

Forward Guidance without Common Knowledge[†]

By GEORGE-MARIOS ANGELETOS AND CHEN LIAN*

How does the economy respond to news about future policies or future fundamentals? Standard practice assumes that agents have common knowledge of such news and face no uncertainty about how others will respond. Relaxing this assumption attenuates the general equilibrium effects of news and rationalizes a form of myopia at the aggregate level. We establish these insights within a class of games which nests, but is not limited to, the New Keynesian model. Our results help resolve the forward-guidance puzzle, offer a rationale for the front-loading of fiscal stimuli, and illustrate more broadly the fragility of predictions that rest on long series of forward-looking feedback loops. (JEL D82, D83, D84, E12, E23, E52, E62)

How does aggregate demand respond to forward guidance, that is, to news about future monetary policy? How does aggregate supply respond to news about future productivity? How do asset prices respond to news about future earnings?

Standard practice in macroeconomics and finance is to address such questions by assuming rational expectations along with common knowledge of the relevant news and of the structure of the economy. This imposes that the agents have a perfect and common understanding of the effects of the news on current and future economic outcomes. For instance, in the context of forward guidance, it is standard to assume not only that everybody is aware of and attentive to the policy announcement, but also that nobody doubts the awareness, the attentiveness, and the responsiveness of other agents. As a result, the agents have no doubts about how inflation and income, the product of the behavior of others, will adjust.

*Angeletos: Department of Economics, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139 (email: angelet@mit.edu); Lian: Department of Economics, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139 (email: lianchen@mit.edu). This paper was accepted to the *AER* under the guidance of John Leahy, Coeditor. We are grateful to the editor, three anonymous referees for extensive feedback, and to Marco Del Negro, Alisdair McKay, Alexandre Kohlhas, Kristoffer Nimark, Paul Pichler, Franck Portier, and Mirko Wiederholt for conference discussions. For helpful comments, we also thank Adrien Auclert, Harris Dellas, Xavier Gabaix, Alessandro Pavan, Karthik Sastry, and seminar participants at the University of Chicago, Duke University, Northwestern University, Princeton University, the 2016 and 2017 NBER Summer Institutes, the 2016 ESEM in Edinburgh, the 2017 NBER EFG and ME meetings, the 2017 SED meeting, the 2018 AEA meeting, the NYU Stern IUB workshop, the ECB NSM workshop, the 2017 Hydra Workshop in Sardinia, the 20th Central Bank Macroeconomic Modeling Workshop at the Bank of France, the 2017 SNB/CEP Workshop at Study Center Gerzensee, the 2017 Workshop on Sentiments and Beliefs at the Bank of Chile. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

[†]Go to <https://doi.org/10.1257/aer.20161996> to visit the article page for additional materials and author disclosure statement(s).

In this paper, we accommodate such doubts by removing common knowledge of the relevant news and introducing higher-order uncertainty, that is, by allowing the agents to be uncertain about the beliefs and the responses of others. This friction can be the by-product of either dispersed private information, as in Lucas (1972) and Morris and Shin (2002), or rational inattention and costly contemplation, as in Sims (2003) and Tirole (2015). It can thus be interpreted interchangeably as a form of coordination failure that is consistent with equilibrium uniqueness, and as a form of bounded rationality that is consistent with the rational expectations equilibrium concept.

Our contribution is to show that this friction attenuates general equilibrium effects, such as those associated with the Keynesian multiplier and the inflation-spending feedback in the New Keynesian model, and causes the economy to respond to news about the future *as if* the agents were myopic. These insights shed new light on topical policy questions such as the power of forward guidance and the optimal timing of fiscal stimuli. More generally, they highlight the fragility of predictions that rest on long series of forward-looking, general equilibrium feedback loops.

The New Keynesian Model without Common Knowledge.—In Sections II and III, we revisit the New Keynesian model, relax its common knowledge assumptions, and show how its building blocks, the modern versions of the IS and Phillips curves, can be embedded into a class of games that we call dynamic beauty contests. This serves three goals. First, it clarifies the three general equilibrium (GE) mechanisms that operate within the New Keynesian model: the spending-income multiplier running within the demand block, the strategic complementarity in price-setting running within the supply block, and the inflation-spending feedback running across the two blocks. Second, it unearths the role of higher-order beliefs that lie underneath these mechanisms. Third, it motivates the more abstract analysis of Section IV.

GE Attenuation and the Horizon Effect in Dynamic Beauty Contests.—Section IV studies a class of games in which optimal decisions today depend positively on expectations of the future decisions of others. It nests the building blocks of the New Keynesian model, but is not limited to them.¹

Within this framework, we ask the following question: how does the aggregate outcome today (at $t = 0$) respond to the news about the fundamental in a future date (at $t = T \geq 1$)? In our leading application (forward guidance), this question refers to the response of aggregate spending and employment to the news about the policy rate T periods later.

Clearly, the answer to this question depends on how precise and credible the news is, because these factors control how much the forecasts of the policy instrument, or other fundamentals, vary in the first place. But this is not what we are after in this paper. Instead, we study the response of the aggregate outcome *relative* to that of the aforementioned forecasts.

¹Note the crucial difference from the *static* beauty contests studied in Morris and Shin (2002), Woodford (2003a), Angeletos and Pavan (2007), and Bergemann and Morris (2013): in these works, behavior is *not* forward-looking.

Denote this relative response by ϕ_T . Next, let ϕ_T^* be the value that obtains in the frictionless, common knowledge benchmark. This can be decomposed into two components as follows:

$$\phi_T^* = \phi_T^{PE} + \phi_T^{GE},$$

where ϕ_T^{PE} captures the direct or partial-equilibrium (PE) effect, namely the response of the typical agent holding constant the response of others, and ϕ_T^{GE} captures the additional, general equilibrium (GE) effect.² Under quite general assumptions on the information structure, we establish the following three properties about the value of ϕ_T that obtains away from the common knowledge benchmark:

- (i) The lack of common knowledge dulls the GE effect. As a result, ϕ_T is bounded between ϕ_T^{PE} and ϕ_T^* , and is closer to ϕ_T^{PE} when there is more higher-order uncertainty. We refer to this property as “GE attenuation.”
- (ii) This attenuation is stronger the longer the horizon of the news or, equivalently, the longer the series of forward-looking, general equilibrium feedback loops. As a result, the ratio ϕ_T/ϕ_T^* decreases with T . We refer to this property as the “horizon effect.”
- (iii) Under a mild condition, the attenuation grows without bound, and ϕ_T becomes vanishingly small relative to ϕ_T^* , as $T \rightarrow \infty$. This is our “limit result.”

Let us explain the reasoning behind these findings. To begin with, note that the GE effect is driven by beliefs of the future behavior of others. In equilibrium, these beliefs can be expressed as functions of the hierarchy of beliefs of the future fundamental (e.g., the future policy rate). Understanding the GE effect is therefore akin to understanding the hierarchy of beliefs.

In the common knowledge benchmark, the role of higher-order beliefs is concealed due to the property that higher-order beliefs coincide with first-order beliefs. But this property is degenerate: away from that benchmark, higher-order beliefs differ from, and tend to move less than, lower-order beliefs, because they are more anchored to the common prior. By the same token, the expectations of future endogenous outcomes adjust relatively little to the news and the GE effect is attenuated. This proves our first result.

Our second result, the horizon effect, follows from the combination of the aforementioned anchoring of higher-order beliefs with a property that is embedded in the class of forward-looking games studied in this paper. In this class, longer horizons raise the relative importance of higher-order beliefs in equilibrium outcomes, thus also reinforcing the attenuation effect. This is because longer horizons involve more iterations of the forward-looking, Euler-like equations of the model, which in turn

²In our framework, the two components have the same sign, meaning that the GE effect reinforces the PE effect. Furthermore, ϕ_T^{PE} can decrease with T , because of the discounting embedded in individual preferences. Yet, ϕ_T^{GE} can increase with T , and may even explode as $T \rightarrow \infty$, because of powerful general equilibrium feedback loops.

map to beliefs of higher order. In a nutshell, iterating on dynamic GE feedback loops is akin to ascending the hierarchy of beliefs.

Finally, our limit result follows from combining the insight above with the fact that infinite-order beliefs are pegged at the common prior and are therefore unresponsive to the news, even if the level of inattentiveness, or the idiosyncratic noise in the observation and interpretation of the news, is arbitrarily small. This result can be seen as a sharp illustration of the horizon effect. But it also contains the following lesson: predictions that hinge on long series of GE feedback loops are particularly fragile to relaxations of the complete information, rational expectations benchmark.³

We complement these results with another, which recasts the informational friction as a form of myopia: accommodating higher-order uncertainty is akin to having the representative agent discount the future more heavily at the aggregate level than what is rational at the individual level. Furthermore, the as-if discounting is larger the stronger the underlying GE effect.⁴

Policy Lessons.—Suppose that the economy is in a liquidity trap and the zero lower bound (ZLB) on the nominal interest rate binds between today, $t = 0$, and some future date, $t = T - 1$ for some known $T \geq 2$. Because of this constraint, the monetary authority is unable to stimulate the economy by reducing its current policy rate. It can nevertheless try to achieve the same goal by committing to low rates at $t = T$ (or later). The baseline New Keynesian model predicts that the effectiveness of this kind of forward guidance increases with T and explodes as $T \rightarrow \infty$. What is more, the effect is quantitatively huge even for modest T .

These predictions constitute the so-called forward guidance puzzle (Del Negro, Giannoni, and Patterson 2015; McKay, Nakamura, and Steinsson 2016). They are at odds not only with the available evidence (Campbell et al. 2012), but also with the logic that news regarding the distant future should be heavily discounted. This logic is based on PE reasoning. The puzzling predictions are driven by GE effects.

As noted earlier, there are three GE mechanisms at work: the modern version of the Keynesian income-spending multiplier, which runs inside the demand block; the dynamic strategic complementarity in the price-setting decisions of the firms, which runs inside the supply block; and the inflation-spending feedback, which runs across the two blocks. These mechanisms act as multipliers of the PE effects of monetary policy. Removing common knowledge arrests these multipliers, thus also bringing the predictions of the model closer to the aforementioned PE logic.

To illustrate, suppose that each agent worries that any other agent is unaware of, or inattentive to, the policy news with a 25 percent probability. Under a textbook parameterization (Galí 2008), this kind of doubt reduces the power of forward guidance by about 90 percent at the five-year horizon. Importantly, this attenuation is *relative* to the movements in *expectations* of interest rates. Our theory therefore

³In the New Keynesian model, examples of such predictions include not only those regarding the power of forward guidance, but also those regarding the indeterminacy of interest-rate pegs and the so-called neo-Fisherian effects.

⁴Apart from offering a sharp representation of how the absence of common knowledge modifies the behavior of the economy, this result also rationalizes the kind of myopia assumed in Gabaix (2016).

helps explain why forward guidance may have a modest effect on expectations of inflation and income, and on actual activity, *relative* to its effect on the yield curve.

Turning to fiscal policy, the standard New Keynesian model predicts that, in the presence of a binding ZLB constraint, a fiscal stimulus of a given size is more effective when it is back-loaded, i.e., when it is announced now but implemented later on. This prediction hinges on the same GE feedback loops as those that govern the power of forward guidance. A variant of the aforementioned results therefore offers a rationale for the front-loading of fiscal stimuli: such front-loading improves coordination in the sense of reducing the bite of higher-order uncertainty.

Remark.—Our work invites the analyst to study the role of higher-order beliefs. It does not, however, require that the agents themselves engage in higher-order reasoning. Instead, as in any rational expectations context, it suffices that their expectations of the relevant economic outcomes, such as inflation and income, are consistent with actual behavior. From this perspective, our analysis is a revision of the predictions that the analyst can make, once she liberates the rational expectations hypothesis from the auxiliary assumption that the agents face no doubts about the information, or the attentiveness, of others.

Layout.—Section I discusses the related literature. Section II introduces our version of the New Keynesian model. Section III nests the demand and supply blocks of that model within the class of dynamic beauty contests that we consider. Section IV, which is self-contained, studies this class of games and develops the key theoretical results. Sections V and VI work out the implications for, respectively, monetary and fiscal policy. Section VII concludes the main text. The online Appendix contains the proofs and additional results.

I. Related Literature

On the theoretical side, our paper builds heavily on the macroeconomic literature on incomplete information and beauty contests.⁵ Some of this literature, including Morris and Shin (2002) and Woodford (2003a), focuses on *static* beauty contests, namely settings in which the agents only have to forecast the concurrent actions of others. By contrast, Allen, Morris, and Shin (2006), Bacchetta and van Wincoop (2006), Morris and Shin (2006), and Nimark (2008, 2017) study *dynamic* beauty contests, namely settings in which the agents must forecast the future actions of others. Our paper shares with the latter set of works the emphasis on forward-looking expectations. But whereas these works focus on the effects of learning, our paper focuses on the horizon of the events that the agents have to forecast.

On the applied side, our paper adds to the recent literature on the forward-guidance puzzle. Del Negro, Giannoni, and Patterson (2015) and McKay, Nakamura, and Steinsson (2016) seek to resolve the puzzle by introducing short horizons and liquidity constraints, and Andrade et al. (2015) do so by letting forward guidance confound good news about the future policy instrument with bad news about other

⁵See Angeletos and Lian (2016c) for a survey.

fundamentals. We make an orthogonal point: we show how higher-order uncertainty anchors income and inflation expectations, thus attenuating the GE effects of forward guidance.

