Myopia and Anchoring

By George-Marios Angeletos and Zhen Huo[*]

We develop an equivalence between the equilibrium effects of incomplete information and those of two behavioral distortions: myopia, or extra discounting of the future; and anchoring of current behavior to past behavior, as in models with habit persistence or adjustment costs. We show how these distortions depend on higher-order beliefs and GE mechanisms, and how they can be disciplined by evidence on expectations. We finally illustrate the use of our toolbox with a quantitative application in the context of inflation, a bridge to the HANK literature, and an extension to networks.
JEL: C72, D80, E30, E7
Keywords: Informational Frictions, Dynamic Beauty Contests, Micro to Macro, DSGE

I. Introduction

What are the macroeconomic effects of informational frictions? How do they depend on general equilibrium (GE) mechanisms, market structures, and agent heterogeneity? And how can they be quantified?

We develop a toolbox for addressing such questions and illustrate its use. On the theoretical front, we offer an illuminating representation result and draw connections to the literatures on networks and HANK models. On the quantitative front, we show how to extract the informational friction from survey evidence on expectations and proceed to argue that it can rationalize sizable sluggishness in the response of inflation and aggregate spending to shocks.

Framework

Our starting point is a representative-agent model, in which an endogenous outcome of interest, denoted by a_t, obeys the following law of motion:

$$(1) \qquad a_t = \varphi \xi_t + \delta \mathbb{E}_t \left[a_{t+1} \right],$$

where ξ_t is the underlying stochastic impulse, or fundamental, $\varphi > 0$ and $\delta \in (0, 1]$ are fixed scalars, and $\mathbb{E}_t[\cdot]$ is the rational expectation of the representative agent.

Condition (1) stylizes a variety of applications. In the textbook New Keynesian model, this condition could be either the New Keynesian Philips Curve (NKPC), with a_t standing for inflation and ξ_t for the real marginal cost, or the Euler condition of the representative consumer (a.k.a. the Dynamic IS curve), with a_t standing for aggregate spending and ξ_t for the real interest rate. Alternatively, this condition can be read as an asset-pricing equation, with ξ_t standing for the asset's dividend and a_t for its price.

We depart from these benchmarks by letting people have a noisy "understanding" of the economy, in the sense of incomplete information. The friction could be the product of either dispersed knowledge (Lucas, 1972) or rational inattention (Sims, 2003). And it is the source of both first- and higher-order uncertainty. Relative to the frictionless, full-information, rational-expectations benchmark, there is therefore not only gradual learning of the exogenous innovations, but also a friction in how people reason about others (Morris and Shin, 1998; Tirole, 2015) and thereby about GE effects (Angeletos and Lian, 2018).

An Observational Equivalence

Our main result is a representation of the equilibrium effects of the informational friction in terms of two behavioral distortions. Under appropriate assumptions, the equilibrium dynamics of the aggregate outcome a_t in the incomplete-information economy are shown to coincide with that of a representative-agent economy in which condition (1) is modified as follows:

$$(2) \qquad a_t = \varphi \xi_t + \delta \omega_f \mathbb{E}_t \left[a_{t+1} \right] + \omega_b a_{t-1},$$

for some $\omega_f < 1$ and $\omega_b > 0$. The first distortion ($\omega_f < 1$) represents myopia towards the future, the second ($\omega_b > 0$) anchors current outcomes to past outcomes. One dulls the forward-looking behavior, the other adds a backward-looking element akin to habit or adjustment costs.

Crucially, both distortions increase not only with the level of noise but also with parameters that regulate the strategic interaction, or the GE feedback in the economy. Economies in which the Keynesian cross is steeper, firms are more strategic, or input-output linkages are stronger behave *as if* they are populated by more impatient and more backward-looking agents.

Underlying insights and marginal contribution

The documented effects encapsulate the role of higher-order beliefs. To fix ideas, consider the response of aggregate demand $(= a_t)$ to a drop in the real interest rate $(= \xi_t)$. A consumer that becomes aware of this event now may nevertheless doubt that others will be aware of the same event in the near future and may therefore also doubt that aggregate spending will go up. As this logic applies for the average consumer, the economy as a whole systematically underestimates the future movements in aggregate income, and behaves like a representative agent that excessively discounts the future. And the larger the dependence of spending on income, or the steeper the Keynesian cross, the larger this discounting.

This explains the documented myopia. The anchoring, on the other hand, has to do with learning. As more times passes since the occurrence of any given shock, consumers become progressively more aware of it. But higher-order beliefs adjust more sluggishly than first-order beliefs—equivalently, the expectations of income adjust more sluggishly than expectations of interest rates. This reduces the speed of adjustment in aggregate spending, or equivalently it increases the apparent dependence of current spending on past spending. And the steeper the Keynesian cross, the larger this effect, too.

Versions of these insights have been documented in the literature before, albeit not in the sharp form offered here.[1] Relative to the state of the art, our theoretical contribution contains: the bypassing of the curse of dimensionality in higher-order beliefs; the existence, uniqueness and analytical characterization of the equilibrium; the aforementioned observational-equivalence result; and an extension to a class of incomplete-information networks. This in turn paves the way to our applied contribution, which we detail below.

DSGE, micro to macro, and bounded rationality

Our observational equivalence offers the sharpest to-date illustration of how informational frictions may substitute for the ad hoc forms of sluggish adjustment employed in the DSGE literature: the backward-looking element in condition (2) is akin to that introduced by habit persistence in consumption, adjustment costs to investment, or indexation of prices to past inflation.

Crucially, the documented distortions increase not only with the level of noise but also with parameters that regulate the strength of GE feedback loops and the associated importance of higher-order beliefs. In the context of the NKPC, examples of such parameters include the frequency of price adjustment, the degree of market concentration, and the input-output matrix; and in the context of the Dynamic IS curve, they include liquidity constraints and consumer heterogeneity.

---

[1]In particular, the role of learning as source of sluggish adjustment in behavior is the common theme of Sims (2003) and Mankiw and Reis (2002); the higher sluggishness of higher-order beliefs relative to first-order beliefs has been emphasized by Woodford (2003) and Morris and Shin (2006); and the role of higher-order beliefs as a source of as-if myopia has been highlighted by Angeletos and Lian (2018).

