}\} \cap R(\pi)$ is empty, but $\mathbf{t} \in R(\pi)$ (or, equivalently, $\mathbf{e} \in R(\pi)$). Then, by definition,

$$E_G(\pi | t) = E(\pi | t)$$

Then, as in the first case, the private sector's inflation forecast is consistent with rational expectations, even if we do not change its signal function.

The only remaining case is that $R(\boldsymbol{\pi})$ is empty. In this case,

$$E_G(\pi | t) = E(\pi)$$

This inflation forecast can be generated in a rational-expectations model, in which the private sector is entirely uninformed of θ, a - i.e., we modify $p(t | \theta, a)$ into $p'(t | \theta, a)$ that is constant in θ, a . ■

This result establishes that nothing “interesting” happens in this model when the private sector's model does not invert the causal link between inflation and output. In fact, the only departure from the rational-expectations benchmark occurs when the private sector's model postulates that none of the other relevant variables are a direct cause of inflation. In this case, the private sector acts as if it is entirely uninformed of θ and a . It is easy to show that in the unique linear equilibrium, $e = 0$ with certainty, while the central bank's policy is

$$a(\theta) = \frac{\theta}{2}$$

That is, the private sector's expectations are entirely rigid, as in the rational-expectations benchmark with $\lambda = 0$. This results in an overly rigid central-bank policy.

4 Examples

In this section I analyze in detail two examples in which the private sector's causal model violates the condition in Proposition 1. As a result, the noise terms in the equations for inflation and output play a non-trivial role in the private sector's inflation forecast, and therefore also in the central bank's equilibrium policy.

In both examples, I assume that $t = a$ - i.e., the private sector observes a before forming its inflation forecast. (The private sector does not observe θ , but this is immaterial.) Since t coincides with a , we can remove it from the

model, such that G^* can be written as:

$$\begin{array}{ccccc} \theta & \rightarrow & a & \rightarrow & \pi \\ & & \downarrow & & \downarrow \\ & & e & \rightarrow & y \end{array}$$

That is, the central bank's action a is an immediate consequence of the exogenous variable θ ; actual inflation π and the private sector's inflation forecast e are conditionally independent consequences of a ; and these two variables are the direct causes of real output y .

We will now analyze two specifications of the private sector's DAG G .

4.1 Case 1: $G : \theta \rightarrow a \rightarrow \pi \leftarrow y$

This DAG tells a causal story, according to which inflation is a consequence of two independent causes: real output and the central bank's policy. This causal model postulates absolute *monetary neutrality*: it admits no path of causal links from a to y . Thus, the private sector's causal model distorts the true model G^* in two ways: it omits the expectations variable e and inverts the causal link between inflation and output.

The private sector's inflation forecast after observing the central bank's action a is

$$E_G(\pi | a) = \int_{\pi} \pi p_G(\pi | a) = \int_{\pi} \int_y \pi p(y) p(\pi | a, y) \quad (4)$$

Compare this with the following way to write the rational-expectations inflation forecast:

$$E(\pi | a) = \int_{\pi} \int_y \pi p(y | a) p(\pi | a, y) \quad (5)$$

The discrepancy arises because $p_G(\pi | a)$ involves an implicit expectation over y *without* conditioning on a . (Note that since $\pi, y \perp \theta | a$ according to both G^* and G , θ does not feature in either of the two expressions.) Because the term $p(y)$ in (4) is actually *not* independent of a , a change in the central bank's strategy can lead to a change in $E_G(\pi | a)$.

Proposition 2 *There is a unique linear equilibrium, given by*

$$\begin{aligned} e(a) &= \frac{1 - \beta}{1 - \beta\lambda} a \\ a(\theta) &= \frac{1}{\left(\frac{1-\lambda}{1-\beta\lambda}\right)^2 + 1} \theta \end{aligned} \tag{6}$$

where

$$\beta = \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \sigma_\eta^2} \tag{7}$$