The insight that lack of common knowledge attenuates GE effects is quite general; see Angeletos and Lian (2016a) for an abstract formulation. The present paper offers a topical application within the context of the New Keynesian model along with another general insight: the horizon effect and the associated form of myopia.

Related forms of GE attenuation and myopia can be found in Farhi and Werning (2017) and Gabaix (2016). Farhi and Werning (2017) obtain such attenuation by replacing rational expectations equilibrium (REE) with Level- k thinking, an approach that follows the lead of García-Schmidt and Woodford (2015). Gabaix (2016) achieves the same goal by assuming that beliefs are biased in the direction of underestimating the persistence of the underlying shocks and of their effects on the economy. We refer the interested reader to Angeletos and Lian (2016a) and Angeletos and Huo (2018) for detailed expositions of how these approaches connect to and differ from ours.⁶

Closely related is the earlier work of Wiederholt (2015). That paper notes that higher-order uncertainty can dull the inflation-spending feedback of the New Keynesian model, thus also reducing the power of forward guidance. However, that paper does not consider how higher-order uncertainty can attenuate the GE effects that run inside each block of the model. Most important, it does not consider how the impact of higher-order uncertainty on equilibrium outcomes and the resulting weakening of forward guidance vary with the horizon of the events that the agents need to forecast, which is the core contribution of our paper.

Finally, Chung, Herbst, and Kiley (2015) and Kiley (2016) argue that some of the paradoxical predictions of the New Keynesian model are resolved once the nominal rigidity is attributed to sticky information as in Mankiw and Reis (2002). But this is largely because these works abstract entirely from the price stickiness seen in the microdata and from the forward-looking aspect in the firms' price-setting decisions. These works also rule out information frictions among the consumers. They therefore do not share our insights regarding the anchoring of forward-looking expectations and the attenuation of the associated GE effects.

II. Framework

In this section, we introduce the framework used for the applied purposes of the paper. This is the same as the textbook New Keynesian model (Woodford 2003b; Galí 2008), except that we allow the agents to face uncertainty about the information, the beliefs, and the behavior of others.

⁶Let us briefly note the key differences. Because our approach is compatible with the REE concept, it bypasses a challenge faced by the aforementioned works, namely what the agents do when they experience again and again a reality that is inconsistent with their beliefs. Our approach is also easier to reconcile with the evidence in Coibion and Gorodnichenko (2012, 2015) and Vellekoop and Wiederholt (2017), which indicate that expectations adjust over time in a manner broadly consistent with Bayesian learning rather than being dogmatic in the manner assumed in the aforementioned works. Finally, while the form of GE attenuation captured by our approach is robust to settings in which GE effects offset PE effects (think of settings featuring strategic substitutability), this is not necessarily the case for Level- k thinking.

Consumers.—There is a measure-1 continuum of ex ante identical consumers in the economy, indexed by $i \in \mathcal{I}_c = [0, 1]$. Preferences are given by

$$(1) \quad \mathcal{U}_{i,0} = \sum_{t=0}^{+\infty} \beta^t U(c_{i,t}, n_{i,t}),$$

where $c_{i,t}$ and $n_{i,t}$ denote the consumer’s consumption and labor supply at period t , $\beta \in (0, 1)$ is the discount factor, and U is the per-period utility function. The latter is specified as

$$U(c, n) = \frac{1}{1 - 1/\sigma} c^{1-1/\sigma} - \frac{1}{1 + \epsilon} n^{1+\epsilon},$$

where $\sigma > 0$ is the elasticity of intertemporal substitution and $\epsilon > 0$ is the inverse of the Frisch elasticity. The budget constraint in period t is given, in real terms, by the following:

$$(2) \quad c_{i,t} + s_{i,t} = \frac{R_{t-1}}{\pi_t} s_{i,t-1} + w_{i,t} n_{i,t} + e_{i,t},$$

where $s_{i,t}$ is the consumer’s saving in period t , R_{t-1} is the nominal gross interest rate between $t - 1$ and t , $\pi_t \equiv p_t/p_{t-1}$ is the gross inflation rate, p_t is the aggregate price level at t , and $w_{i,t}$ and $e_{i,t}$ are the real wage and the real dividend received by the consumer. The wage and the dividend are allowed to be consumer-specific for reasons that will be explained shortly. We finally denote aggregate consumption by $c_t = \int_{\mathcal{I}_c} c_{i,t} di$, aggregate labor supply by $n_t = \int_{\mathcal{I}_c} n_{i,t} di$, and so on.

Firms.—There is a measure-1 continuum of ex ante identical firms, indexed by $j \in \mathcal{I}_f = (1, 2]$. Each of these firms is a monopolist that produces a differentiated intermediate-good variety. The output of firm j is denoted by y_t^j , its nominal price is denoted by p_t^j , and its real profit is denoted by e_t^j . The technology is assumed to be linear in labor and productivity is fixed to 1, so that

$$(3) \quad y_t^j = l_t^j,$$

where l_t^j is the labor input. These intermediate goods are used by a competitive sector as inputs in the production of the final good. The technology is constant elasticity of substitution (CES) with elasticity $\varsigma > 1$. Aggregate output is thus given by

$$(4) \quad y_t = \left(\int_{\mathcal{I}_f} (y_t^j)^{\frac{\varsigma-1}{\varsigma}} dj \right)^{\frac{\varsigma}{\varsigma-1}}.$$

The corresponding price index, that is, the nominal price level, is denoted by p_t .

Sticky Prices.—Nominal rigidity takes the familiar, Calvo-like, form: in each period, a randomly selected fraction $\theta \in (0, 1]$ of the firms must keep prices unchanged, while the rest can reset them.

Monetary Policy.—For the time being, we do not need to specify how monetary policy is conducted. To fix ideas, however, it is useful to think of a situation where the agents know what the policy will be in the short run (say, because the ZLB is binding) but face uncertainty about the policy in the more distant future. In this context, the question of interest is how outcomes today vary with the expectations of the future policy.

Rational Expectations and Common Knowledge.—Throughout, we impose rational expectations equilibrium (REE). We nevertheless depart from standard practice by removing common knowledge of the state of nature, accommodating higher-order uncertainty, and preventing the agents from reaching a consensus on the future path of the economy.⁷

In abstract games, such as those studied in Morris and Shin (2002, 2003) or in Section IV, lack of common knowledge can be *directly* imposed by endowing each agent with exogenous private information (or different Harsanyi types). In macro and finance applications, however, there is the complication that market signals, such as prices, aggregate information. To preserve the absence of common knowledge, we either have to make sure there is “noise” in the available market signals (Grossman and Stiglitz 1980), or assume that the observation of all the relevant variables, including the available market signals, is contaminated by idiosyncratic noise due to rational inattention (Sims 2003).

We follow the first approach in our New Keynesian application in order to make clear that our policy lessons are robust to allowing each consumer to observe her own income, each firm to observe her own demand and supply conditions, and everybody to have common knowledge of the *current* monetary policy and the *current* price level. We nevertheless like the second approach, too, because it bypasses the need for the auxiliary shocks described below and because it allows the reinterpretation of the assumed friction as “costly contemplation” (Tirole 2015). Section IV therefore allows for a flexible interpretation of the friction under consideration and online Appendix B sketches in more detail how our insights can be recast under the lenses of rational inattention.

The Auxiliary Shocks.—Take any period t and let w_t , e_t , and μ_t denote, respectively, the average real wage, the average firm profit, and the average markup in the economy. The real wage and the dividend received by consumer i at t are given by, respectively, $w_{i,t} = w_t \xi_{i,t}$ and $e_{i,t} = e_t \zeta_{i,t}$, where $\xi_{i,t}$ and $\zeta_{i,t}$ are i.i.d. across i and t , independent of one another, and independent of any other random variable in the economy. On the other hand, the real wage paid by firm j is $w_t^j = w_t u_t^j$, and the markup charged by it is $\mu_t^j = \mu_t \nu_t^j$, where u_t^j and ν_t^j are i.i.d. across both j and t , independent of one another, and independent of any other random variable in the economy. One can interpret u_t^j and ν_t^j as idiosyncratic shocks to a firm’s marginal cost and her optimal markup, and $\xi_{i,t}$ and $\zeta_{i,t}$ as idiosyncratic shocks to a consumer’s

⁷By definition, the state of nature contains the entire profile of the information sets in the population. Furthermore, in any given equilibrium, information sets pin down beliefs of all the endogenous outcomes. It follows that the state of nature can encode arbitrary news about the future, and that the agents can lack a common belief about the future outcomes only if they lack common knowledge of the state of nature.

labor and financial income.⁸ Finally, the aggregate markup shock, μ_t , is also i.i.d. over time and independent of any other variable.

As anticipated, the *sole* modeling role of all these shocks is to “noise up” the information that each agent can extract from the available market signals. We can thus accommodate the desired friction while allowing, in every t , each consumer i to have private knowledge of $(w_{i,t}, e_{i,t})$, each monopolist j to have private knowledge of (w_t^j, μ_t^j) , and everybody to have common knowledge of the current interest rate, R_t , of the current prices, $(p_t^j)_{j \in [0, 1]}$, and therefore also of the price level.

III. The New Keynesian Model as a Pair of Beauty Contests

In this section, we develop our beauty contest representation of the New Keynesian model. Apart from motivating the class of games studied in Section IV, this representation reveals how the GE effects of that model are related to higher-order beliefs.

To keep the analysis tractable, we work with the log-linearization of the model around a steady state in which inflation is 0 and the nominal interest rate equals the natural rate (i.e., $\beta R = 1$). With abuse of notation, we henceforth reinterpret all the variables as the log-deviations from their steady-state counterparts and concentrate on the joint determination of π_t and c_t (or, equivalently, of π_t and y_t).

PROPOSITION 1 (Beauty Contests): *In any equilibrium, and regardless of the level of higher-order uncertainty, the following properties are true. First, aggregate spending satisfies*

$$(5) \quad c_t = -\sigma \left\{ \sum_{k=1}^{+\infty} \beta^{k-1} \bar{E}_t^c [r_{t+k}] \right\} + \frac{1-\beta}{\beta} \left\{ \sum_{k=1}^{+\infty} \beta^k \bar{E}_t^c [c_{t+k}] \right\},$$

where $\bar{E}_t^c[\cdot]$ denotes the average expectation of the consumers and $r_t \equiv R_{t-1} - \pi_t$ denotes the real interest rate between $t - 1$ and t . Second, inflation satisfies

$$(6) \quad \pi_t = \varkappa \left\{ mc_t + \sum_{k=1}^{+\infty} (\beta\theta)^k \bar{E}_t^f [mc_{t+k}] \right\} + \frac{1-\theta}{\theta} \left\{ \sum_{k=1}^{+\infty} (\beta\theta)^k \bar{E}_t^f [\pi_{t+k}] \right\} + \varkappa \mu_t,$$

where $\bar{E}_t^f[\cdot]$ denotes the average expectation of the firms, mc_t denotes the real marginal cost in period t , and $\varkappa \equiv (1 - \theta)(1 - \beta\theta)/\theta$. Finally, $mc_t = (\epsilon + 1/\sigma)y_t$ and $y_t = c_t$.

Condition (5) is a GE variant of the textbook version of the permanent income hypothesis. It follows from log-linearizing the individual consumption functions, aggregating them, and using market clearing in the product, labor, and debt markets. It gives aggregate consumption in any given period as an increasing function

⁸ Although we treat these shocks as exogenous, it is not hard to fill in the missing microfoundations. For instance, firm-specific markup shocks can be microfounded by allowing for good-specific shocks to the elasticity of the demand faced by each monopolist. Similarly, consumer- and firm-specific wage shocks can be justified by introducing idiosyncratic taste and productivity shocks, letting labor markets be segmented, and allocating the agents in such a way that there is market-specific random variation in the equilibrium wage, translating to idiosyncratic variation in the wage received by a consumer or the wage paid by a firm.

of the consumers' expectations of future aggregate income, which in equilibrium coincides with aggregate consumption, and as a decreasing function of their expectations of future real interest rates. Condition (6), on the other hand, follows from the optimal price-resetting behavior of the firms. It gives the rate of inflation as a function of the current real marginal cost and of the firms' expectations of the future marginal costs and the future inflation.

In the rest of this section, we first show that conditions (5) and (6) reduce to, respectively, the representative consumer's Euler condition and the New Keynesian Phillips curve (NKPC) once higher-order uncertainty is assumed away. This clarifies the sense in which our result "opens up" these familiar equations away from the common knowledge benchmark. We then expand on the three GE mechanisms that are embedded in the model and on their relation to higher-order beliefs.