At the same time, our analysis also yields the following, seemingly paradoxical, conclusion: more responsiveness at the micro level often comes together with more sluggishness at the macro level. For instance, a smaller Calvo friction maps to more sluggishness in aggregate inflation, and a higher marginal propensity to consume (MPC) maps to more habit-like persistence in aggregate consumption. In both cases, the reason is that the larger micro-level responsiveness is tied to a larger bite of higher-order uncertainty.

At the same time, our result builds a bridge to a recent literature that emphasizes how lack of common knowledge (Angeletos and Lian, 2018) and related kinds of bounded rationality (Farhi and Werning, 2019; Gabaix, 2020; Garcıa-Schmidt and Woodford, 2019) make agents behave *as if* they are myopic. But whereas this prior literature has restricted the belief error triggered by any shock to be time-invariant, our analysis lets it decay with the age of the shock, thanks to the accommodation of learning. This explains why our approach yields not only $\omega_f < 1$ but also $\omega_b > 0$, which is exactly what the data want.

Connection to evidence on expectations

Our results facilitate a simple quantitative strategy. We show how estimates of ω_f and ω_b can be obtained by combining knowledge about GE parameters with an appropriate moment of the average forecasts. Such a moment is estimated in Coibion and Gorodnichenko (2015), or CG for short: it is the the coefficient of the regression of the average forecast errors on past forecast revisions.

The basic intuition is that a higher value for this moment indicates a larger informational friction. But both the structural interpretation of this moment and its mapping to the macroeconomic dynamics is modulated by the GE feedback. When this feedback is strong enough, a modest friction by the CG metric may camouflage a large friction in terms of the values for ω_f and ω_b.

At the same time, we explain why the evidence on the under-reaction of average forecasts provided in CG is more "reliable" for our purposes than the conflicting evidence on the over-reaction of individual forecasts provided in Bordalo et al. (2020) and Broer and Kohlhas (2019). In an extension that adds a behavioral element as in those papers (a form of overconfidence), we can vary the theory's implications about individual forecasts without varying the structural relation between average forecasts and aggregate outcomes. The values of ω_f and ω_b are thus pinned down solely by the CG moment.

Applications: NKPC, HANK, and Asset Pricing

Our first application (Section VI) concerns inflation. Using our toolbox, we show that the friction implicit in surveys of expectations is large enough to rationalize existing estimates of the Hybrid NKPC. This complements Nimark (2008), which articulated the basic idea but did not discipline the theory with expectations data. To the best of our knowledge, ours is indeed the first estimate of what the available evidence of expectations means for inflation dynamics.

Echoing a core theme of our paper, we show that most of the documented effect regards the expectations of the behavior of others (inflation) rather than the expectations of the fundamental (real marginal cost). We finally put forward three ideas, all of which stem from the endogeneity of the Hybrid NKPC under the prism of our analysis. The first two draw a possible causal link from the increase in market concentration and the conduct of monetary policy to the reduction in inflation persistence. The third highlights that the economy's production network may influence not only the slope of the Philips curve (as in Rubbo, 2020; La'O and Tahbaz-Salehi, 2020) but also its backward-looking element.

Our second application (Section VII) turns to aggregate demand. As already mentioned, our theory provides a micro-foundation of habit-like persistence in aggregate spending. For a plausible calibration, this persistence is quantitatively comparable to that assumed in the DSGE literature, but requires no actual habit at the micro level. This helps reconcile the gap between the levels of habit required to match the macroeconomic time series and the much smaller levels estimated in microeconomic data (Havranek, Rusnak and Sokolova, 2017).

Relatedly, because the as-if myopia and habit increase with the MPC, our results help reconcile the high responsiveness of consumer spending to income shocks at the micro level with the sluggishness of aggregate spending to interest-rate shocks at the macro level.[2] This hints at a link between our contribution and the emerging HANK literature. We take a step in this direction by studying a heterogeneous-agent extension of our setting and showing the following property in it: a positive cross-sectional correlation between MPC and income cyclicality, like that documented empirically in Patterson (2019), amplifies the expectations-driven sluggishness in the response of aggregate spending to monetary policy.

Other applications include investment (Online Appendix Section **??**) and asset pricing (Online Appendix Section **??**). In the latter context, our results illustrate how higher-order uncertainty may be the source of both momentum and excessive discounting. They also suggest that both distortions may be greater at the level of the entire stock market than at the level of the stock of a particular firm, which in turn may help rationalize Samuelson's dictum (Jung and Shiller, 2005).[3]

Networks

Our HANK application is an example of how our toolbox can be extended to a class of networks. In this context, we offer a tractable characterization of the equilibrium dynamics as functions of the network and information structures. This builds a bridge to a growing literature that emphasizes the network structure of the economy but often ignores informational frictions.[4]

---

[2] A similar point has been made recently by Auclert, Rognlie and Straub (2020).

[3] Choi, Rondina and Walker (2020) also attempt to rationalize the discrepancies between aggregate and individual asset prices based on incomplete information and segmented markets, but their work focuses on pricing efficiency and volatility instead of momentum and discounting.

[4] A few notable exemptions are Bergemann, Heumann and Morris (2017) and Golub and Morris (2019) on the abstract front, and Nimark, Chahrour and Pitschner (2019), Auclert, Rognlie and Straub (2020)

II. Framework

In this section we set up our framework and illustrate its applicability.

A. Basic ingredients

Time is discrete, indexed by $t \in \{0, 1, ...\}$, and there is a continuum of agents, indexed by $i \in [0, 1]$. At any t, each agent chooses an action $a_{i,t} \in \mathbb{R}$. Let a_t be the average action. Best responses admit the following recursive formulation:

$$(3) \qquad a_{i,t} = \mathbb{E}_{i,t} \left[\varphi \xi_t + \beta a_{i,t+1} + \gamma a_{t+1} \right],$$

where ξ_t is an underlying fundamental, $\mathbb{E}_{i,t}[\cdot]$ is the agent's expectation in period t, and (φ, β, γ) are parameters, with $\varphi > 0$, $\gamma \in [0, \delta)$, and $\beta \equiv \delta - \gamma$, for some $\delta \in (0, 1)$. As it will become clear, δ parameterizes the agent's overall concern about the future and γ the GE, or strategic, considerations.