Proof. Let us begin by deriving a formula for the private sector's inflation forecast. Denote

$$e(a) = E_G(\pi | a) = \int_y p(y) E(\pi | a, y)$$

Since $\pi = a + \varepsilon$,

$$E(\pi | a, y) = a + E(\varepsilon | a, y)$$

By the Phillips Curve,

$$\varepsilon + \eta = y - a + \lambda e(a)$$

For given a and y , the R.H.S is a constant, whereas the L.H.S is a sum of two independent variables that are normally distributed with mean zero. Therefore, to calculate $E(\varepsilon | a, y)$, we can apply the standard formula for $E(X | X + Z)$ when X and Z are independent normal variables, and obtain

$$E(\varepsilon | a, y) = \beta(y - a + \lambda e(a))$$

where

$$\beta = \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + \sigma_\eta^2}$$

We can now write

$$e(a) = \int_y p(y) [a + \beta y - \beta a + \beta \lambda e(a)] = (1 - \beta)a + \beta \lambda e(a) + \beta E(y)$$

Since $\pi = a + \varepsilon$ and $E(\varepsilon) = 0$, $E(\pi) = E(a)$. Plugging the Phillips curve, we

obtain

$$E(y) = E(a) - \lambda E(e(a))$$

and therefore,

$$e(a) = (1 - \beta)a + \beta\lambda e(a) + \beta[E(a) - \lambda E(e(a))]$$

This functional equation defines $e(a)$. Taking expectations, we obtain

$$\begin{aligned} E(e(a)) &= (1 - \beta)E(a) + \beta\lambda E(e(a)) + \beta E(a) - \beta\lambda E(e(a)) \\ &= E(a) = E(\pi) \end{aligned}$$

Plug this into the expression for $e(a)$, and get:

$$e(a) = \frac{1 - \beta}{1 - \beta\lambda}a + \frac{\beta - \beta\lambda}{1 - \beta\lambda}E(a)$$

Plugging the expression for β , we obtain

$$e(a) = \frac{\sigma_\eta^2}{\sigma_\eta^2 + (1 - \lambda)\sigma_\varepsilon^2}a + \frac{(1 - \lambda)\sigma_\varepsilon^2}{\sigma_\eta^2 + (1 - \lambda)\sigma_\varepsilon^2}E(a)$$

Plugging the expression we obtained for $e(a)$, the expected output conditional on a is

$$E(y | a) = \frac{1 - \lambda}{1 - \beta\lambda}a - \lambda \frac{\beta - \beta\lambda}{1 - \beta\lambda}E(a)$$

We can now write down the central bank's expected payoff from a given θ as follows:

$$\left[\frac{1 - \lambda}{1 - \beta\lambda}a - \lambda \frac{\beta - \beta\lambda}{1 - \beta\lambda}E(a) \right]^2 + (a - \theta)^2$$

Our definition of linear equilibrium adopts the interim approach - i.e., when finding the optimal a given θ , we need to take $E(a)$ as *given*. (This does not matter in fact; we could also take the ex-ante approach and the result would be the same, although this is an artifact of the linearity of the model.) The first-order condition with respect to a is

$$\left(\frac{1-\lambda}{1-\beta\lambda}\right) \left[\frac{1-\lambda}{1-\beta\lambda} a - \lambda \frac{\beta-\beta\lambda}{1-\beta\lambda} E(a) \right] + a - \theta = 0$$

Taking the expectation with respect to a and recalling that $E(\theta) = 0$, it is immediate that $E(a) = 0$. Therefore,

$$a(\theta) = \frac{1}{\left(\frac{1-\lambda}{1-\beta\lambda}\right)^2 + 1} \theta$$

We have now obtained the precise formulas for the unique linear equilibrium, given by (6). ■

Thus, when $\lambda = 1$, $e(a) = a = E(\pi | a)$ - i.e., the private sector's inflation forecast coincides with the rational-expectations benchmark. Despite the private sector's misspecified Phillips Curve, it ends up having correct expectations and therefore the central bank's strategy coincides with the rational-expectations benchmark.

However, when $\lambda < 1$ - i.e., when anticipated inflation has real effects - the inflation forecast is "rigid" in the sense of being a convex combination of the correct conditional expected inflation $E(\pi | a)$ and the ex-ante, unconditional expected inflation $E(\pi)$. As a result, the central bank's equilibrium strategy is less responsive to θ than in the rational-expectations benchmark.