Relation to the Euler Condition and the NKPC.—To begin with, rearrange condition (5) as follows:

$$(7) \quad c_t = -\sigma \bar{E}_t^c[r_{t+1}] + (1 - \beta) \bar{E}_t^c[c_{t+1}] \\ + \beta \left\{ -\sigma \sum_{k=1}^{+\infty} \beta^{k-1} \bar{E}_t^c[r_{t+1+k}] + \frac{1 - \beta}{\beta} \sum_{k=1}^{+\infty} \beta^k \bar{E}_t^c[c_{t+1+k}] \right\}.$$

Next, compute c_{t+1} by applying condition (5) in period $t + 1$, and take the period- t average expectation of it:

$$(8) \quad \bar{E}_t^c[c_{t+1}] = -\sigma \sum_{k=1}^{+\infty} \beta^{k-1} \bar{E}_t^c[\bar{E}_{t+1}^c[r_{t+1+k}]] + \frac{1 - \beta}{\beta} \sum_{k=1}^{+\infty} \beta^k \bar{E}_t^c[\bar{E}_{t+1}^c[c_{t+1+k}]].$$

Clearly, the term inside the big brackets in condition (7) is the same as the one in the right-hand side of condition (8), except for the difference in the expectation operators. Can we use this observation to recover the Euler equation of a representative consumer?

In the common knowledge benchmark, the answer is yes. Because, in any period, all agents are assumed to share the same information and the same belief, we have $\bar{E}_t^c[\cdot] = E_t[\cdot]$, where $E_t[\cdot]$ denotes the rational expectation conditional on the common information set, i.e., the information of the representative agent. Using this fact and applying the law of iterated expectations for the representative agent, we infer that, for every $k \geq 1$,

$$\bar{E}_t^c[\bar{E}_{t+1}^c[r_{t+1+k}]] = E_t[E_{t+1}[r_{t+1+k}]] = E_t[r_{t+1+k}] = \bar{E}_t^c[r_{t+1+k}],$$

and similarly $\bar{E}_t^c[\bar{E}_{t+1}^c[c_{t+1+k}]] = \bar{E}_t^c[c_{t+1+k}]$. That is, second-order beliefs coincide with first-order beliefs. It follows that $\bar{E}_t^c[c_{t+1}]$ coincides with the term inside the brackets in condition (7). Therefore, this condition reduces to the familiar, representative-agent Euler condition:

$$(9) \quad c_t = -\sigma E_t[r_{t+1}] + E_t[c_{t+1}].$$

A similar argument establishes that, in the common knowledge benchmark, condition (6) reduces to

$$(10) \quad \pi_t = \varkappa mc_t + \beta E_t[\pi_{t+1}] + \varkappa \mu_t,$$

which is the familiar NKPC.⁹

Away from the common knowledge benchmark, the argument above breaks because the law of iterated expectations no longer holds for the average beliefs. In particular, as it will become clear in the next section, the second-order beliefs of either the exogenous fundamentals or the endogenous outcomes do not coincide with the corresponding first-order beliefs. In the present context, this means that

$$\bar{E}_t^c [\bar{E}_{t+1}^c [r_{t+1+k}]] \neq \bar{E}_t^c [r_{t+1+k}] \quad \text{and} \quad \bar{E}_t^c [\bar{E}_{t+1}^c [c_{t+1+k}]] \neq \bar{E}_t^c [c_{t+1+k}].$$

It follows that the term in the bracket in condition (7) does not coincide with $\bar{E}_t^c [c_{t+1}]$ and therefore condition (5) does not reduce to the familiar Euler condition. The same logic explains why condition (6) cannot be reduced to the familiar NKPC.

To recap, the key lesson so far is that the familiar representation of the New Keynesian model hinges on the assumption that the agents share the same beliefs, not only about the exogenous impulses, such as news about current or future monetary policy, but also about the future path of aggregate spending and inflation. When this assumption is imposed, one can understand the equilibrium of the economy as the solution to a single-agent decision problem: think of a fictitious agent, who controls the vector (c_t, π_t) and whose optimal behavior is described by the dynamic system (9)–(10). When, instead, this assumption is relaxed, the equilibrium of the economy is better understood as the solution to a complex, multilayer game among a large number of differentially informed players.¹⁰

The Economy as a Pair of Dynamic Beauty Contests.—Consider first the behavior of the consumers and, momentarily, treat the process for the real interest rate as exogenous. Condition (5) then defines a game in which the players are the consumers, the actions are the consumers’ spending levels at different periods, and the payoff-relevant fundamental is the path of real interest rates. This game is a dynamic beauty contest in the sense that each consumer has an incentive to spend more now when she expects the other consumers to spend more in the future. This is because higher aggregate consumption translates to higher income, which in turn justifies more consumption today. Condition (5) therefore isolates the GE mechanism that runs inside the demand block of the model that is known as the Keynesian multiplier.

Consider next the behavior of the firms and, momentarily, treat the process of mc_t as exogenous. Condition (6) then defines a beauty contest game in which the players are the firms, the actions are the prices, and the relevant fundamental is the average real marginal cost. A dynamic strategic complementary is present in this

⁹Using the facts that $mc_t = (\epsilon + 1/\sigma)y_t$ and $c_t = y_t$, the NKPC can be restated as $\pi_t = \kappa y_t + \beta E_t[\pi_{t+1}] + u_t$, where $\kappa \equiv (\epsilon + 1/\sigma)\varkappa$ is the slope of the NKPC with respect to output and $u_t \equiv \varkappa \mu_t$ is the so-called cost-push shock.

¹⁰These points explain the sense in which higher-order uncertainty can be interpreted as frictional coordination.

game because the optimal reset price of a firm depends on her expectation of the future path of the nominal price level, which in turn depends on the decisions of other firms. Condition (6) therefore encapsulates the GE mechanism that runs inside the supply block.

Finally, consider the interaction of the demand and the supply blocks, that is, the interaction of conditions (5) and (6). This interaction can be understood as a metagame between the consumers and the firms. This metagame features dynamic strategic complementarity in the following sense: the consumers find it optimal to spend more now when they expect the firms to raise their prices in the future, because high inflation means low real returns to saving; and, symmetrically, the firms find it optimal to raise their prices now when they expect the consumers to spend more in the future, because high aggregate output means high marginal costs. This GE mechanism is known as the inflation-spending feedback.

PE versus GE.—Let us momentarily treat the process for the real rate as exogenous and concentrate on the consumption beauty contest defined by condition (5).

Next, fix a $t \geq 0$ and a $k \geq 2$, and ask the following question: how does c_t vary with the period- t beliefs of r_{t+k} , holding constant the beliefs of $r_{t+k'}$ for all $k' \neq k$? This formalizes the question of how much aggregate spending moves in response to news about future real rates.

From condition (5), we readily see that, if we also hold constant $\bar{E}_t^c[c_{t+k}]$ for all $k' \geq 1$, a one-unit reduction in $\bar{E}_t^c[r_{t+k}]$ raises c_t by $\beta^{k-1}\sigma$. This can be interpreted as a PE effect because it measures the response of the typical consumer to the interest rate faced by himself, holding constant his beliefs about the future behavior of all other consumers, thus also holding constant his expectations of future income.¹¹ By the same token, the GE effect is tied to how much the typical consumer expects the other consumers to adjust their spending in the future.

From GE to Higher-Order Beliefs.—Understanding the adjustment in the expectations of aggregate spending and the associated GE effect is formally equivalent to understanding the hierarchy of beliefs.

For instance, consider $\bar{E}_t^c[c_{t+1}]$, the average expectation of the next-period aggregate spending. From condition (5), c_{t+1} can be expressed as a function of the period- $(t+1)$ *first-order* beliefs of the future real rates and the future aggregate spending. It follows that $\bar{E}_t^c[c_{t+1}]$ is a function of the period- t *second-order* beliefs of the future rates and the future aggregate spending, a fact already shown in condition (8). Using the same logic again and again, we can express the second-order beliefs of future spending, the last term in condition (8), as functions of third- and higher-order beliefs of the future real rates. We conclude that understanding how the rational expectations of future aggregate spending and income adjust to news

¹¹ Strictly speaking, this is the PE effect plus the “within-period” GE effect, namely, the feedback effect between aggregate spending and the *contemporaneous* level of income. In the version of the model we have considered, this within-period GE effect is not attenuated because each consumer knows perfectly her current income and each firm knows perfectly her current marginal cost. Had we relaxed these assumptions, we could have attenuated this effect as well. In any event, this effect is of little interest since it vanishes as the length of the time interval shrinks to 0, which justifies the interpretation of $\beta^{k-1}\sigma$ as the PE effect.

about future real rates is formally the same as understanding how the first- and higher-order beliefs of the future rates adjust to the news.

A similar point applies if we focus on the beauty contest defined by condition (6), or if we consider the metagame between the demand and the supply blocks of the economy: understanding the GE adjustment to any news is equivalent to understanding the associated movement in the relevant higher-order beliefs. This explains where we are heading next: by recasting the GE effects of news about the future in terms of higher-order beliefs, we explain why these effects lose steam once one departs from the familiar but unrealistic common knowledge benchmark.

IV. Dynamic Beauty Contests

This section contains the broader theoretical contribution of our paper. In this section, we sidestep the microfoundations and establish our main insights within a more abstract class of dynamic beauty contests. This class nests not only the two blocks of the New Keynesian model, but also other applications, such as incomplete-information asset-pricing models in the tradition of Singleton (1987). Moreover, we remain agnostic about the origins of higher-order uncertainty. This helps clarify that our lessons do not hinge on whether such uncertainty is the product of dispersed private information or of cognitive constraints along the line of Sims (2003) and Tirole (2015).¹²

A. Setup

Time is discrete, indexed by $t \in \{0, 1, \dots\}$, and there is a continuum of players, indexed by $i \in [0, 1]$. In each period t , each agent i chooses an action $a_{i,t} \in \mathbb{R}$. We denote the corresponding average action by a_t and let $\Theta_t \in \mathbb{R}$ denote an exogenous fundamental that becomes commonly known in period t .

We specify the best response of player i in period t as follows:

$$(11) \quad a_{i,t} = \Theta_t + \gamma E_{i,t}[a_{i,t+1}] + \alpha E_{i,t}[a_{t+1}],$$

where $\alpha, \gamma > 0$ are fixed parameters. Note that a player’s optimal action in any given period depends on her expectation of both her own and the aggregate action in the next period. The former effect is parameterized by γ , the latter by α .

By iterating condition (11) forward, aggregating across i , and letting $\bar{E}_t[\cdot]$ denote the average expectation in period t , we reach the following representation of the beauty contest under consideration:

$$(12) \quad a_t = \Theta_t + \gamma \left\{ \sum_{k=1}^{+\infty} \gamma^{k-1} \bar{E}_t[\Theta_{t+k}] \right\} + \alpha \left\{ \sum_{k=1}^{+\infty} \gamma^{k-1} \bar{E}_t[a_{t+k}] \right\} \quad \forall t \geq 0.$$

¹²By sidestepping the microfoundations, we not only speak to a larger class of environments, but also abstract from the kind of auxiliary shocks that were necessary before in order to limit the endogenous aggregation of information through markets. This rules out the possibility that agents confuse one kind of fundamental for another (say, the news about future monetary policy with news about the idiosyncratic components of wage and markups), isolates the role of higher-order uncertainty about the fundamental of interest, and clarifies why our contribution is orthogonal to that of Lucas (1972). It also rules out the amplification mechanism studied in Chahrour and Gaballo (2017), which rests on such confusion.

This is the key equation we work with in this section. It relates the aggregate outcome in any given period to the concurrent fundamental, the average forecasts of the future fundamentals, and the average forecasts of the future aggregate outcomes. The first two terms capture the direct or PE effect of the current and future fundamentals; the GE effect is captured by the last term, which regards the actions of others. Under this interpretation, α parameterizes the importance of the GE effect.

When all agents share the same information, so that $\bar{E}_t[\cdot] = E_t[\cdot]$ is the rational expectation of a representative agent, we can use the law of iterated expectations to restate condition (12) as follows:

$$(13) \quad a_t = \Theta_t + \delta E_t[a_{t+1}] \quad \forall t \geq 0,$$

where $\delta \equiv \gamma + \alpha$. The above can be thought of as an aggregate-level Euler condition and δ as the effective discount factor that governs how much current outcomes depend on expectations of future outcomes. This dependence reflects the combination of PE and GE effects, as indicated by the fact that δ is given by the sum of γ and α . Yet, by looking at equation (13) alone, it is impossible to tell these effects apart. This explains why it is important to “open up” this equation to (12), just as we did with the representative agent’s Euler condition and the NKPC in the previous section: the distinction between PE and GE effects becomes crucial once the common knowledge assumption is dropped.

Interpretation.—Condition (12) directly nests the supply block of our New Keynesian model: just interpret Θ_t as the real marginal cost scaled by \varkappa , interpret a_t as inflation, and set $\gamma = \beta\theta$ and $\alpha = \beta(1 - \theta)$. Similarly, to nest the demand block, interpret Θ_t as the real interest rate scaled by $-\sigma$, interpret a_t as the aggregate level of spending, and set $\gamma = \beta$ and $\alpha = 1 - \beta$.¹³ Moving to a different application, consider the class of incomplete-information asset-pricing models studied in Singleton (1987), Allen, Morris, and Shin (2006), Bacchetta and van Wincoop (2006), and Nimark (2017). These models feature a single asset (“stock market”) and overlapping generations of differentially informed traders. The equilibrium asset price is shown to satisfy the following condition:

$$p_t = d_t - s_t + \beta \bar{E}_t[p_{t+1}],$$

where p_t , d_t , and s_t denote the period- t price, dividend, and supply, respectively, and $\bar{E}_t[\cdot]$ is the average expectation of the period- t traders. The above can be nested in condition (12) by letting $\Theta_t = d_t - s_t$, $a_t = p_t$, $\alpha = \beta$, and $\gamma = 0$.¹⁴ This

¹³A minor qualification is needed here. In our version of the New Keynesian model, the consumers do not know the *real* interest rate between today and tomorrow because they have to forecast tomorrow’s inflation. It follows that we cannot simply let $\Theta_t = -\sigma r_{t+1}$. That said, we can nest the demand block in condition (12) *as it is* if we assume that prices are completely rigid ($\theta = 1$), fix inflation at 0, and let $\Theta_t = -\sigma R_t$. For the more general case in which inflation is variable, we can either modify the model so that r_{t+1} is observed, or nest the existing version of the demand block in the variant of (12) that replaces Θ_t with $\bar{E}_t[\Theta_t]$. The results developed in the sequel are robust to such modifications.