Iterating on condition (3) yields the following representation of i's best response:

$$(4) \qquad a_{i,t} = \sum_{k=0}^{\infty} \beta^k \mathbb{E}_{i,t} \left[\varphi \xi_{t+k} \right] + \gamma \sum_{k=0}^{\infty} \beta^k \mathbb{E}_{i,t} \left[a_{t+k+1} \right].$$

While the recursive form (3) is more convenient for certain derivations, the extensive form given above is more precise because it embeds appropriate "boundary" conditions for $t \to \infty$.[5] It also makes salient how a agent's optimal behavior at any given point of time depends on her expectations of the *entire* future paths of the fundamental and of the average action.

Aggregating condition (4) yields the following equilibrium restriction:

$$(5) \qquad a_t = \varphi \sum_{k=0}^{\infty} \beta^k \overline{\mathbb{E}}_t \left[\xi_{t+k} \right] + \gamma \sum_{k=0}^{\infty} \beta^k \overline{\mathbb{E}}_t \left[a_{t+k+1} \right],$$

where $\overline{\mathbb{E}}_t[.]$ denotes the average expectation in the population. This condition highlights the fixed-point relation between the equilibrium value of a_t and the expectations of it. As it will become clear, this condition also allows us to nest a variety of applications.[6]

---

and La'O and Tahbaz-Salehi (2020) on the applied front. None of these papers, however, share either our analytical results or our emphasis on forward-looking behavior.

[5]Namely, we have imposed that, for any date τ and history, $\lim_{t \to \infty} \beta^t \mathbb{E}_{i,\tau}[a_{i,t}] = 0$, $\lim_{t \to \infty} \beta^t \mathbb{E}_{i,\tau}[\xi_t] = 0$, and $\lim_{t \to \infty} \beta^t \mathbb{E}_{i,\tau}[a_t] = 0$. The first property can be understood as the transversality condition. The second represents a restriction on the fundamental process, trivially satisfied when ξ_t is bounded. The third represents an equilibrium refinement.

[6]The same best-response structure is assumed in Angeletos and Lian (2018). But whereas that paper considers a non-stationary setting where ξ_t is fixed at zero for all $t \neq T$, for some given $T \geq 1$, we consider a stationary setting in which ξ_t varies in all t and, in addition, there is gradual learning over time. Our framework also reminds the *static* beauty contests studied in Morris and Shin (2002), Woodford (2003), Angeletos and Pavan (2007), and Huo and Pedroni (2020). There, agents try to predict the concurrent behavior of others. Here, they try to predict the future behavior of others.

B. *Complete information and beyond*

Suppose that information is complete, meaning that all agents share the *same* information and this fact itself is common knowledge. The economy then admits a representative agent. That is, $a_{i,t} = a_t$ and $\mathbb{E}_{i,t} = \mathbb{E}_t$, where $\mathbb{E}_t$ stands for the representative agent's expectation, and condition (3) reduces to

$$
(6) \qquad\qquad a_t = \mathbb{E}_t[\varphi \xi_t + \delta a_{t+1}].
$$

This may correspond to the textbook versions of the Dynamic IS and New Keynesian Philips curves, or an elementary asset-pricing equation. By the same token, the equilibrium outcome is given by

$$
(7) \qquad\qquad a_t = \varphi \sum_{h=0}^{\infty} \delta^h \mathbb{E}_t[\xi_{t+h}].
$$

This can be read as "inflation equals the present discounted value of real marginal costs" or "the asset's price equals the present discounted value of its dividends."

Clearly, only the composite parameter $\delta = \beta + \gamma$ enters the determination of the equilibrium outcome: its decomposition between β and γ is irrelevant. As made clear in Section III.A below, this underscores that the decomposition between PE and GE considerations is immaterial in this benchmark. Furthermore, the outcome is pinned down by the expectations of the fundamental alone.

These properties hold because this benchmark imposes that agents can reason about the behavior of *others* with the same ease and precision as they can reason about their *own* behavior. Conversely, introducing incomplete (differential) information and higher-order uncertainty, as we shall do momentarily, amounts to accommodating a friction in how agents reason about the behavior of others, or about GE.

C. *Two Examples: Dynamic IS and NKPC*

Before digging any further into the theory, we illustrate how our setting can nest the two building blocks of the New Keynesian model, the Dynamic IS curve and the New Keynesian Philips curve (NKPC). The familiar, log-linearized, representative-agent versions of these equations are given by, respectively,

$$
c_t = \mathbb{E}_t[-\varsigma r_t + c_{t+1}] \qquad \text{and} \qquad \pi_t = \mathbb{E}_t[\kappa \mathtt{mc}_t + \chi \pi_{t+1}],
$$

where c_t is aggregate consumption, r_t is the real interest rate, π_t is inflation, $\mathtt{mc}_t$ is the real marginal cost, $\varsigma > 0$ is the elasticity of intertemporal substitution, $\kappa \equiv \frac{(1-\chi\theta)(1-\theta)}{\theta}$ is the slope of the Philips curve, $\theta \in (0,1)$ is the Calvo parameter, $\chi \in (0,1)$ is the representative agent's discount factor, and $\mathbb{E}_t$ is her expectation. Clearly, both of these conditions are nested in condition (6).

Relaxing the common-knowledge foundations of the New Keynesian model along the lines of Angeletos and Lian (2018) yields the following incomplete-information extensions of these equations:

$$(8) \qquad c_t = -\varsigma \sum_{k=0}^{\infty} \chi^k \overline{\mathbb{E}}_t[r_{t+k}] + (1 - \chi) \sum_{k=1}^{\infty} \chi^{k-1} \overline{\mathbb{E}}_t[c_{t+k}],$$

$$(9) \qquad \pi_t = \kappa \sum_{k=0}^{\infty} (\chi\theta)^k \overline{\mathbb{E}}_t [\mathtt{mc}_{t+k}] + \chi(1 - \theta) \sum_{k=0}^{\infty} (\chi\theta)^k \overline{\mathbb{E}}_t [\pi_{t+k+1}],$$

where $\overline{\mathbb{E}}_t$ denotes the average expectation of the consumers in condition (8) and that of the firms in condition (9). The first equation is nested in condition (5) by letting $a_t = c_t$, $\xi_t = r_t$, $\varphi = -\varsigma$, $\beta = \chi$, $\gamma = 1 - \chi$, and $\delta = 1$; the second by letting $a_t = \pi_t$, $\xi_t = \mathtt{mc}_t$, $\varphi = \kappa$, $\beta = \chi\theta$, $\gamma = \chi(1 - \theta)$ and $\delta = \chi$.