To understand the reason for the private sector's partially responsive expectations, recall that according to the DAG $G : \theta \rightarrow a \rightarrow \pi \leftarrow y$, the private sector erroneously regards y as an *exogenous* variable that affects π , and therefore assigns some weight to the ex-ante expected value of y when forming its inflation forecast. Because y is in fact a consequence of a , the private sector ends up assigning weight to the ex-ante expectation of a , thus failing to fully condition on the actual realization of a .

The *extent* of this failure depends on the relative magnitudes of σ_ε^2 and σ_η^2 . As the Phillips relation becomes less noisy (relative to the relation between the central bank's action and inflation), the erroneous weight on $E(a)$ increases and the deviation from rational expectations is exacerbated. This "expectational rigidity" leads to a policy-attenuation effect that increases

with $\sigma_\varepsilon^2/\sigma_\eta^2$.

The intuition behind the comparative statics is as follows. The private sector tries to account for fluctuations in π (conditional on a) by the variation in y , mistakenly treating the latter as exogenous. The variables π and y are normally distributed. When $\sigma_\eta^2/\sigma_\varepsilon^2$ is low, the definition of conditional expectation of normal variables implies that y gets a large weight in the private sector's inflation forecast, and therefore the model misspecification error generates a larger forecast error - which in turn impels the central bank to adopt a more rigid policy.

4.2 Case 2: $G : \theta \rightarrow a \rightarrow y \rightarrow \pi$

This DAG tells a causal story, according to which the central bank's action has a causal effect on real output, which in turn is the sole direct cause of inflation. As in Case 1, this specification of G inverts the true causal link between output and inflation. However, unlike Case 1, this causal model regards y as an endogenous variable that mediates the effect of monetary policy on inflation.

The private sector's inflation forecast after observing the central bank's action a is

$$E_G(\pi | a) = \int_\pi \pi p_G(\pi | a) = \int_\pi \int_y \pi p(y | a) p(\pi | y) \quad (8)$$

Compare this expression with (5). The error in (8) lies in the term $p(\pi | y)$, which fails to condition π on a . This failure reflects the assumption, embedded in G , that π is independent of a conditional on y . This assumption is inconsistent with the true model G^* .

A linear equilibrium is *strongly linear* if $e(a)$ is proportional to a and $a(\theta)$ is proportional to θ . The next result restricts attention to strongly linear equilibria, in order to avoid the minor technical distraction of proving that every linear equilibrium must be strongly linear (this could be demonstrated transparently in Case 1, but unfortunately not in the present case).

Proposition 3 *There is a unique strongly linear equilibrium, given by*

$$\begin{aligned} e(a) &= \gamma a \\ a(\theta) &= \delta \theta \end{aligned} \tag{9}$$

where γ and δ are uniquely given by the pair of equations:

$$\begin{aligned} \gamma &= \frac{(1 - \gamma\lambda)^2 \delta^2 \sigma_\theta^2 + (1 - \gamma\lambda) \sigma_\varepsilon^2}{(1 - \gamma\lambda)^2 \delta^2 \sigma_\theta^2 + \sigma_\varepsilon^2 + \sigma_\eta^2} \\ \delta &= \frac{1}{(1 - \gamma\lambda)^2 + 1} \end{aligned} \tag{10}$$

Proof. Let us begin by deriving a formula for the private sector's inflation forecast. Denote

$$e(a) = E_G(\pi | a) = \int_y p(y | a) E(\pi | y) \tag{11}$$

Guessing a strongly linear equilibrium, suppose $e(a) = \gamma a$ and $a(\theta) = \delta \theta$. Therefore,

$$y = a - \lambda \gamma a + \varepsilon + \eta$$

Given θ , the central bank effectively chooses a to minimize

$$E(y^2 | a) + (a - \theta)^2 = (1 - \gamma\lambda)^2 a^2 + (a - \theta)^2$$

The first-order condition with respect to a pins down the expression for δ given by (10).