¹⁴The restriction $\gamma = 0$ means that the aggregate outcome today depends only on the average expectation of the aggregate outcome tomorrow, as opposed to the entire path of the aggregate outcome in the future. Relative to the more general case we study here, this restriction decreases dramatically the dimensionality of the higher-order

indicates, not only the broader applicability of the insights we develop in the rest of this section, but also that the aggregate outcome a_t can be a market-clearing price in a Walrasian setting rather than the average action of a set of players.

The Question of Interest.—How does aggregate spending respond to news about future real interest rates? How does inflation respond to news about future real marginal costs? How do asset prices respond to news about future dividends? In the rest of this section, we seek to answer these kinds of questions without taking a specific stand on how precise or credible the available news is, or how exactly it has to be modeled. We achieve this by focusing on how the aggregate outcome covaries with the concurrent forecasts of future fundamentals, and by abstracting from the exact source and magnitude of the variation in these forecasts.

More specifically, we fix an arbitrary $T \geq 2$ and isolate the role of news about the period- T fundamental, Θ_T , by shutting down the uncertainty about the fundamentals in any other period. Without any loss, we then let $\Theta_t = 0$ for all $t \neq T$ and treat Θ_T as the *only* random fundamental.¹⁵ We also anchor the beliefs of the aggregate outcome “at infinity” by imposing that $\lim_{k \rightarrow \infty} \gamma^k E_{i,t}[a_{t+k}] = 0$ with probability 1. This is akin to ruling out infinite-horizon bubbles and can be justified either by letting the game end at *any* finite period $T' > T$, or by imposing that $\gamma < 1$ and that a_t is bounded. We finally state the question of interest as follows: how does a_0 , the current outcome, covary with $\bar{E}_0[\Theta_T]$, the current average forecast of the future fundamental?

B. Complete versus Incomplete Information

In what follows, we study how the answer to the aforementioned question varies as we move from the frictionless benchmark typically studied in applied work to the more realistic scenario in which the agents are uncertain about one another’s beliefs and actions. To fix language, and to rule out trivial cases in which the informational friction relates only to variables that are immaterial for our purposes, we introduce the following definitions.

DEFINITION 1: *We say that information is complete, or that the agents have common knowledge of the available news, if $E_{i,t}[\Theta_T] = E_{j,t}[\Theta_T]$ with probability 1 for all (i, j, t) such that $i \neq j$ and $t \leq T - 1$. (And when the converse is true, we say that information is incomplete.)*

DEFINITION 2: *We say that the agents are able to reach a consensus about the future trajectory of the economy if $E_{i,t}[a_\tau] = E_{j,t}[a_\tau]$ with probability 1 for all (i, j, t, τ) such that $i \neq j$ and $t < \tau \leq T$.*

beliefs that the aforementioned works had to deal with. The most essential difference, however, between these works and ours is that they focus on the effects of learning whereas we focus on the effects of different horizons.

¹⁵Setting $\Theta_t = 0$ for all $t \neq T$ is equivalent to isolating the variation in hierarchy of beliefs about the period- T fundamental that is orthogonal to variations in belief hierarchy about the fundamentals in other periods. Our result can thus be read as an “orthogonalization” of the effects of different horizons. Without any loss, we also normalize the unconditional mean of Θ_T to be 0.

Note that the first notion regards the beliefs the agents hold about the *exogenous* fundamental, whereas the second notion regards the beliefs the agents form about the *endogenous* outcomes. While conceptually distinct, the two notions are tied together under the rational expectations hypothesis.

PROPOSITION 2: *Along any rational expectations equilibrium, lack of consensus about the future outcomes is possible only when information is incomplete.*

This clarifies the modeling role, and our preferred interpretation, of the introduced friction. What is at stake here is not how much the agents know about the fundamentals, but rather the ability to coordinate their beliefs and their responses to the exogenous impulses (the news). Accordingly, the comparisons we develop in the sequel between complete- and incomplete-information economies hold true even if the agents in the incomplete-information economy are *better* informed than their counterparts in the complete-information economy, in the sense of facing less first-order uncertainty about Θ_T .¹⁶

The Frictionless Benchmark.—As a reference point for our subsequent analysis, we first answer the question of interest in the absence of higher-order uncertainty, that is, under complete information.

LEMMA 1: *Suppose that information is complete. For all states of nature,*

$$(14) \quad a_0 = \phi_T^* \cdot E_0[\Theta_T],$$

where $\phi_T^* = \delta^T \equiv (\gamma + \alpha)^T$.

The scalar ϕ_T^* measures how much the aggregate outcome covaries with the average forecast of the fundamental T periods later. This scalar is easily computed by iterating on condition (13), which is the aggregate-level Euler condition of the complete-information economy. Indeed, since $\Theta_t = 0$ for all $t \neq T$, we have $a_T = \Theta_T$ and $a_t = \phi_{T-t}^* E_t[\Theta_T]$ for every $t \leq T-1$, where $\{\phi_{T-t}^*\}_{t=0}^{T-1}$ solves the following recursion:

$$(15) \quad \phi_{T-t}^* = \delta \phi_{T-t-1}^* \quad \forall t \leq T-1,$$

with terminal condition $\phi_0^* = 1$. It follows that $\phi_T^* = \delta^T$, where, recall, $\delta \equiv \gamma + \alpha$ and where γ and α parameterize, respectively, PE and GE effects.

The exact interpretation of ϕ_T^* and its magnitude depend, of course, on the application under consideration. Consider, for example, the demand block of the New Keynesian model, in isolation of the supply block. In this context, ϕ_T^* measures the response of aggregate spending in period 0 to the concurrent expectation of the

¹⁶In line with this point, the obtained lessons are also robust to replacing Θ_t in condition (11) with $\theta_{i,t} = \Theta_t + v_{i,t}$, where $v_{i,t}$ is i.i.d. across agents, and letting each agent i know at $t = 0$ the *entire* path of $\theta_{i,t}$; that is, we could have eliminated the uncertainty the typical agent faces about the fundamentals that matter for her *own* decisions (e.g., the uncertainty a consumer faces about her own interest rates), and yet preserve her uncertainty about the beliefs and the decisions of *others* (e.g., the consumer's uncertainty about aggregate spending and inflation).

real interest rate in period T , holding constant the expectations of the real interest rate in all other periods. In the textbook version of the New Keynesian model, this object is the same regardless of T .¹⁷ In the variant studied by McKay, Nakamura, and Steinsson (2017), on the other hand, the corresponding object is decaying with T , due to liquidity constraints.

What is of interest to us, however, is not how ϕ_T^* varies as we move from one complete-information application to another, but rather how it compares to its incomplete-information counterpart. We address this question in the rest of the section. But let us first note three facts about ϕ_T^* . First, ϕ_T^* is invariant to the precision of the representative agent’s information that is available at $t = 0$. Second, ϕ_T^* is the same regardless of whether the representative agent expects to receive additional information between $t = 0$ and $t = T$ or not. Finally, the decomposition between PE and GE effects is irrelevant; all that matters is the combined effect, herein parameterized by $\delta \equiv \gamma + \alpha$.

The Frictional Case.—We henceforth focus on the scenario in which the agents face higher-order uncertainty and are therefore unable to reach a perfect consensus about the response of the economy to any given news about Θ_T . This raises the delicate question of whether and how the level of higher-order uncertainty evolves over time. For the time being, we bypass this complication by making the following simplifying assumption.

ASSUMPTION 1: *There is no learning between $t = 0$ and $t = T$, at which point Θ_T becomes commonly known.*

As noted a moment ago, learning is inconsequential in the complete-information benchmark. This is no more true when information is incomplete, for reasons discussed later. Nevertheless, Assumption 1 is a useful starting point for three reasons. First, it affords an otherwise general specification of the information structure and yields the sharpest version of our results. Second, it separates our results from prior works that study the implications of learning. Third, it can be relaxed without compromising our main lessons.

Let us fill in the remaining details. Notwithstanding Assumption 1, we accommodate an otherwise arbitrary information structure. We first let a random variable s encapsulate the realization of the fundamental Θ_T along with *any* other aggregate shock that influences the information and the beliefs of the agents; we refer to s as the underlying state of nature. We next represent the information of agent i by a random variable ω_i , which is itself correlated with s . Along with the fact that there is no learning, this means that, for all $t \leq T - 1$, $E_{i,t}[\cdot] = E_{i,0}[\cdot] = \mathbb{E}[\cdot | \omega_i]$, where $\mathbb{E}[\cdot | \omega_i]$ is the rational expectation conditional on ω_i . We refer to s as the aggregate state of nature and to ω_i as the signal of agent i . Conditional on s , this signal is an i.i.d. draw from a fixed distribution, whose PDF is given by $\phi(\omega | s)$. This signal therefore encodes the information that agent i has, not only about the fundamental Θ_T , but also about

¹⁷This is readily seen by iterating the Euler condition of the representative consumer. It also follows from our earlier observation that the demand block is nested with $\gamma = \beta$ and $\alpha = 1 - \beta$, which means that $\delta = 1$ and hence that ϕ_T^* does not vary with T .

the entire state of nature and thereby about the beliefs of other agents. In short, ω_i is the Harsanyi type of agent i , that is, an object that governs the agent's expectations of all the endogenous outcomes along any given equilibrium. We finally assume that a law of large number applies in the sense that $\phi(\omega|s)$ is also the cross-sectional distribution of the signals in the population when the aggregate state is s .

Because the state s can be an arbitrary random variable, the specification we have introduced above allows for rich first- and higher-order uncertainty and nests a variety of examples that can be found in the literature. For instance, we may let the aggregate state be $s = (\Theta_T, u)$, where $\Theta_T \sim N(0, \sigma_\theta^2)$ and $u \sim N(0, \sigma_u^2)$ are independent of one another, and specify the signal as $\omega_i = (z, x_i)$, where $z = \Theta_T + u$, $x_i = \Theta_T + v_i$, and $v_i \sim N(0, \sigma_v^2)$ is independent of s and i.i.d. across i . This case nests the information structure assumed in Morris and Shin (2002): each agent receives two signals about Θ_T , a private one given by x_i , and a public one given by z . Alternatively, we may modify the aforementioned case by letting $\omega_i = (z_i, x_i)$, with x_i as before and $z_i = z + \eta_i$, $\eta_i \sim N(0, \sigma_\eta^2)$. In this case, there is a public signal z , whose observation is, however, contaminated by idiosyncratic noise, perhaps due to rational inattention à la Sims (2003). As yet another example, we could let the aggregate state be $s = (\Theta_T, \sigma)$, where $\Theta_T \sim N(0, \sigma_\theta^2)$ and $\sigma \sim U[\underline{\sigma}, \bar{\sigma}]$, and specify the signal as $\omega_i = (x_i, \sigma_i)$, where $x_i = \Theta_T + v_i$, $v_i \sim N(0, \sigma_v^2)$, $\sigma_i \sim U[\sigma - \Delta, \sigma + \Delta]$, and $\Delta, \underline{\sigma}, \bar{\sigma}$ are known scalars such that $0 < \Delta < \underline{\sigma} < \bar{\sigma}$. In this case, some agents are better informed than others, and each agent is uncertain about how informed or uninformed the other agents are.

By adopting this level of generality, we seek, not only to clarify the robustness of our insights, but also to bypass the need for taking a specific stand on what the available signals are and how the higher-order uncertainty is generated. This also explains why most of our results in this section are formulated in terms of how outcomes depend on the belief hierarchy, as opposed to how they depend on a specific set of signals. A concrete example, however, will also be considered.

C. Attenuation and the Horizon Effect

An instrumental step toward addressing the question of interest is to understand the role of higher-order beliefs. Since Θ_t is fixed at 0 for all $t \neq T$ and Θ_T becomes commonly known at $t = T$, we have $a_t = 0$ for all $t \geq T + 1$, $a_T = \Theta_T$, and, for all $t \leq T - 1$,

$$(16) \quad a_t = \gamma^{T-t} \bar{E}_t[\Theta_T] + \alpha \sum_{k=1}^{T-t} \gamma^{k-1} \bar{E}_t[a_{t+k}].$$

By Assumption 1 (no learning), $\bar{E}_t[\cdot] = \bar{E}_0[\cdot]$ for all $t \leq T - 1$. Using this fact, and iterating the condition above backwards, we reach the following lemma, which represents the period-0 outcome as a linear function of the concurrent average first- and higher-order expectations of Θ_T .