To understand condition (8), recall that the Permanent Income Hypothesis gives consumption as a function of the present discounted value of income. Incorporating variation in the real interest rate and heterogeneity in beliefs, and using the fact that aggregate income equals aggregate spending in equilibrium, yields condition (8). Finally, note that $1 - \chi$ measures the marginal propensity to consume (MPC) out of income. The property that $\gamma = 1 - \chi$ therefore means that, in this context, γ captures the slope of the Keynesian cross, or the GE feedback between spending and income.

To understand condition (9), recall that a firm's optimal reset price is given by the present discounted value of its nominal marginal cost. Aggregating across firms and mapping the average reset price to inflation yields condition (9). When all firms share the same, rational expectations, this condition reduces to the familiar, textbook version of the NKPC. Away from that benchmark, condition (9) reveals the precise manner in which expectations of future inflation (the behavior of firms) feed into current inflation. Note in particular that $\gamma = \chi(1 - \theta)$, which means that the effective degree of strategic complementarity increases with the frequency of price adjustment. This is because the feedback from the expectations of future inflation to current inflation increases when a higher fraction of firms are able to adjust their prices today on the basis of such expectations.

III. The Equivalence Result

This section contains the core of our contribution. We motivate the requisite assumptions, solve for the rational-expectations fixed point, develop our observation-equivalence result, and discuss the main insights encapsulated in it.

A. *Higher-Order Beliefs: The Wanted Essence and the Unwanted Complexity*

Higher-order beliefs are synonymous to how agents reason about GE effects. To see this, revisit condition (5), which allows the following decomposition of the

aggregate outcome:

$$(10) \qquad a_t = \varphi \underbrace{\sum_{k=0}^{\infty} \beta^k \overline{\mathbb{E}}_t \left[\xi_{t+k} \right]}_{\text{PE component}} + \gamma \underbrace{\sum_{k=0}^{\infty} \beta^k \overline{\mathbb{E}}_t \left[a_{t+k+1} \right]}_{\text{GE component}}.$$

We label the first term as the PE component because it captures the agents' response to any innovation holding constant their expectations about the endogenous outcome; the additional change triggered by any adjustment in these expectations, or the second term above, represents the GE component.

Consider now two economies, labeled A and B, that share the same $\delta \equiv \beta + \gamma$ but have a different mixture of β and γ. Economy A features $\beta = \delta$ and $\gamma = 0$, which means that GE considerations are entirely absent. Economy B features $\beta = 0$ and $\gamma = \delta$, which corresponds to "maximal" GE considerations.

In economy A, condition (5) reduces to $a_t = \varphi \sum_{k=0}^{\infty} \delta^k \overline{\mathbb{E}}_t [\xi_{t+k}]$, that is, only the first-order beliefs of the fundamental matter. This is similar to the representative-agent benchmark, except that the representative agent's expectations are replaced by the average expectations in the population. In economy B, instead, condition (5) reduces to $a_t = \varphi \overline{\mathbb{E}}_t [\xi_t] + \delta \overline{\mathbb{E}}_t [a_{t+1}]$ and recursive iteration yields

$$(11) \qquad a_t = \varphi \sum_{h=1}^{\infty} \delta^h \overline{\mathbb{F}}_t^h \left[\xi_{t+h-1} \right],$$

where, for any variable X, $\overline{\mathbb{F}}_t^1 [X] \equiv \overline{\mathbb{E}}_t [X]$ denotes the average first-order forecast of X and, for all $h \geq 2$, $\overline{\mathbb{F}}_t^h [X] \equiv \overline{\mathbb{E}}_t \left[\overline{\mathbb{F}}_{t+1}^{h-1} [X] \right]$ denotes the corresponding h-th order forecast. The key difference from both the representative-agent benchmark and economy A is the emergence of such higher-order beliefs. These represent GE considerations, or the agents' reasoning about the behavior of others.

The logic extends to the general case, in which both β and γ are positive. The only twist is that the relevant set of higher-order beliefs is significantly richer than that seen in condition (11). Indeed, let $\zeta_t \equiv \sum_{\tau=0}^{\infty} \beta^\tau \xi_{t+\tau}$ and consider the following set of forward-looking, higher-order beliefs:

$$\overline{\mathbb{E}}_{t_1} [\overline{\mathbb{E}}_{t_2} [\cdots [\overline{\mathbb{E}}_{t_h} [\zeta_{t+k}] \cdots]],$$

for any $t \geq 0$, $k \geq 2$, $h \in \{2, ..., k\}$, and $\{t_1, t_2, ..., t_h\}$ such that $t = t_1 < t_2 < ... < t_h = t + k$. As behavior depends on all these higher-order beliefs, this adds considerable complexity relative to the $\beta = 0$ case. For instance, when $k = 10$ (thinking about the outcome 10 periods later), there are 210 beliefs of the fourth order that are relevant when $\beta > 0$ compared to only one such belief when $\beta = 0$.[7]

---

[7]More generally, for any t and any $k \geq 2$, there are now $k - 1$ types of second-order beliefs, plus $(k - 1) \times (k - 2)/2$ types of third-order beliefs, and so on.

An integral part of our contribution is the bypassing of this complexity. The assumptions that permit this bypassing are spelled out below. They come at the cost of some generality, in particular we abstract from the possible endogeneity of information.[8] But they also bear significant gains on both the theoretical and the quantitative front, which will become evident as we proceed.

B. Specification

We henceforth make two assumptions. First, we let the fundamental ξ_t follow an AR(1) process:

$$(12) \qquad \xi_t = \rho \xi_{t-1} + \eta_t = \frac{1}{1 - \rho L} \eta_t,$$

where $\eta_t \sim \mathcal{N}(0, 1)$ is the period-t innovation, L is the lag operator, and $\rho \in (0, 1)$ parameterizes the persistence of the fundamental. Second, we assume that agent i receives a new private signal in each period t, given by

$$(13) \qquad x_{i,t} = \xi_t + u_{i,t}, \quad u_{i,t} \sim \mathcal{N}(0, \sigma^2),$$

where $\sigma \geq 0$ parameterizes the informational friction (the level of noise). The agent's information in period t is the history of signals up to that period.