Our remaining task is to derive the expression for γ . To do so, let us first obtain an expression for $E(\pi | y)$. By the true equations for inflation and output,

$$\begin{aligned} E(\pi | y) &= E(a + \varepsilon | a - \lambda e(a) + \varepsilon + \eta = y) \\ &= E(a + \varepsilon | a - \lambda \gamma a + \varepsilon + \eta = y) \end{aligned}$$

which is equal to

$$\frac{1}{1 - \gamma\lambda} E[(1 - \gamma\lambda)a \mid (1 - \gamma\lambda)a + \varepsilon + \eta = y] + E[\varepsilon \mid (1 - \gamma\lambda)a + \varepsilon + \eta = y]$$

The expression $(1 - \gamma\lambda)a + \varepsilon + \eta$ is a sum of independent, normally distributed variables with mean zero. Denote

$$\sigma_a^2 = \delta^2 \sigma_\theta^2$$

Then, using the conditional-expectation rule for such variables,

$$E[a \mid (1 - \gamma\lambda)a + \varepsilon + \eta = y] = \frac{(1 - \gamma\lambda)\delta^2 \sigma_\theta^2}{(1 - \gamma\lambda)^2 \delta^2 \sigma_\theta^2 + \sigma_\varepsilon^2 + \sigma_\eta^2} \cdot y$$

and

$$E[\varepsilon \mid (1 - \gamma\lambda)a + \varepsilon + \eta = y] = \frac{\sigma_\varepsilon^2}{(1 - \gamma\lambda)^2 \delta^2 \sigma_\theta^2 + \sigma_\varepsilon^2 + \sigma_\eta^2} \cdot y$$

Therefore,

$$E(\pi \mid y) = \frac{(1 - \gamma\lambda)\delta^2 \sigma_\theta^2 + \sigma_\varepsilon^2}{(1 - \gamma\lambda)^2 \delta^2 \sigma_\theta^2 + \sigma_\varepsilon^2 + \sigma_\eta^2} \cdot y$$

Now, plug this expression in (11) and obtain

$$e(a) = \frac{(1 - \gamma\lambda)\delta^2 \sigma_\theta^2 + \sigma_\varepsilon^2}{(1 - \gamma\lambda)^2 \delta^2 \sigma_\theta^2 + \sigma_\varepsilon^2 + \sigma_\eta^2} \cdot E(y \mid a)$$

By the inflation and output equations,

$$\begin{aligned} E(y \mid a) &= E(\pi - \lambda e(a) + \eta \mid a) \\ &= E(a - \lambda \gamma a + \varepsilon + \eta \mid a) \\ &= (1 - \gamma\lambda)a \end{aligned}$$

Therefore,

$$e(a) = \frac{(1 - \gamma\lambda)\delta^2 \sigma_\theta^2 + \sigma_\varepsilon^2}{(1 - \gamma\lambda)^2 \delta^2 \sigma_\theta^2 + \sigma_\varepsilon^2 + \sigma_\eta^2} \cdot (1 - \gamma\lambda)a = \gamma a$$

where γ is given by (10).

It remains to establish uniqueness of (γ, δ) . Clearly, δ is a monotone function of γ in the relevant range. Therefore, it suffices to show that γ is unique. Plugging the equation for δ in the equation for γ , we obtain

$$\gamma = \frac{\frac{(1-\gamma\lambda)^2}{[(1-\gamma\lambda)^2+1]^2}\sigma_\theta^2 + (1-\gamma\lambda)\sigma_\varepsilon^2}{\frac{(1-\gamma\lambda)^2}{[(1-\gamma\lambda)^2+1]^2}\sigma_\theta^2 + \sigma_\varepsilon^2 + \sigma_\eta^2}$$

It can be checked that the R.H.S of this equation is continuously and monotonically decreasing in γ in the range $[0, 1]$. Moreover, it is strictly positive when $\gamma = 0$ and strictly below 1 when $\gamma = 1$. Therefore, there is a unique intersection between the curve given by the R.H.S and the line given by the L.H.S. ■

As in Case 1, the variance of the noise terms in the inflation and output equations plays a role in the strongly linear equilibrium in the present case. However, this role is different in a number of ways.