LEMMA 2: For every (α, γ, T) , there exist positive scalars $\{\chi_{h,T}\}_{h=1}^T$ such that, regardless of the information structure,

$$(17) \quad a_0 = \sum_{h=1}^T \{\chi_{h,T} \cdot \bar{E}_0^h[\Theta_T]\},$$

where $\bar{E}_0^h[\cdot]$ is defined recursively by $\bar{E}_0^1[\cdot] = \bar{E}_0[\cdot]$ and $\bar{E}_0^h[\cdot] = \bar{E}_0[\bar{E}_0^{h-1}[\cdot]]$ for every $h \geq 2$.

The weights $\chi_{h,T}$ can be constructed recursively, as functions of α , γ , h , and T alone; see online Appendix A for details. Here, we focus on the interpretation and the implications of Lemma 2.

Applied to the demand block of the New Keynesian model, this lemma means that today’s aggregate spending is determined by the hierarchy of beliefs about future real interest rates. Applied to the supply block, it means that today’s inflation is determined by the hierarchy of beliefs about future real marginal costs. In both the textbook version of the New Keynesian model and in the richer dynamic stochastic general equilibrium (DSGE) versions used for quantitative policy evaluation, this kind of higher-order beliefs is “swept under the carpet” because the complete-information assumption lets higher-order beliefs collapse to first-order beliefs.

To see this point in the abstract context under consideration, note that when higher-order beliefs coincide with first-order beliefs, condition (17) reduces to $a_0 = (\sum_{h=1}^T \chi_{h,T}) \cdot \bar{E}_0[\Theta_T]$. Along with Lemma 1, this also means that the complete-information outcome satisfies the following restriction:

$$(18) \quad \phi_T^* = \sum_{h=1}^T \chi_{h,T}.$$

By the same token, the weights $\{\chi_{h,T}\}$ do not matter per se: all that matters for equilibrium behavior is their sum, ϕ_T^* .

When, instead, information is incomplete, higher-order beliefs no longer coincide with first-order beliefs. To understand how a_0 covaries with $\bar{E}_0[\Theta_T]$, we must therefore understand two things: first, how higher-order beliefs covary with first-order beliefs; and second, how the beliefs of different order load into a_0 , that is, what the structure of the weights $\{\chi_{h,T}\}$ is. We complete these tasks in the sequel: but first we state the ultimate lesson.

THEOREM 1: *Suppose that information is incomplete. There exists a positive scalar ϕ_T , which depends on both T and the information structure, such that the following properties hold:*

(i) *For all states of nature,*

$$(19) \quad a_0 = \phi_T \cdot \bar{E}_0[\Theta_T] + \epsilon,$$

where ϵ is either identically 0 or is random but orthogonal to $\bar{E}_0[\Theta_T]$.

(ii) *The ratio ϕ_T/ϕ_T^* is strictly less than 1.*

(iii) *The ratio ϕ_T/ϕ_T^* is strictly decreasing in T .*

Part (i) identifies ϕ_T as the incomplete-information counterpart of ϕ_T^* . Part (ii) establishes that the absence of common knowledge reduces the extent to which the

aggregate outcome covaries with the average forecast of the fundamental, regardless of how precise the latter is. We refer to this finding as the “attenuation effect.” Part (iii) establishes that this kind of attenuation increases with T : the longer the horizon, the larger the reduction in the responsiveness of the aggregate outcome to the forecasts of the fundamental relative to the frictionless benchmark. We refer to this finding as the “horizon effect.”

Let us now sketch the proof of the result. Part (i) is trivial: it follows directly from projecting a_0 on $\bar{E}_0[\Theta_T]$ and letting ϕ_T be the coefficient of this projection and ϵ the residual. The latter captures any variation in higher-order beliefs that is orthogonal to the variation in first-order beliefs, such as the kind of “sentiment shocks” studied in Angeletos and La’O (2013).

To prove parts (ii) and (iii), note first that Lemma 2 implies that the following condition holds regardless of the information structure:

$$(20) \quad \phi_T = \sum_{h=1}^T \chi_{h,T} \beta_h,$$

where $\{\chi_{h,T}\}$ are the same scalars as those appearing in Lemma 2 and

$$(21) \quad \beta_h \equiv \frac{\text{cov}(\bar{E}_0^h[\Theta_T], \bar{E}_0^1[\Theta_T])}{\text{var}(\bar{E}_0^1[\Theta_T])}$$

is the coefficient of the projection of $\bar{E}_0^h[\Theta_T]$ on $\bar{E}_0^1[\Theta_T]$, for any $h \geq 1$. When information is complete, $\bar{E}_0^h[\Theta_T]$ coincides with $\bar{E}_0^1[\Theta_T]$ for all h and all realizations of uncertainty, which in turn means that β_h is identically 1 for all h . Away from this fragile benchmark, the following is true.

PROPOSITION 3: *With incomplete information, β_h is bounded in $(0, 1)$ for all $h \geq 2$ and is strictly decreasing in h .*¹⁸

In words, higher-order beliefs comove less with first-order beliefs than lower-order beliefs. This result is easy to establish for the specific information structure studied in Morris and Shin (2002), but requires more work for the more general structure allowed here. It has a similar flavor as the result found in Samet (1998), which essentially states that higher-order beliefs are anchored to the common prior, but is distinct from it, because Samet’s result regards only the asymptotic properties of higher-order beliefs as $h \rightarrow \infty$ and contains no information on whether the comovement between first- and h -order forecasts varies monotonically with h .

Let us now go back to the proof of parts (ii) and (iii) of Theorem 1. Part (ii) now follows directly from the fact that $\beta_h \in (0, 1)$, along with conditions (18) and (20). In a nutshell, ϕ_T is less than ϕ_T^* simply because higher-order beliefs move less than one-to-one with first-order beliefs when, and only when, information is incomplete. Importantly, this is true regardless of how much the first-order beliefs themselves comove with the fundamental. The following is therefore also true, underscoring

¹⁸Obviously, for $h = 1$, we have $\beta_1 = 1$.

once again that the introduced friction has to do *only* with the accommodation of higher-order uncertainty and frictional coordination, not with the first-order uncertainty about Θ_T .

COROLLARY 1: *Fix (α, γ, T) . An incomplete-information economy features a lower response to expectations of future fundamentals (i.e., a lower ϕ_T) than a complete-information economy, even if the agents have more precise information about Θ_T in the former than in the latter.*

What remains to prove is part (iii), namely, the property that ϕ_T decreases with T relative to ϕ_T^* . This follows from combining our result that the comovement between first- and higher-order beliefs decreases with the belief order h (Proposition 3), with another result, which we establish next and which sheds light on how the horizon T affects the relative importance of higher-order beliefs in the period-0 outcome.

THEOREM 2: *Fix (α, γ) . For any (h, T) such that $1 \leq h \leq T$, let $s_{h,T}$ be the total weight on beliefs of order up to, and including, h ; that is, $s_{h,T} \equiv \sum_{r=1}^h \chi_{r,T}$, where the χ s are the same coefficients as those appearing in condition (17). The ratio $s_{h,T}/s_{T,T}$, which measures the relative contribution of the first h orders of beliefs to the aggregate outcome, strictly decreases with the horizon T and converges to 0 as $T \rightarrow \infty$.*

This result is crucial (which is why it qualifies, at least in our eyes, to be called a theorem). Longer horizons increase the number of loops from future aggregate actions to current actions. But when one increases the number of loops, one is effectively walking down the hierarchy of beliefs: forecasting outcomes further and further into the future maps to forecasting the forecasts of others at higher and higher orders. This in turn explains why longer horizons increase the relative importance of higher-order beliefs. Part (iii) of Theorem 1 then follows from combining this finding with the fact that high-order beliefs themselves covary less with first-order beliefs than do lower-order beliefs.

To sum up: the attenuation effect ($\phi_T/\phi_T^* < 1$) follows merely from the fact that incomplete information dampens the comovement of higher-order beliefs with first-order beliefs ($\beta_h < 1$ for all h). The horizon effect (ϕ_T/ϕ_T^* decreases with T), on the other hand, follows from the fact that this dampening increases with the order of belief (β_h decreases with h), together with the fact that longer horizons raise the relative importance of higher-order beliefs (Theorem 2).

D. The Limit as $T \rightarrow \infty$

We complement the preceding results by establishing that, as long as higher-order beliefs are sufficiently anchored in the sense made precise below, ϕ_T becomes vanishingly small relative to its complete-information counterpart as the horizon gets larger and larger.

PROPOSITION 4 (Limit): *If $\lim_{h \rightarrow \infty} \text{var}(\bar{E}^h[\Theta_T]) = 0$, then*

$$\lim_{T \rightarrow \infty} \frac{\phi_T}{\phi_T^*} = 0.$$

To understand this result, note first that, by Theorem 1, the ratio ϕ_T/ϕ_T^* is strictly decreasing in T and bounded in $(0, 1)$. It follows that this ratio necessarily converges to a constant $\varphi \in [0, 1)$ as $T \rightarrow \infty$. In online Appendix A, we show that $\varphi = \lim_{h \rightarrow \infty} \beta_h$, where β_h is defined as before. This fact is, essentially, a corollary of Theorem 2: in the limit as $T \rightarrow \infty$, only the infinite-order beliefs matter. Establishing that ϕ_T/ϕ_T^* converges to 0 as $T \rightarrow \infty$ is therefore equivalent to establishing that β_h converges to 0 as $h \rightarrow \infty$. A sufficient (in fact, also a necessary) condition for this to be the case is that $\lim_{h \rightarrow \infty} \text{var}(\bar{E}^h[\Theta_T]) = 0$.

This condition need not hold for every information structure, but can be said to be generic in the following sense. Take, as a reference point, the familiar case in which Θ_T is Normal and each agent observes a noisy private signal and a noisy public signal, as in Morris and Shin (2002). In this case, $\lim_{h \rightarrow \infty} \beta^h$ is strictly positive and is pinned down by the precision of the public signal. But now let us perturb this economy by allowing each agent's observation of the public signal to be contaminated by an idiosyncratic noise, perhaps due to rational inattention or some other cognitive limitation. Then, $\lim_{h \rightarrow \infty} \beta^h$ is necessarily 0, even if the aforementioned noise is *arbitrarily* small.¹⁹ In short, it takes a tiny perturbation to break the common knowledge nature of a public signal and to guarantee the limit result in Proposition 4 holds.

As noted in the introduction, this result illustrates how predictions that depend on long series of forward-looking, general equilibrium, feedback loops can be particularly fragile to relaxations of common knowledge. It also builds a bridge between our results and the uniqueness result in global games (Morris and Shin 1998, 2003): the common thread is the disproportionate effect of the (very) high orders of belief and the resulting discontinuity in small perturbations of common knowledge. Finally, it is worth noting that the result holds true even if ϕ_T^* itself explodes to infinity as $T \rightarrow \infty$. This case is relevant in the context of forward guidance, which we study in Section V.

E. The Friction as Extra Discounting at the Aggregate Level

The preceding analysis established our main results under a general specification of the information structure. We now use a tractable example to recast the friction under consideration as a certain form of myopia. This sheds further light on how the absence of common knowledge interacts with the underlying GE mechanisms and how it modifies forward-looking behavior at the aggregate level.

For this example, we let $\Theta_T = z + \eta$, where $z \sim N(0, \sigma_z^2)$ and $\eta \sim N(0, \sigma_\eta^2)$ are independent of one another. We interpret z as the component of Θ_T that is realized at $t = 0$ and η as a residual that is realized at $t = T$. Finally, we let the signal of agent i received at $t = 0$ be $\omega_i = z + v_i$, where $v_i \sim N(0, \sigma_v^2)$ is i.i.d across i and orthogonal to both z and η . We can then think of z as news about Θ_T and v_i as an idiosyncratic noise in the observation, or interpretation, of this news by agent i .²⁰

¹⁹This case is formalized by letting agent i observe $z_i = z + \eta_i$, where $z = \Theta_T + u$ is the underlying public signal and η_i is the contaminating idiosyncratic noise. Our statement is that $\lim_{h \rightarrow \infty} \beta_h = 0$ as long as $\text{var}(\eta_i)$ is not exactly 0.

²⁰The interpretation of this noise as the product of rational inattention is discussed at online Appendix B. For another complementary interpretation, see the discussion of Assumption 3 in Section V.

Regardless of its interpretation, the key property of this example is that it imposes an exponential structure in the belief hierarchy. By this we mean that, for every $h \geq 2$,

$$\bar{E}_0^h[\Theta_T] = \lambda \bar{E}_0^{h-1}[\Theta_T],$$

where $\lambda \equiv \sigma_v^{-2}/(\sigma_v^{-2} + \sigma_z^{-2}) \in (0, 1)$. This scalar therefore controls the speed with which the comovement between first and h -order beliefs decay with h . A lower λ captures a lower degree of common knowledge; the frictionless, complete-information benchmark is nested in the limit as $\lambda \rightarrow 1$.

The following result can then be shown.

PROPOSITION 5 (Discounting): *The elasticity ϕ_T in the incomplete-information economy described above is the same as the elasticity ϕ_T^* in a representative-agent economy in which, for all $t \leq T - 2$, the Euler condition (13) holds with $\delta \equiv \gamma + \alpha$ replaced by $\delta' \equiv \gamma + \alpha\lambda$.*

This result illustrates that, under appropriate conditions, the object of interest can be calculated with the same ease as in the complete-information benchmark: all one has to do is to hold the PE effect constant and to discount the GE effect by the factor $\lambda \in (0, 1)$. It also offers a sharp formalization of the idea that removing common knowledge causes the economy to act as if the representative agent is myopic and discounts the future more heavily than in the frictionless benchmark (the effective δ is reduced). The extent of this kind of myopia is inversely related to λ , the degree of common knowledge.