As anticipated in the previous subsection, these assumptions aim at minimizing complexity without sacrificing essence. Borrowing from the literature on rational inattention, we also invite a flexible interpretation of our setting as one where fundamentals and outcomes are observable but cognitive limitations makes agents act *as if* they observe the entire state of nature with idiosyncratic noise. But instead of endogenizing the noise, we fix it in a way that best serves our purposes.

C. Solving the Rational-Expectations Fixed Point

Consider momentarily the frictionless benchmark ($\sigma = 0$), in which case the outcome is pinned down by first-order beliefs, as in condition (7). Thanks to the AR(1) specification for the fundamental, we have $\mathbb{E}_t[\xi_{t+k}] = \rho^k \xi_t$, for all $t, k \geq 0$. We thus reach the following result, which states that the complete-information outcome follows the same, up to a rescaling, AR(1) process as the fundamental.[9]

PROPOSITION 1: *In the frictionless benchmark ($\sigma = 0$), the equilibrium outcome is given by*

$$(14) \qquad a_t = a_t^* \equiv \frac{\varphi}{1 - \rho\delta} \xi_t = \frac{\varphi}{1 - \rho\delta} \frac{1}{1 - \rho L} \eta_t.$$

---

[8]This abstraction is the right benchmark for our purposes, including the connections built to the evidence on expectations: this evidence helps discipline the theoretical mechanisms we are concerned with, but contains little guidance on the degree or manner in which information may be endogenous.

[9]All proofs are delegated to Online Appendix Section **??**.

Consider next the case in which information is incomplete ($\sigma > 0$). As already explained, the outcome is then a function of an infinite number of higher-order beliefs. Despite the assumptions made here about the process of ξ_t and the signals, these beliefs remain exceedingly complex.

Let us illustrate this point. Using the Kalman filter, one can readily show that the first-order belief $\overline{\mathbb{E}}_t[\xi_t]$ obeys the following AR(2) dynamics:

$$(15) \qquad \overline{\mathbb{E}}_t[\xi_t] = \left(1 - \frac{\lambda}{\rho}\right)\left(\frac{1}{1 - \lambda L}\right)\xi_t,$$

where $\lambda = \rho(1 - g)$ and $g \in (0, 1)$ is the Kalman gain, itself a decreasing function of the level of noise.[10] It follows that the second-order belief $\overline{\mathbb{E}}_t[\overline{\mathbb{E}}_{t+1}[\xi_{t+1}]]$ follows an ARMA(3,1). By induction, for any $h \geq 1$, the h-th order belief $\overline{\mathbb{E}}_t[\overline{\mathbb{E}}_{t+1}[...\overline{\mathbb{E}}_{t+h}[\xi_{t+h}]]$ follows an ARMA($h + 1, h - 1$). Beliefs of higher order thus exhibit increasingly complex dynamics.

As explained in Section III.A, the above set of higher-order beliefs is the relevant one when $\beta = 0$. The general case with $\beta > 0$ is subject to an even greater curse of dimensionality in terms of higher-order beliefs. And yet, this complexity vanishes once we focus on the rational-expectations fixed point: under our assumptions, the fixed point turns out to be merely an AR(2) process, whose exact form is characterized below.

PROPOSITION 2 (Solution):   *The equilibrium exists, is unique and is such that the aggregate outcome obeys the following law of motion:*

$$(16) \qquad a_t = \left(1 - \frac{\vartheta}{\rho}\right)\left(\frac{1}{1 - \vartheta L}\right)a_t^*,$$

where a_t^* is the frictionless counterpart, obtained in Proposition 1, and where ϑ is a scalar that satisfies $\vartheta \in (0, \rho)$ and that is given by the reciprocal of the largest root of the following cubic:

$$(17) \qquad C(z) \equiv -z^3 + \left(\rho + \frac{1}{\rho} + \frac{1}{\rho\sigma^2} + (\delta - \gamma)\right)z^2$$
$$- \left(1 + (\delta - \gamma)\left(\rho + \frac{1}{\rho}\right) + \frac{\delta}{\rho\sigma^2}\right)z + (\delta - \gamma).$$

Condition (16) gives the incomplete-information dynamics as a transformation of the complete-information counterpart. This transformation is indexed by ϑ. Relative to the frictionless benchmark (herein nested by $\vartheta = 0$), a higher ϑ means both a smaller impact effect, captured by the factor $1 - \frac{\vartheta}{\rho}$ in condition (16), and a more sluggish build up over time, captured by the lag term ϑL.

---

[10]The Kalman gain is given by the unique $g \in (0, 1)$ such as that $(1 - g) = (1 - \rho^2(1 - g))g\sigma^2$. This yields g as a continuous and decreasing function of σ, with $g = 1$ when $\sigma = 0$ and $g \to 0$ when $\sigma \to \infty$.

To understand the math behind the result, let $\beta = 0$ momentarily. In this case, the outcome obeys

$$(18) \qquad a_t = \varphi\overline{\mathbb{E}}_t[\xi_t] + \gamma\overline{\mathbb{E}}_t[a_{t+1}].$$

If we guess that a_t follows an AR(2), we have that $\overline{\mathbb{E}}_t[a_{t+1}]$ follows an ARMA(3,1). As already noted, $\overline{\mathbb{E}}_t[\xi_t]$ follows the AR(2) given in (15). The right-hand side of the above equation is therefore the sum of an AR(2) and an ARMA(3,1). If the latter was arbitrary, this sum would have returned an ARMA(5,3), contradicting our guess that a_t follows an AR(2). But the relevant ARMA(3,1) is *not* arbitrary.