First, the effect of the noisiness of the Phillips Curve (i.e. σ_η^2) on expectational rigidity goes in the opposite direction. Expectations become more rigid as the Phillips relation becomes *less* precise. The intuition is that in Case 2, G posits that the effect of a on π is mediated by y . As the statistical relation between π and y becomes noisier, the perceived mediation effect is attenuated, and this means that the perceived indirect effect of a on π becomes weaker.

Second, unlike Case 1, the private sector's expectations depart from the rational-expectations benchmark even when $\lambda = 1$. That is, even when anticipated inflation fully offsets the effect of inflation on output, the private sector's causal model leads to incorrect, rigid inflation forecasts. In particular, when $\lambda = 1$ and σ_ε^2 (which measures the noisiness of the mapping from a to π) gets large, γ approaches $\frac{1}{2}$ - i.e., the inflation forecast is only "half-responsive" to changes in a .

Finally, unlike Case 1, the variance of the exogenous inflation target θ plays a role in the characterization of the linear equilibrium. As σ_θ^2 becomes larger, the private sector's expectations become less rigid, and as a result the

central bank’s policy approaches the rational-expectations benchmark. The intuition is that larger fluctuations in θ translate to larger fluctuations in a . This in turn implies that the private sector’s interpretation of fluctuations in inflation and output becomes dominated by the fluctuations in a , which is the correct interpretation; the misperception of the role of σ_η^2 and σ_ε^2 becomes less important for the private sector’s expectations.

5 Discussion

The comparison between Cases 1 and 2 shows that the details of the private sector’s causal misinterpretation of the Phillips Curve matter for the qualitative properties of the equilibrium in the interaction between the central bank and the private sector. However, the broad lesson that σ_η^2 and σ_ε^2 matter for the equilibrium holds in both cases because they share the reverse-causality error that misperceives the direction of causality between output and inflation.

Note, however, that not every G that exhibits this reverse causality will give rise to this effect. For example, suppose $R(\boldsymbol{\pi}) = \{\mathbf{y}\}$ and $R(\mathbf{y}) = \emptyset$. Then, G admits no causal path from a or e to π . As in earlier examples that exhibited this feature, the private sector’s equilibrium expectations will be entirely rigid: $e = 0$ with probability one, leading to the most rigid central-bank policy that is possible in this model, without any sensitivity to σ_η^2 and σ_ε^2 . Another example is $R(\boldsymbol{\pi}) = \{\mathbf{a}, \mathbf{y}\}$ and $R(\mathbf{y}) = \{\mathbf{a}\}$. In this case, G induces rational expectations. This can be shown algebraically, or using basic equivalence results from the Bayesian-network literature (see Spiegler (2016,2020b) for details). Thus, reverse causality is a necessary condition for the noise-sensitivity of central-bank policy, but not a sufficient one.

The noise-sensitivity of linear equilibria is of interest in light of the “*caution principle*” attributed to Brainard (1967), which has been suggested as an explanation of monetary-policy attenuation. A high σ_ε^2 captures a situation in which the central bank is highly uncertain about the effect of monetary policy on inflation. However, given our conventional linear-quadratic parameterization of the model, this uncertainty has no effect on the central

bank's policy under rational expectations. However, when the private sector's model reverses the causal link between inflation and output, high σ_ε^2 can lead to an attenuated policy - although, as we saw, the extent of this effect depends on details of the private sector's model. The reason is that σ_ε^2 affects the private sector's interpretation of inflation fluctuations, and therefore its inflation forecast conditional on its information. The lesson is that Phillips-Curve disagreements between the central bank and the private sector can rationalize policy distortions due to uncertainty about the consequences of monetary policy - even in settings that otherwise could not give rise to policy distortions.

The policy distortion that emerged in this paper takes the form of *attenuation* - i.e., the policy is less responsive to exogenous shocks than in the rational-expectations benchmark. However, I do not want to stress this point too strongly. The important effect that our analysis has highlighted is the attenuated response of the private sector's inflation forecast, as a result of its causally misperceived Phillips Curve. The fact that this effect translated to a policy-attenuation effect is an artifact of how the model treated the central bank's private information θ - namely, as an exogenous inflation target. Under alternative specifications of the central bank's preferences, rigid private-sector forecasts would lead to *over-responsive* central-bank policy. One could also imagine a model in which θ has no direct payoff relevance, and instead plays a role as a parameter in the inflation/output equations.² Depending on how the private sector incorporates this dependency into its causal model, such a specification could lead to a more complicated, less transparent relation between the private sector's forecast errors and the central bank's policy distortion. Since paper is a first attempt to incorporate causal disagreements into familiar models of monetary policy, I prioritized pedagogical considerations and refrained from adding these complications. However, I believe they merit further study.