The logic behind this result is simple. In the example we consider, all the average higher-order forecasts are linear transformations of the average first-order forecast. This guarantees that, for all $t \in \{1, \dots, T - 1\}$, the aggregate outcome can be expressed as a multiple of the average first-order forecast: $a_t = \phi_{T-t} \bar{E}_0[\Theta_T]$, for some known scalar ϕ_{T-t} . It follows that $\bar{E}_{t-1}[a_t] = \phi_{T-t} \bar{E}_0^2[\Theta_T]$. And since $\bar{E}_0^2[\Theta_T] = \lambda \bar{E}_0[\Theta_T]$, we conclude that $\bar{E}_{t-1}[a_t] = \lambda a_t$. It is therefore *as if* the average agent systematically underestimates the variation in the future aggregate outcome, which in turn explains the result in Proposition 5. Clearly, the exact result relies on the particular information structure assumed above. The logic, however, applies more generally. As already mentioned, the key feature of the example we consider is that λ controls the speed with which β_h decays with h . The kind of myopia, or discounting, reported above is merely a manifestation of the more general property that β_h decays with h . Theorem 1 and Proposition 5 are therefore mirror images of each other: saying that there is more attenuation at longer horizons is the same as saying that the agents are, effectively, myopic.

Finally, note that, for any $\lambda < 1$, the gap between δ and δ' increases with α , which proves the following.

COROLLARY 2: *The documented GE attenuation and the associated as-if discounting are larger, the stronger the underlying GE effect.*

This result offers a sharp illustration of how the friction accommodated in this paper interacts with GE mechanisms. It also suggests that the applied lessons

delivered in Sections V and VI are likely to be reinforced if one extends the analysis to more realistic versions of the New Keynesian model that strengthen the relevant GE mechanisms by introducing short horizons, borrowing constraints, and hand-to-mouth consumers.²¹

F. Allowing for Learning

We now return to the role played by Assumption 1. Relaxing this assumption does not appear to invalidate the insights we have developed, but complicates the analysis and precludes us from establishing Theorem 1 for arbitrary information structures. This is due to the increased complexity of the kind of higher-order beliefs that emerge once information is changing between 0 and T .

To appreciate what we mean by this, note that the absence of Assumption 1 guarantees that the following properties hold for any $h \in \{2, \dots, T\}$ and any $\{t_2, t_3, \dots, t_h\}$ such that $0 < t_2 < \dots < t_h < T$:

$$\bar{E}_0[\bar{E}_{t_2}[\dots[\bar{E}_{t_h}[\Theta_T]\dots]]] = \bar{E}_0^h[\cdot].$$

That is, Assumption 1 helps collapse the “cross-period” higher-order beliefs to the “within-period” higher-order beliefs. This is the key step for obtaining the convenient representation of a_0 in Lemma 2 and, thereby, for applying Proposition 3. Without Assumption 1, we must instead express a_0 as a function of all the aforementioned kind of “cross-period” higher-order beliefs, which greatly complicates the analysis. Note, in particular, that there are $T - 1$ types of second-order beliefs (namely, $\bar{E}_0[\bar{E}_t[\cdot]]$ for all t such that $1 \leq t \leq T - 1$), plus $(T - 1) \times (T - 2)/2$ types of third-order beliefs, plus $(T - 1) \times (T - 2) \times (T - 3)/6$ types of fourth-order beliefs, and so on. What is more, the correlation structure between first- and higher-order beliefs is more intricate, reflecting the anticipation of learning. Intuitively, the first-order belief $\bar{E}_0[\Theta_{10}]$ can be *less* correlated with the second-order belief $\bar{E}_0[\bar{E}_1[\Theta_{10}]]$ than with the third-order belief $\bar{E}_0[\bar{E}_8[\bar{E}_9[\Theta_{10}]]]$ if there is little or no learning between periods 0 and 1 but a lot of learning between periods 1 and 8. While this possibility does not invalidate our lessons,²² it precludes us from extending Theorem 1 to arbitrary forms of learning.

It is possible, however, to overcome this caveat for two leading forms of learning studied in the literature. In the one, we let the agents become gradually aware of Θ_T , as in Mankiw and Reis (2002) and Wiederholt (2015). In the other, we allow the agents to receive a new private signal about Θ_T in each period prior to T , as in Woodford (2003a), Nimark (2008), and Mackowiak and Wiederholt (2009). In these two cases, we can not only recover Theorem 1, but also obtain a variant of Proposition 5. See online Appendix B for details.

²¹ Recall that the demand block of the New Keynesian model is nested with $\gamma = \beta$ and $\alpha = 1 - \beta$. The results of Galí, López-Salido, and Vallés (2007), Del Negro, Giannoni, and Patterson (2015), Werning (2015), Kaplan, Moll, and Violante (2016) and McKay, Nakamura, and Steinsson (2016, 2017) indicate that the aforementioned frictions map to a lower γ (weaker PE effect) but also a higher α (stronger GE effect). Corollary 2 then implies that these frictions are likely to reinforce our attenuation and horizon effects.

²² This possibility is indeed allowed in the cases studied in online Appendix B, where the theorem continues to apply.

Also note that Theorem 2, which establishes that longer horizons raise the relative importance of higher-order beliefs, follows directly from the best-response structure and is therefore independent of the information structure. This is suggestive of why our horizon effect may extend beyond the aforementioned cases. In line with this idea, we are able to show that the result of Proposition 4, namely the property that ϕ_T becomes vanishingly small relative to ϕ_T^* as $T \rightarrow \infty$, extends to *arbitrary* forms of learning as long as the higher-order uncertainty remains bounded away from 0, in a sense made precise in online Appendix B.

On the basis of these findings, we conclude that our lessons regarding the interaction of horizons and lack of common knowledge are robust to the introduction of learning. Two additional lessons, however, emerge once we take into account learning.

The first is that, holding the period-0 information constant, ϕ_T is closer to ϕ_T^* in the scenario in which subsequent learning is allowed relative to the scenario in which such learning is ruled out. Intuitively, the anticipation of more (less) information in the future eases (exacerbates) the friction in the present. Note that this effect hinges on the agents being forward-looking and is therefore absent in *static* beauty contests, such as those studied in Morris and Shin (2002), Woodford (2003a), and Angeletos and La'O (2010).

The second lesson is that, as time passes and agents accumulate more information, higher-order beliefs converge to first-order beliefs, causing the anchoring of the expectations and the attenuation of GE effects to decay with the lag of time since the news has arrived.²³ Although this prediction is not a core theme of our paper, it is worth noting that it is consistent with the available evidence on the impulse responses of expectations to identified shocks (Coibion and Gorodnichenko 2012; Vellekoop and Wiederholt 2017). By contrast, this evidence is hard to reconcile with the theories developed in Gabaix (2016) and Farhi and Werning (2017): by design, these theories attenuate the response of expectations to shocks but do not allow this attenuation to decay with the lag since the shock has hit the economy.

V. Revisiting Forward Guidance

In this section, we return to the New Keynesian model and study the implications of our insights for the effectiveness of forward guidance.²⁴

To begin with, suppose either that the monetary authority commits on implementing a specific path for the *real* interest rate, or that prices are infinitely rigid ($\theta = 1$) so that inflation is identically zero and the real interest rate coincides with the nominal one. In either case, we can characterize the response of the economy

²³Formally, consider the setting with learning studied in online Appendix B, pick any $\tau \in \{1, \dots, T-2\}$, and look at the response of a_t at $t = \tau$ rather than at $t = 0$. It is easy to check that this response increases with the precision of the information that arrives between $t = 0$ and $t = \tau$. A variant of this property is at the core of Woodford (2003a), Bacchetta and van Wincoop (2006), Nimark (2008, 2017), and Angeletos and La'O (2010): these papers show how, following persistent shocks to the underlying fundamentals, there is initially a large “wedge” between first- and higher-order beliefs, but this wedge decays as time passes and learning occurs.

²⁴Campbell et al. (2012) distinguish two types of forward guidance, “Delphic” and “Odyssean,” depending on whether the central bank communication is interpreted partly as a signal of the underlying state of the economy or exclusively as a commitment about future policy. Throughout our analysis, we are concerned with the second kind of forward guidance.

to news about future rates by studying the demand block alone. The results of the previous section then directly imply that the absence of common knowledge reduces this response, and the more so the further into the future the change in the real rate is. In this sense, the effectiveness of forward guidance is reduced.

The obvious caveat is that this argument treats the path of the *real* interest rate as the exogenous impulse or, equivalently, that it shuts down inflation. While this is useful for understanding how monetary policy controls aggregate demand holding constant inflation expectations, it is inappropriate for understanding how forward guidance matters during a liquidity trap. In that context, it is essential to treat the real interest rate as endogenous and to capture the feedback loop between aggregate spending and inflation.²⁵

As noted before, this feedback loop is a GE mechanism that runs *between* the two blocks of the model. This means that the analysis of Section IV, which effectively dealt with each block in isolation, is not directly applicable. To address this complication, let us assume momentarily that information is complete among the firms, even though it is incomplete among the consumers. In this special case, the standard NKPC remains valid, implying that inflation can be expressed as the present value of the future real marginal costs and, thereby, of future aggregate spending. Using this fact to substitute away inflation from condition (5), and letting $y_t = c_t$, we arrive at the following representation of the equilibrium.

LEMMA 3: *Suppose the firms have complete information. The equilibrium level of aggregate output (also, spending) satisfies the following condition at every t :*

$$(22) \quad y_t = -\sigma R_t - \sigma \sum_{k=1}^{\infty} \beta^k \bar{E}_t^c [R_{t+k}] + \sum_{k=1}^{\infty} (1 - \beta + k\sigma\kappa) \beta^{k-1} \bar{E}_t^c [y_{t+k}],$$

where $\kappa \equiv (\epsilon + 1/\sigma)\varkappa = (\epsilon + 1/\sigma)(1 - \theta)(1 - \beta\theta)/\theta$.

This can, once again, be understood as a beauty contest among the consumers. But unlike the one seen earlier in condition (5), the one obtained here subsumes the feedback loop between inflation and spending. This in turn explains why the expectations of income show up with different weights from those in condition (5), as well as why these weights depend on κ , which is the slope of the NKPC with respect to aggregate output.

Although the beauty contest obtained in Lemma 3 is not directly nested in the framework of Section IV, all the lessons continue to hold. Proposition 3, which regards the belief hierarchy, remains intact, while Lemma 2 and Theorem 2, which hinge on the dynamic structure, are extended in online Appendix C. It follows that Theorem 1 and all the other results of Section IV can be directly applied to the present context.

A Concrete Example.—To make things more concrete, and to accommodate the possibility that the firms, too, have incomplete information, we henceforth impose the following assumptions.

²⁵ Accordingly, for the rest of the analysis, we let $\theta < 1$ (equivalently, $\kappa > 0$).

ASSUMPTION 2 (Monetary Policy): *There exists a known $T \geq 2$ such that:*

- (i) *At any $t < T$, the nominal interest rate is pegged at 0.²⁶*
- (ii) *At any $t > T$, monetary policy replicates flexible-price outcomes.*
- (iii) *The period- T nominal rate is such that*

$$(23) \quad R_T = z + \eta,$$

where z and η are random variables, independent of one another and of any other shock in the economy, with $z \sim N(0, \sigma_z^2)$ and $\eta \sim N(0, \sigma_\eta^2)$. The former is realized at $t = 0$; the latter is realized at $t = T$ and is unpredictable prior to that point.

ASSUMPTION 3 (Information):

- (i) *At $t = 0$, each agent i , be it a consumer or a firm, observes a private signal of z , given by*

$$\omega_i = z + v_i,$$

where v_i is an idiosyncratic noise, Normally distributed, with mean 0 and variance $\sigma_c^2 \geq 0$ or $\sigma_f^2 \geq 0$, depending on whether the agent is, respectively, a consumer or a firm.

- (ii) *No other exogenous information arrives until $t = T$, at which point R_T becomes common knowledge.*
- (iii) *The volatilities of the markup and idiosyncratic shocks are arbitrarily large relative to that of z .*

Assumption 2 specifies the monetary policy and identifies T as the horizon of forward guidance. Part (i) is motivated by the idea that the ZLB constraint is binding during a liquidity trap. Part (ii), on the other hand, permits us to concentrate on expectations of R_T , the nominal interest rate that the monetary authority will implement in the first period after the ZLB has ceased to bind.²⁷ Finally, part (iii) splits the randomness of R_T into two orthogonal components, along the lines of the example considered in the end of the last section. The first component, z , can be interpreted as the anticipated component of R_T , or as the news about future monetary policy; such news could be the product of a policy announcement. The second component, η , captures the residual uncertainty, or the unanticipated component, which is revealed at $t = T$.

²⁶Since R_t is the log-deviation from steady-state gross nominal rate, this means $R_t = -\rho$, where ρ is the discount rate defined by $\beta = e^{-\rho}$.

²⁷Needless to say, we let the central bank keep the interest rate low for just one period after the ZLB has ceased to bind, as opposed to many periods, only to simplify the exposition.