Because the impulse behind a_t is ξ_t, one can safely guess that a_t inherits the root of ξ_t. That is, $(1 - \vartheta L)(1 - \rho L)a_t = b\eta_t$, for some scalars b and ϑ. This in turn implies that the AR roots of the ARMA(3,1) process for $\overline{\mathbb{E}}_t[a_{t+1}]$ are the reciprocals of ρ, ϑ and λ. As seen in (15), the roots of $\overline{\mathbb{E}}_t[\xi_t]$ are the reciprocals of ρ and λ. It follows that the sum in the right-hand side of (18) is at most an ARMA(3,1) of the following form:

$$(19) \qquad a_t = \frac{c(1 - dL)}{(1 - \vartheta L)(1 - \rho L)(1 - \lambda L)}\eta_t,$$

where c and d are functions of b and ϑ. For our guess to be correct, it has to be that $d = \lambda$ and $c = b$. The first equation, which lets the MA part and the last AR part cancel out so as to reduce the above to an AR(2), and yields (17). The second equation, which pins down the scale, yields $b = \left(1 - \frac{\vartheta}{\rho}\right)\left(\frac{\varphi}{1 - \rho\delta}\right)$.

This is the crux of how the rational expectations fixed point works. The proof presented in Online Appendix Section **??** follows a somewhat different path, which is more constructive, accommodates $\beta > 0$, and can be extended to richer settings along the lines of Huo and Takayama (2018).

When $\gamma = 0$, GE considerations are absent, the outcome is pinned down by first-order beliefs, and Proposition 2 holds with $\vartheta = \lambda$, where λ is the same root as that seen in (15). When instead $\gamma > 0$, GE considerations and higher-order beliefs come into play. As already noted, these beliefs follow complicated ARMA processes of ever increasing orders. And yet, the equilibrium continues to follow an AR(2) process. The only twist is that $\vartheta > \lambda$, which, as mentioned above, means that the equilibrium outcome exhibits less amplitude and more persistence than the first-order beliefs. This is the empirical footprint of higher-order uncertainty, or of the kind of imperfect GE reasoning accommodated in our analysis.

Below, we translate these properties in terms of our observational-equivalence result (Propositions 3 and 5). The following corollary, which proves useful when connecting the theory to evidence on expectations, is also immediate.

COROLLARY 1 (Forecasts): *Any moment of the joint process of the aggregate outcome, a_t, and of the average forecasts, $\overline{\mathbb{E}}_t[a_{t+k}]$ for all $k \geq 1$, are functions of only the triplet $(\vartheta, \lambda, \rho)$, or equivalently of $(\gamma, \delta, \rho, \sigma)$.*

D.   *The Equivalence Result*

Momentarily put aside our incomplete-information economy and, instead, consider a "behavioral" economy populated by a representative agent whose aggregate Euler condition (6) is as follows:

$$(20) \qquad a_t = \varphi \xi_t + \delta \omega_f \mathbb{E}_t [a_{t+1}] + \omega_b a_{t-1},$$

for some scalars ω_f, ω_b. It is easy to verify that the equilibrium process of a_t in this economy is an AR(2) whose coefficients are functions of (ω_f, ω_b) and (φ, δ, ρ). In comparison, the equilibrium process of a_t in our incomplete-information economy is an AR(2) whose coefficients determined as in Proposition 2. Matching the coefficients of the two AR(2) processes, and characterizing the mapping from the latter to the former, we reach the following result.

PROPOSITION 3 (Observational Equivalence):   *Fix $(\varphi, \delta, \gamma, \rho)$. For any noise level $\sigma > 0$ in the incomplete-information economy, there exists a unique pair (ω_f, ω_b) in the behavioral economy such that the two economies generate the same joint dynamics for the fundamental and the aggregate outcome. Furthermore, this pair satisfies $\omega_f < 1$ and $\omega_b > 0$.*

This result allows one to recast the informational friction as the combination of two behavioral distortions: extra discounting of the future, or myopia, in the form of $\omega_f < 1$; and backward-looking behavior, or anchoring of the current outcome to past outcome, in the form of $\omega_b > 0$.

This representation is, of course, equivalent to the closed-form solution provided in Proposition 2. We prefer the new representation not only because it serves the applied purposes of our paper, but also because the main insights about myopia and anchoring extend to richer settings, while the specific AR(2) solution provided in Proposition 2 does not. This idea is formalized in Online Appendix Section **??**.

E.   *The Roles of Noise and GE Considerations*

As one would expect, both distortions increase with the level of noise.

PROPOSITION 4 (Noise):   *A higher σ maps to a lower ω_f and a higher ω_b.*

What this result, however, fails to highlight is the dual meaning of "noise" in our setting: a higher σ represents not only less accurate information about the fundamental (larger first-order uncertainty) but also more friction in how agents reason about others (larger higher-order uncertainty). The latter, strategic or GE, channel is highlighted by the next result.

PROPOSITION 5 (GE):   *Consider an increase in the relative importance of GE considerations, as captured by an increase in γ holding $\delta \equiv \beta + \gamma$, as well as σ and ρ, constant. This maps to both greater myopia (lower ω_f) and greater anchoring (higher ω_b).*

This result circles backs to our discussion in Section III.A regarding the interpretation of higher-order uncertainty as a distortion in agents' GE reasoning. It also anticipates a point we make in Section V. While the kind of evidence on informational frictions provided by Coibion and Gorodnichenko (2015) is an essential ingredient for the quantitative evaluation of the assumed friction, it is not sufficient. One must combine such evidence with knowledge of how important the GE feedback from expectations to actual behavior is.

F. Robustness

The results presented above depend on stark assumptions about the process of ξ_t and the information structure. But the key insights regarding myopia, anchoring, and the role of higher-order beliefs are more general. Online Appendix Section **??** shows how to generalize these insights in a setting that allows ξ_t to follow an essentially arbitrary MA process, as well as information to diffuse in a flexible manner.[11] The elegance of our observational-equivalence result is lost, but the essence remains.

Another extension, better suited for applied purposes, is offered in Section VIII. There, we consider a multi-variate analogue of condition (4). This allows one to handle the full, three-equation New Keynesian model, the HANK variant considered in Section VII, and a large class of linear networks.

IV. Connection to DSGE, Bounded Rationality, and beyond

In the end of Section II we sketched how our framework nests incomplete-information extensions of the Dynamic IS curve and the NKPC. We also discussed how γ relates to the slope of the Keynesian cross, or the income-spending multiplier, in the first context and to the frequency of price adjustment in the second. The following translations of our abstract results are thus immediate.