²I thank Stephane Dupraz for making these observations.

References

- [1] Athey, S., A. Atkeson and P. Kehoe (2005), “The Optimal Degree of Discretion in Monetary Policy,” *Econometrica* 73, 1431-1475.
- [2] Barro, R. and D. Gordon (1983), “Rules, Discretion and Reputation in a Model of Monetary Policy,” *Journal of Monetary Economics* 12, 101-121.
- [3] Brainard, W. (1967), “Uncertainty and the Effectiveness of Policy,” *American Economic Review* 57, 411-425.
- [4] Cho, I., N. Williams and T. Sargent (2002), “Escaping Nash Inflation,” *Review of Economic Studies*, 69, 1-40.
- [5] Cowell, R., P. Dawid, S. Lauritzen and D. Spiegelhalter (1999), *Probabilistic Networks and Expert Systems*, Springer, London.
- [6] Dupraz, S., S. Guilloux-Nefussi and A. Penalver (2020), “A Pitfall of Cautiousness in Monetary Policy,” (2020), mimeo.
- [7] Edge, R., M. Kiley and J.P. Laforde (2008), “Natural Rate Measures in an Estimated DSGE Model of the US Economy,” *Journal of Economic Dynamics and Control* 32, 2512-2535.
- [8] Esponda. I. and D. Pouzo (2016), “Berk-Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models,” *Econometrica* 84, 1093-1130.
- [9] Evans, G. and S. Honkapohja (2001), *Learning and Expectations in Macroeconomics*, Princeton University Press.
- [10] Eyster, E. and M. Rabin (2005), “Cursed Equilibrium,” *Econometrica*, 73, 1623-1672.
- [11] Garcia-Schmidt, M. and M. Woodford (2015), “Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis,” NBER Working Paper no. w21614.

- [12] Giacomini, R. (2015), “Economic Theory and Forecasting: Lessons from the Literature,” *Econometrics Journal* 18, C22-C41.
- [13] Hoover, K. (2001), *Causality in macroeconomics*, Cambridge University Press.
- [14] Jehiel, P. (2005), “Analogy-Based Expectation Equilibrium,” *Journal of Economic Theory*, 123, 81-104.
- [15] Koller, D. and N. Friedman (2009), *Probabilistic Graphical Models: Principles and Techniques*. MIT press.
- [16] Kydland, F. and E. Prescott (1977), “Rules rather than Discretion: The Inconsistency of Optimal Plans,” *Journal of Political Economy* 85, 473-491.
- [17] Orphanides, A. and J. Williams, “Robust Monetary Policy with Imperfect Knowledge,” *Journal of monetary Economics* 54, 1406-1435.
- [18] Pearl, J. (2009), *Causality: Models, Reasoning and Inference*, Cambridge University Press, Cambridge.
- [19] Piccione, M. and A. Rubinstein (2003), “Modeling the Economic Interaction of Agents with Diverse Abilities to Recognize Equilibrium Patterns,” *Journal of the European Economic Association* 1, 212-223.
- [20] Sargent, T. (2001), *The Conquest of American inflation*, Princeton University Press.
- [21] Spiegel, R. (2016), “Bayesian Networks and Boundedly Rational Expectations,” *Quarterly Journal of Economics* 131, 1243-1290.
- [22] Spiegel, R. (2020a), “Can Agents with Causal Misperceptions be Systematically Fooled?”, *Journal of the European Economic Association* 18, 583-617.
- [23] Spiegel, R. (2020b), “Behavioral Implications of Causal Misperceptions”, *Annual Review of Economics* 12, 81-106.

- [24] 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, 911-949.
- [25] Woodford, M. (2013), "Macroeconomic Analysis without the Rational Expectations Hypothesis," *Annual Review of Economics* 5, 303-346.