Assumption 3 turns to the information structure. Part (i) lets each agent's observation of z be contaminated by an idiosyncratic noise, thereby introducing higher-order uncertainty. Part (ii) shuts down any *exogenous* learning. Finally, part (iii) shuts down the *endogenous* learning that obtains through the observation of the realized market outcomes (wages, prices, etc.) in every period.

Part (i) of Assumption 3 may appear to be at odds with the fact that central bank communications enjoy ample coverage in the financial press; one may instead be tempted to treat forward guidance as a public signal. But it is one thing to say that something is "public news" in the real world and it is a different thing to assume that something is a "public signal" in the theory. Doing the latter requires, not only that every agent herself is aware of and attentive to the news, but also that she is fully confident that every other agent is also aware of and attentive to the news, and so on. Clearly, that kind of common knowledge does not have an obvious counterpart in the real world. What is more, even if one insists on modeling the central bank communications as a perfectly public signal, we can recast the required idiosyncratic noise as idiosyncratic variation in the *interpretation* of such communications.²⁸

To sum up, part (i) of Assumption 3 captures the friction we are interested in. Parts (ii) and (iii), on the other hand, are mostly for convenience: they facilitate a sharp characterization of the belief hierarchy.²⁹ In particular, similarly to the example studied in Section IV, the higher-order beliefs of, respectively, the consumers and the firms satisfy, for all $h \geq 2$,

$$\bar{E}_0^{c,h}[R_T] = \lambda_c^{h-1} \cdot \bar{E}_0^c[R_T] \quad \text{and} \quad \bar{E}_0^{f,h}[R_T] = \lambda_f^{h-1} \cdot \bar{E}_0^f[R_T],$$

where

$$\lambda_c \equiv \frac{\sigma_c^{-2}}{\sigma_c^{-2} + \sigma_z^{-2}} \in (0, 1] \quad \text{and} \quad \lambda_f \equiv \frac{\sigma_f^{-2}}{\sigma_f^{-2} + \sigma_z^{-2}} \in (0, 1].$$

Note that λ_c and λ_f control how much higher-order beliefs comove with first-order beliefs; they therefore parameterize the friction of interest. The frictionless, complete-information benchmark is nested with $\lambda_c = \lambda_f = 1$, and a larger friction corresponds to lower values for the λ s.

We are now ready to study how variation in expectations of R_T (triggered by variation in z) translates into variation in equilibrium output. We start with the common knowledge (CK) benchmark.

PROPOSITION 6 (Forward Guidance with CK): *Suppose $\lambda_c = \lambda_f = 1$, which means that z is common knowledge. There exists a scalar $\phi_T^* > \sigma$ such that, for all realizations of uncertainty,*

$$(24) \quad y_0 = y_0^{trap} - \phi_T^* \cdot E_0[R_T],$$

²⁸Formally, this can be done by modeling the central bank communication itself as a public signal and by letting each agent have private information about the noise in that signal.

²⁹Learning can be accommodated along the lines of online Appendix B and was indeed allowed in the first NBER version of our paper (Angeletos and Lian 2016b).

where y_0^{trap} is the liquidity-trap level of output (i.e., the one obtained when the period- T nominal interest rate is fixed at the steady-state value). Furthermore, ϕ_T^* is strictly increasing in T and $\phi_T^* \rightarrow \infty$ as $T \rightarrow \infty$.

This result contains the predictions of the textbook New Keynesian model regarding the power of forward guidance during a liquidity trap: this power, as measured by ϕ_T^* , is predicted to increase without bound as the time of action is pushed further and further into the future. It is this prediction, along with its quantitative evaluation, that constitutes the forward-guidance puzzle.³⁰

Contrast this prediction with the PE effect of future interest rates: as evident in condition (5), the PE effect is decreasing in T , simply because of the discounting embedded in intertemporal preferences. The reason that ϕ_T^* exhibits the opposite pattern is because of the GE effects that run within and between the two blocks of the model. In particular, the feedback loop between aggregate spending and income implies that ϕ_T^* would stay constant with T even if we were to shut down the inflation response. The feedback loop between aggregate spending and inflation then explains why ϕ_T^* actually increases with T .

Let us elaborate. Reducing the interest rate at $t = T$ increases spending and causes inflation at $t = T$. Because the nominal interest rate is pegged prior to T , this translates to a low *real* interest rate between $T - 1$ and T . This stimulates demand at $T - 1$, contributing to even higher inflation at $T - 1$, which feeds to even higher demand at $T - 2$, and so on. A longer horizon therefore maps to a larger number of iterations in this feedback loop and, thereby, to a stronger cumulative effect at $t = 0$. Finally, as $T \rightarrow \infty$, this feedback loop explodes, which explains why ϕ_T^* increases without bound.

The fact that the forward-guidance puzzle is driven by GE effects is already recognized in the literature: see, for example, the discussions in Del Negro, Giannoni, and Patterson (2015) and McKay, Nakamura, and Steinsson (2016). Our contribution is to explain how these GE effects depend on higher-order beliefs, which in turn explains why these effects are not as potent once one moves away from the common knowledge benchmark.

PROPOSITION 7 (Forward Guidance without CK): *Suppose that $\lambda_c < 1$ and/or $\lambda_f < 1$, which means that at least one group of agents lacks common knowledge of the news. There exists a scalar ϕ_T such that, for all realizations of uncertainty,*

$$(25) \quad y_0 = y_0^{trap} - \phi_T \cdot \bar{E}_0^c[R_T].$$

Furthermore, the following properties hold:

- (i) ϕ_T is bounded between the PE effect and the complete-information counterpart: $\sigma\beta^T < \phi_T < \phi_T^*$.

³⁰It is worth noting that the scalar ϕ_T^* is invariant to the ratio σ_z/σ_η , which parameterizes the precision of the available news: varying the ratio σ_z/σ_η affects how much $E_0[R_T]$ varies with R_T , but does not affect how much either y_0 or π_0 vary with $E_0[R_T]$. This property extends to the incomplete-information scenario studied next and explains the sense in which both the puzzle and the resolution we offer are orthogonal to the question of how accurate or credible the news about the future interest rates might be.

- (ii) ϕ_T is strictly increasing in both λ_c and λ_f ; the ratio ϕ_T/ϕ_T^* is strictly decreasing in T and converges to 0 as $T \rightarrow \infty$; finally, when λ_c is sufficiently low, ϕ_T also converges to 0 as $T \rightarrow \infty$.

Part (i) formalizes the sense in which the standard model “maximizes” the power of forward guidance: ϕ_T^* is an upper bound to the prediction that the analyst can make if she maintains the rational expectations hypothesis but is agnostic about the degree of common knowledge in the economy. By varying the degree of common knowledge, we effectively span the range between this upper bound and the underlying PE effect.

Part (ii) is a version of our earlier horizon effect: by attenuating the feedback loop between inflation and spending, as well as the GE effects that run within each block of the model, lack of common knowledge reduces the power of forward guidance by a factor that increases with T . This attenuation increases without bound in the sense that ϕ_T becomes vanishingly small relative to ϕ_T^* as $T \rightarrow \infty$, even if the friction is small (i.e., if λ_c and λ_f are arbitrarily close to 1). Finally, if the friction is large enough, the documented effect can be strong enough that ϕ_T is decreasing in T , not only relative to ϕ_T^* , but also in absolute value.

A Numerical Illustration.—We now use a numerical example to illustrate our findings. We interpret the period length as a quarter and adopt the calibration of the textbook New Keynesian model found in Galí (2008). That is, we set the discount factor to 0.99, the Frisch elasticity to 1, the elasticity of intertemporal substitution to 1, and the price revision rate to 1/3.

What remains is to pick the values of λ_c and λ_f , that is, the departure from common knowledge. The existing literature offers little guidance on how to make this choice. In want of a better alternative, we let $\lambda_c = \lambda_f = 0.75$ and interpret this as a situation in which every agent who has heard the policy announcement believes that any other agent has failed to hear, or “trust,” the announcement with a probability equal to 25 percent. This is arguably a modest “grain of doubt” in the minds of people about their ability to coordinate the adjustment in their beliefs and their behavior.^{31,32}

The solid line in Figure 1 plots the resulting attenuation effect, as measured by the ratio ϕ_T/ϕ_T^* , against the horizon length, T . By setting $\lambda_c = \lambda_f$, this line assumes that the consumers and the firms are subject to the same informational friction. The dashed line isolates the friction in the consumer side ($\lambda_c = 0.75$) by shutting it down in the firm side ($\lambda_f = 1$). The dotted line does the converse.

The attenuation is strongest when the friction is present in both sides. Furthermore, the effect is quantitatively significant. For example, at a horizon of five years ($T = 20$), the power of forward guidance is only *one-tenth* of its

³¹ The proposed reinterpretation of the informational friction is exact: if we let the signal of an agent be binary, revealing the true z with probability λ , we obtain exactly the same characterization for ϕ_T as the one under the signal structure assumed above.

³² Although Coibion and Gorodnichenko (2012) provide evidence in support of the kind of informational friction we have accommodated here, that evidence does not directly relate to forward guidance. Yet, taking that evidence at face value could justify an even larger departure from common knowledge: whenever a shock hits the economy, the average forecast error in the expectations of inflation appears to be one-half as large as the actual response in inflation, which translates to a value for λ close to 0.5.




FIGURE 1. ATTENUATION EFFECT AT DIFFERENT HORIZONS

common knowledge counterpart. This is on top of any mechanical effect that the noise may have on the size of the shift of the expectations of future interest rates: by construction, the documented attenuation effect is normalized by the size of the variation in the first-order beliefs of R_T .³³

This illustrates the following broader point: by anchoring the movements in the expectations of economic outcomes *relative* to the movements in the expectations of the policy instrument, our paper helps operationalize the idea that policymakers may have a harder time managing the former kind of expectations than the latter. This may help explain, for example, why forward guidance may trigger a large movement in the term structure without a commensurate movement in expectations of inflation and income.

Online Appendix C elaborates on the mechanics behind Figure 1. We first show how the documented effects can be understood through the lenses of a discounted Euler condition and a discounted NKPC, along the lines of our earlier, more abstract, result in Proposition 5. This representation helps connect our paper to McKay, Nakamura, and Steinsson (2016, 2017) and Gabaix (2016). It also helps explain why almost all of the effect seen in Figure 1 comes from the attenuation of two GE mechanisms: the price-setting complementarity inside the supply block, and the inflation-spending feedback between the two blocks. By contrast, the attenuation of the Keynesian multiplier that runs inside the demand block plays a small role.

³³Estimated DSGE models typically assume a higher degree of price stickiness than the one assumed here. This tends to reduce the attenuation effect. For instance, raising the value of θ from $2/3$, the value in Galí (2008), to 0.85, the value in Christiano, Eichenbaum, and Rebelo (2011), implies that the ratio ϕ_T/ϕ_T^* increases from about 0.1 to about 0.3 at $T = 20$. That said, note that menu-cost models calibrated to microdata indicate that the “right” value for θ is probably even smaller than $2/3$. In any event, our goal here is only to illustrate; a more comprehensive quantitative evaluation is left for future work.

Let us explain the last point. In the textbook version of the New Keynesian model, consumers have infinite horizons and their spending depends on aggregate spending *only* through the present value of permanent income. What is more, the discount rate is close to 0. This means that varying the expectations of income in the next, say, five years has a small effect on current spending. It follows that the attenuation of this particular GE mechanism cannot possibly be quantitatively significant in the textbook version of the model. But now note that the Keynesian multiplier can be reinforced by short horizons, hand-to-mouth consumers, countercyclical precautionary motives, and feedback effects between housing prices and consumer spending. By Corollary 2, we expect this to reinforce also the associated attenuation effect.³⁴

The following clarification is also worth making. In our analysis, we have abstracted incomplete markets in the sense of borrowing constraints, but have allowed incomplete markets in the sense of ruling out risk sharing. Had we allowed the consumers to pool their resources at the end of the liquidity trap and to enjoy the same level of consumption thereafter, we would have effectively eliminated the coordination friction among them, thus also abstracting from the attenuation of the GE effect that runs inside the demand block of the model. This is the scenario considered in Wiederholt (2015).

VI. Revisiting Fiscal Multipliers

We now introduce government spending in the model and study how aggregate income responds to news about government spending in the future. As usual, we treat the level of government spending as exogenous and assume that it is financed with lump sum taxation. But unlike the standard practice, we remove common knowledge of the future level of government spending and of its macroeconomic effects.

The equilibrium characterization given in Proposition 1 still holds, except that the real marginal cost is now given by $mc_t = \Omega_c c_t + (1 - \Omega_c) g_t$, where g_t is the level of government spending and $\Omega_c \in (0, 1)$ is a constant that increases with the steady-state fraction of GDP absorbed by private consumption. Note then that g_t enters the joint dynamics of c_t and π_t only through the aforementioned formula for mc_t . An exogenous shock to government spending is therefore akin to an exogenous shock in real marginal costs.

Pick now any pair (T, T') such that $2 \leq T \leq T'$. We want T' to capture the length of the liquidity trap and T the period during which a fiscal stimulus is going to be enacted.³⁵ Accordingly, we let monetary policy satisfy Assumption 2, with T replaced by T' . We also fix $R_{T'} = 0$, so as to focus on fiscal policy. Finally, we assume that $g_t = 0$ for all $t \neq T$ so as to focus on beliefs about g_T , and specify the information structure in the same way as in Assumption 3. We can then show the following.