COROLLARY 2: *Applying our result to condition (9) yields the following NKPC:*

$$\pi_t = \kappa mc_t + \omega_f \chi \mathbb{E}_t[\pi_{t+1}] + \omega_b \pi_{t-1}. \tag{21}$$

In this context, the distortions increase with the frequency of price adjustment.

COROLLARY 3: *Applying our result to condition (8) yields the following Dynamic IS curve:*

$$c_t = -\varsigma r_t + \omega_f \mathbb{E}_t[c_{t+1}] + \omega_b c_{t-1}. \tag{22}$$

In this context, the distortions increase with the MPC, or the slope of the Keynesian cross.

---

[11]Such richness is prohibitive in general. We cut the Gordian knot by orthogonalizing the information about the innovations occurring at different points of time. Although this modeling approach is unusual, it nests "sticky information" (Mankiw and Reis, 2002) as a special case and clarifies the theoretical mechanisms.

Condition (21) looks like the Hybrid NKPC. Condition (22) looks like the Euler condition of representative consumer who exhibits habit persistence plus myopia. Online Appendix Section **??** offers a related result for investment: we take a model in which adjustments cost depend on the change in the stock of capital, as in traditional Q theory; add incomplete information; and show that this model looks like a model in which adjustment costs depend on the change in the rate of investment.

Together, these results illustrate how informational frictions can substitute for the more ad hoc sources of sluggishness in *all* the equations of DSGE models. The basic idea is familiar from previous works (e.g., Sims, 2003; Mankiw and Reis, 2002; Woodford, 2003; Nimark, 2008). The added value here is the sharpness of the provided representation and the following, complementary lessons.

First, we build a bridge to a recent literature that shows how lack of common knowledge and related forms of bounded rationality make agents behave *as if* they are myopic. These works help generate $\omega_f < 1$ but restrict $\omega_b = 0$. In Angeletos and Lian (2018), this is because there is no learning. In Farhi and Werning (2019), Garcıa-Schmidt and Woodford (2019) and Iovino and Sergeyev (2017), it is a direct implication of the adopted solution concept: level-k thinking amounts to equating beliefs of order $h \leq k$ to their complete-information counterparts, and beliefs of order $h > k$ to zero. This makes agents underestimate GE effects, which maps to $\omega_f < 1$, but precludes the mistake in beliefs to be corrected over time, which maps to $\omega_b = 0$. Our approach, instead, naturally delivers both $\omega_f < 1$ and $\omega_b > 0$, which is what the macroeconomic data want.[12] By the same token, our approach allow both for under-reaction and momentum in average expectations, which is what the available survey evidence want.

Second, we offer a new rationale for why the information-driven sluggishness may loom large at the macro level even if is absent at the micro level. Previous work has emphasized that agents may naturally have less information about aggregate shocks than about idiosyncratic shocks (Maćkowiak and Wiederholt, 2009). We add that higher-order uncertainty effectively amplifies the friction at the macro level. We further clarify these points in Online Appendix Section **??** by considering an extension of our framework with idiosyncratic shocks. And in Online Appendix Section **??**, we discuss how the exact same logic transported to an asset-pricing context may help rationalize larger momentum at the macro level than at the micro level, or what is known as Samuelson's dictum (Jung and Shiller, 2005).

Third, by tying the macro-level distortions to strategic complementarity and GE feedbacks, we highlight how the former can be endogenous to market structures and policies that regulate the latter. We come back to this point in Section VI.

Fourth, in the context of the NKPC, we show that higher price flexibility contributes to more sluggishness in inflation by intensifying the role of higher-order

---

[12]This point applies to dynamic settings. In static games such as Morris and Shin (2002), the three approaches are observationally equivalent vis-a-vis the macroeconomic time series.

beliefs. This seems an intriguing, new addition to the "paradoxes of flexibility." And in the context of the Dynamic IS curve, we tie the habit-like persistence in consumption to the MPC, or the slope of the Keynesian cross. This hints at the promise of incorporating incomplete information in the HANK literature, an idea we expand on in Section VII.

Finally, we offer a simple strategy for quantifying the distortions of interest. We spell out the elements of this strategy in the next section and put it at work in our subsequent applications to inflation and consumption dynamics.

V. Connection to Evidence on Expectations

Proposition 3 ties the documented distortions to σ. This parameter may not be a priori known to the analyst ("econometrician"). Surveys of expectations, however, can help identify it. In this section, we use our results to map readily available evidence on expectations to the macroeconomic distortions of interest. We also clarify which subset of such evidence is best suited for our purposes (moments of average forecasts) and provide two examples of robustness for the offered mapping (one regarding overconfidence and another regarding public signals).

A. Calibrating the friction

Consider Coibion and Gorodnichenko (2015), or CG for short. This paper runs the following regression on data from the Survey of Professional Forecasters:

$$(23) \qquad a_{t+k} - \overline{\mathbb{E}}_t[a_{t+k}] = K_{\mathrm{CG}} \left(\overline{\mathbb{E}}_t[a_{t+k}] - \overline{\mathbb{E}}_{t-1}[a_{t+k}] \right) + v_{t+k,t},$$

where a_t is an economic outcome such as inflation and $\overline{\mathbb{E}}_t[a_{t+k}]$ is the average ("consensus") forecast of the value of this outcome k periods later. CG's main finding is that K_{CG}, the coefficient of the above regression, is positive. That is, a positive revision in the average forecast between $t-1$ and t predicts a positive average forecast error at t.

What does this mean under the lenses of the theory? Insofar as agents are rational, an agent's forecast error ought to be orthogonal to his *own* past revision, itself an element of the agent's information set. But this does not have to be true at the aggregate level, because the past average revision may not be commonly known. More succinctly, $K_{\mathrm{CG}} \neq 0$ is possible because the forecast error of one agent can be predictable by the past information of another agent.

Furthermore, because forecasts adjust sluggishly towards the truth, the theory suggests that K_{CG} ought to be positive and increasing in the informational friction. To illustrate this, CG treat a_t as an exogenous AR(1) process, assume the same Gaussian signals as we do, and show that in this case $K_{\mathrm{CG}} = \frac{1-g}{g}$, where $g \in (0,1)$ is the Kalman gain, itself a decreasing function of σ. They therefore argue that their estimate of K_{CG} offers a measure of the informational friction.