³⁴Farhi and Werning (2017) make a similar point with regard to the interaction of borrowing constraints with Level- k thinking.

³⁵Similarly to Christiano, Eichenbaum, and Rebelo (2011), Woodford (2011), and Werning (2011), the fiscal multipliers we are concerned with are those that obtain when the ZLB constraint binds. The investigation of how our insights extend, when the economy is away from this constraint and monetary policy follows a Taylor rule, is the subject of ongoing work.

PROPOSITION 8: *Propositions 6 and 7 hold true with ϕ_T^* and ϕ_T reinterpreted as, respectively, the complete- and the incomplete-information slope of y_0 with respect to $\bar{E}_0^c[g_T]$.*

The part of this result that regards ϕ_T^* is the fiscal analogue of the forward-guidance puzzle: in the common knowledge benchmark, the current impact of a fiscal stimulus of a given size increases as the stimulus is pushed further into the future. The underlying logic is essentially the same as before: the PE effect of a fiscal stimulus, which is its direct effect on the concurrent real marginal costs, decays with the horizon T , but the GE effects that run within and between the two blocks of the model overturn this property. The other part of the result, which regards ϕ_T , contains the policy lesson of this section: once we relax the assumption that there is common knowledge of the size and the consequences of the fiscal stimulus, the fiscal multiplier is reduced at every T , and the more so the larger T is.

In short, while the standard model predicts that fiscal stimuli must be back-loaded in order to pile up the feedback effects between aggregate spending and inflation, our approach offers a rationale for front-loading them. Such front-loading eases the coordination friction in the economy: the sooner the fiscal stimulus is enacted, the less concerned the agents would be about the beliefs and the responses of others at long horizons, and, thereby, the easier it is to coordinate on a large response today.³⁶ An interesting question, which we leave open for future research, is how this insight interacts with borrowing constraints and hand-to-mouth consumers, which is the more standard rationale for the front-loading of fiscal stimuli.

VII. Conclusion

Modern macroeconomics assigns a crucial role to forward-looking expectations, such as consumer expectations of future income and future real interest rates or firm expectations of future inflation and future real marginal costs. This property seems desirable and realistic. However, by assuming common knowledge along with rational expectations, the dominant modeling practice hardwires a certain kind of perfection in the ability of economic agents to understand what happens around them, to align their forward-looking beliefs, and to coordinate their responses to any exogenous impulse. In so doing, it also maximizes the general equilibrium multipliers on fiscal and monetary policies.

Conversely, allowing for higher-order uncertainty helps accommodate a realistic friction in the ability of economic agents to forecast, or comprehend, the macroeconomic effects of policy news. This in turn arrests the underlying general equilibrium feedback loops and reduces the ability of policymakers to steer the economy, especially when the policy news regards longer horizons.

We first formalized these ideas within an abstract framework, which was flexible enough to nest the demand and the supply block of the New Keynesian model along

³⁶Of course, whether the provided rationale for front-loading is sufficiently strong to offset the standard back-loading property depends on the severity of the informational friction and all the parameters that govern the magnitude of the GE feedback loops, such as κ , the slope of the NKPC. Letting λ or κ be small enough guarantees that the fiscal multiplier decreases with T , not only relative to its common knowledge counterpart, but also in absolute value.

with other applications. We next showed how these ideas lessen the forward-guidance puzzle and offer a rationale for the front-loading of fiscal stimuli. As shown in online Appendix C, the same logic also helps moderate the paradox of flexibility, namely the prediction that, under certain conditions, a higher degree of price flexibility can amplify demand shocks and raise the effectiveness of monetary policy.

By anchoring the movements in the expectations of economic outcomes *relative* to the movements in the expectations of the policy instrument, our paper helps operationalize the idea that policymakers may have a harder time managing the former kind of expectations than the latter. But it also hints to the following possibility: especially when it comes to longer horizons, it may be more important for the policymaker to communicate her intended path for the economy (e.g., for output and inflation) than her intended path for the policy instrument. What is more, these communications must be loud and clear, not only in the ears of Wall Street (financial markets), but also in the ears of Main Street (firms and consumers). We leave the exploration of these ideas open for future research.

Another interesting research question, which we are pursuing in ongoing work, is the investigation of the implications of our insights for equilibrium selection, for the determinacy properties of Taylor rules, and for the “neo-Fisherian” effects discussed in, *inter alia*, García-Schmidt and Woodford (2015) and Cochrane (2016, 2017). As already noted, these issues hinge on the same kind of higher-order beliefs as the ones we have unearthed in this paper.

Finally, it is worth clarifying the following point. In the present paper, we focused on environments in which GE effects reinforce PE effects (equivalently, they are the source of strategic complementarity). In this case, GE attenuation translates to underreaction. In other contexts, GE effects mitigate PE effects (they are the source of strategic substitutability). For example, this is the case when agents compete for limited resources and their actions are strategic substitutes. In this case, our insights remain valid, but their empirical implication is reversed: GE attenuation now translates to *over*reaction. In a companion paper (Angeletos and Lian 2016a), we discuss how this observation offers a unified explanation to seemingly disparate phenomena.

REFERENCES

- Allen, Franklin, Stephen Morris, and Hyun Song Shin. 2006. “Beauty Contests and Iterated Expectations in Asset Markets.” *Review of Financial Studies* 19 (3): 719–52.
- Andrade, Philippe, Gaetano Gaballo, Eric Mengus, and Benoit Mojon. 2015. “Forward Guidance and Heterogeneous Beliefs.” Unpublished.
- Angeletos, George-Marios, and Zhen Huo. 2018. “Myopia and Anchoring.” National Bureau of Economic Research Working Paper 24545.
- Angeletos, George-Marios, and Jennifer La’O. 2010. “Noisy Business Cycles.” In *NBER Macroeconomics Annual 2009*, Vol. 24, edited by Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, 319–78. Chicago: University of Chicago Press.
- Angeletos, George-Marios, and Jennifer La’O. 2013. “Sentiments.” *Econometrica* 81 (2): 739–79.
- Angeletos, George-Marios, and Chen Lian. 2016a. “Dampening General Equilibrium: From Micro to Macro.” National Bureau of Economic Research Working Paper 23379.
- Angeletos, George-Marios, and Chen Lian. 2016b. “Forward Guidance without Common Knowledge.” National Bureau of Economic Research Working Paper 22785.
- Angeletos, George-Marios, and Chen Lian. 2016c. “Incomplete Information in Macroeconomics: Accommodating Frictions in Coordination.” In *Handbook of Macroeconomics*, Vol. 2, edited by John B. Taylor and Harald Uhlig, 1065–240. Amsterdam: Elsevier.

- Angeletos, George-Marios, and Alessandro Pavan.** 2007. "Efficient Use of Information and Social Value of Information." *Econometrica* 75 (4): 1103–42.
- Bacchetta, Philippe, and Eric van Wincoop.** 2006. "Can Information Heterogeneity Explain the Exchange Rate Determination Puzzle?" *American Economic Review* 96 (3): 552–76.
- Bergemann, Dirk, and Stephen Morris.** 2013. "Robust Predictions in Games with Incomplete Information." *Econometrica* 81 (4): 1251–308.
- Campbell, Jeffrey R., Charles L. Evans, Jonas D. M. Fisher, and Alejandro Justiniano.** 2012. "Macroeconomic Effects of Federal Reserve Forward Guidance." *Brookings Papers on Economic Activity* 43 (1): 1–80.
- Chahrour, Ryan, and Gaetano Gaballo.** 2017. "Learning from Prices: Amplification and Business Fluctuations." European Central Bank Working Paper 2053.
- Christiano, Lawrence, Martin Eichenbaum, and Sergio Rebelo.** 2011. "When Is the Government Spending Multiplier Large?" *Journal of Political Economy* 119 (1): 78–121.
- Chung, Hess, Edward Herbst, and Michael T. Kiley.** 2015. "Effective Monetary Policy Strategies in New Keynesian Models: A Reexamination." In *NBER Macroeconomics Annual 2014*, Vol. 29, edited by Jonathan A. Parker and Michael Woodford, 289–344. Chicago: University of Chicago Press.
- Cochrane, John H.** 2016. "Do Higher Interest Rates Raise or Lower Inflation?" Unpublished.
- Cochrane, John H.** 2017. "The New-Keynesian Liquidity Trap." *Journal of Monetary Economics* 92: 47–63.
- Coibion, Olivier, and Yuriy Gorodnichenko.** 2012. "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy* 120 (1): 116–59.
- Coibion, Olivier, and Yuriy Gorodnichenko.** 2015. "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts." *American Economic Review* 105 (8): 2644–78.
- Del Negro, Marco, Marc Giannoni, and Christina Patterson.** 2015. "The Forward Guidance Puzzle." Federal Reserve Bank of New York Staff Report 574.
- Farhi, Emmanuel, and Iván Werning.** 2017. "Monetary Policy, Bounded Rationality, and Incomplete Markets." National Bureau of Economic Research Working Paper 23281.
- Gabaix, Xavier.** 2016. "A Behavioral New Keynesian Model." National Bureau of Economic Research Working Paper 22954.
- Gali, Jordi.** 2008. *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*. Princeton, NJ: Princeton University Press.
- Gali, Jordi, J. David López-Salido, and Javier Vallés.** 2007. "Understanding the Effects of Government Spending on Consumption." *Journal of the European Economic Association* 5 (1): 227–70.
- García-Schmidt, Mariana, and Michael Woodford.** 2015. "Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis." National Bureau of Economic Research Working Paper 21614.
- Grossman, Sanford J., and Joseph E. Stiglitz.** 1980. "On the Impossibility of Informationally Efficient Markets." *American Economic Review* 70 (3): 393–408.
- Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante.** 2016. "Monetary Policy According to HANK." National Bureau of Economic Research Working Paper 21897.
- Kiley, Michael T.** 2016. "Policy Paradoxes in the New Keynesian Model." *Review of Economic Dynamics* 21: 1–15.
- Lucas, Robert E., Jr.** 1972. "Expectations and the Neutrality of Money." *Journal of Economic Theory* 4 (2): 103–24.
- Mackowiak, Bartosz, and Mirko Wiederholt.** 2009. "Optimal Sticky Prices under Rational Inattention." *American Economic Review* 99 (3): 769–803.
- Mankiw, N. Gregory, and Ricardo Reis.** 2002. "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics* 117 (4): 1295–328.
- McKay, Alisdair, Emi Nakamura, and Jón Steinsson.** 2016. "The Power of Forward Guidance Revisited." *American Economic Review* 106 (10): 3133–58.
- McKay, Alisdair, Emi Nakamura, and Jon Steinsson.** 2017. "The Discounted Euler Equation: A Note." *Economica* 84 (336): 820–31.
- Morris, Stephen, and Hyun Song Shin.** 1998. "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks." *American Economic Review* 88 (3): 587–97.
- Morris, Stephen, and Hyun Song Shin.** 2002. "Social Value of Public Information." *American Economic Review* 92 (5): 1521–34.
- Morris, Stephen, and Hyun Song Shin.** 2003. "Global Games: Theory and Applications." In *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*, Vol. 1, edited by Mathias Dewatripont, Lars Peter Hansen, and Stephen J. Turnovsky, 56–114. New York: Cambridge University Press.

- Morris, Stephen, and Hyun Song Shin.** 2006. "Inertia of Forward-Looking Expectations." *American Economic Review* 96 (2): 152–57.
- Nimark, Kristoffer.** 2008. "Dynamic Pricing and Imperfect Common Knowledge." *Journal of Monetary Economics* 55 (2): 365–82.
- Nimark, Kristoffer.** 2017. "Dynamic Higher Order Expectations." Unpublished.
- Samet, Dov.** 1998. "Iterated Expectations and Common Priors." *Games and Economic Behavior* 24 (1–2): 131–41.
- Sims, Christopher A.** 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50 (3): 665–90.
- Singleton, Kenneth J.** 1987. "Asset Prices in a Time Series Model with Disparately Informed, Competitive Traders." In *New Approaches to Monetary Economics*, edited by William Barnett and Kenneth Singleton, 249–72. Cambridge, MA: Cambridge University Press.
- Tirole, Jean.** 2015. "Cognitive Games and Cognitive Traps." Unpublished.
- Vellekoop, Nathanael, and Mirko Wiederholt.** 2017. "Inflation Expectations and Choices of Households." Unpublished.
- Werning, Iván.** 2011. "Managing a Liquidity Trap: Monetary and Fiscal Policy." National Bureau of Economic Research Working Paper 17344.
- Werning, Iván.** 2015. "Incomplete Markets and Aggregate Demand." National Bureau of Economic Research Working Paper 21448.
- Wiederholt, Mirko.** 2015. "Empirical Properties of Inflation Expectations and the Zero Lower Bound." Unpublished.
- Woodford, Michael.** 2003a. "Imperfect Common Knowledge and the Effects of Monetary Policy." In *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, edited by Philippe Aghion et al., 25–58. Princeton, NJ: Princeton University Press.
- Woodford, Michael.** 2003b. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, NJ: Princeton University Press.
- Woodford, Michael.** 2011. "Simple Analytics of the Government Expenditure Multiplier." *American Economic Journal: Macroeconomics* 3 (1): 1–35.