In our context, a_t is endogenous to expectations. This complicates the structural interpretation and use of this measure. The level of noise now influences

not only the agents' forecasting of a_t, but also its own stochastic process. Furthermore, because the level of noise interacts with the GE feedback in shaping the process for a_t, the GE parameter γ enters the mapping between σ and K_{CG}. The next result shows what exactly is going on.

PROPOSITION 6 (K_{CG}):   *The theoretical counterpart of the coefficient of regression (23) for $k = 1$ is given by*

$$(24) \qquad K_{CG} = \lambda \frac{\vartheta + \rho - \rho\vartheta(\lambda + \vartheta) - \rho\lambda\vartheta(1 - \lambda\vartheta)}{(\rho - \lambda)(1 - \lambda\vartheta)(\rho + \vartheta - \lambda\rho\vartheta)},$$

where λ and ϑ are defined as in Section III.C. It follows that
   (i) K_{CG} is increasing in σ, the level of noise; and
   (ii) K_{CG} is decreasing in γ, the GE feedback.

The formula for K_{CG} is not particularly intuitive. However, in combination with our closed-form characterizations for λ and ϑ, it allows us to prove the two illuminating comparative statics stated above. The first verifies that CG's logic that a high value for K_{CG} signals a high informational friction extends from their PE context, where a_t follows an exogenous process, to our GE context, where the process for a_t is influenced by the informational friction. The second comparative static highlights the limits of this logic: a small value for K_{CG} could conceal a large value for σ if the GE feedback is large enough.

At first glance, this may appear to contradict our result in Proposition 5 that a higher γ translates to larger distortions in the equilibrium dynamics. But the underlying logic for both results is actually the same. When γ is higher, agents are more willing to coordinate their behavior. This reduces the reliance of behavior on private information and increases the reliance on the prior or higher-order beliefs. As this happens, the equilibrium outcome becomes less responsive to innovations. But precisely because of this reason, the reliance of the forecasts of the outcome on private information is also reduced, which means that the forecast error of one agent is less predictable by the information of another agent, and hence that the K_{CG} coefficient is closer to zero.

What does this mean for the structural interpretation and use of the available expectations evidence? When the GE effect increases, both of the aforementioned channels work in the same direction: for given σ, a higher γ means both larger distortions in terms of (ω_f, ω_b) and a smaller observable footprint in terms of K_{CG}. The following is therefore true:

COROLLARY 4:   *As γ increases, the same value for K_{CG} maps to both more myopia (smaller ω_f) and more anchoring (larger ω_b) in the aggregate outcome.*

This is illustrated in Figure 1. On the horizontal axis, we vary the value of K_{CG} that may be recovered from running regression (23) on the applicable expectations data. On the vertical axis, we report the predicted values for ω_f and ω_b. For given

γ, a higher K_{CG} maps to a higher σ and thereby to larger distortions. But a higher γ maps to larger distortions for given K_{CG} not only because it amplifies the effect of noise, but also because it requires a larger σ to match the given K_{CG}.
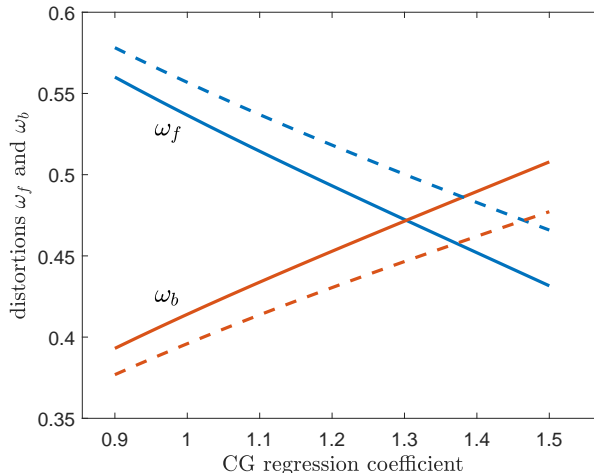

FIGURE 1. MYOPIA AND ANCHORING

Note: The distortions as functions of the proxy offered in Coibion and Gorodnichenko (2015). The solid lines correspond to a stronger degree of strategic complementarity, or GE feedback, than the dashed one.

B. *Individual forecasts and overconfidence*

So far, we have emphasized how one could make use of the moment estimated in CG, along with our tools, to obtain an estimate of ω_f and ω_b. Other moments of the *average* forecasts, such as the persistence of the average forecast errors estimated in Coibion and Gorodnichenko (2012), could serve a similar role. But what about moments of the *individual* forecasts? We next explain why such moments can be largely ignored for our purposes (but not for other purposes).

Consider, in particular, the individual-level counterpart of the CG regression, that is, the regression of one's forecast errors on one's *own* past revisions. As noted earlier, rational expectations requires that the coefficient of this regression be zero. Bordalo et al. (2020) and Broer and Kohlhas (2019) argue that this coefficient is negative in the data, supporting the presence of overconfidence. Our own take is that the evidence is inconclusive: the relevant coefficient switches signs across variables and samples (inflation vs. unemployment, pre- vs post-Volker, etc), making it hard to reject rational expectations. But even if we take for granted those papers' preposition of systematic bias in beliefs, this does not necessarily upset either our theoretical results or our proposed quantitative strategy.

We illustrate this point by augmenting our model with the same kind of over-confidence as Broer and Kohlhas (2019): whereas the actual level of noise is σ,

agents perceive it to be $\hat{\sigma}$, for some $\hat{\sigma} < \sigma$. (For completeness, under-confidence, or $\hat{\sigma} > \sigma$, is also allowed.) In this extension, the gap between $\hat{\sigma}$ and σ, or the degree of overconfidence, emerges as the essential determinant of the aforementioned individual-level moment.[13] But this moment and its determinant "drop out of the picture" for our purposes:

PROPOSITION 7: *Propositions 2–6 and Corollary 1 remain valid, modulo the replacement of σ with $\hat{\sigma}$ throughout. By implication, the mapping from K_{CG} to (ω_f, ω_b) is invariant to the degree of overconfidence.*

To understand this result, note that the perceived $\hat{\sigma}$ alone determines how much each agent's beliefs and choices vary with his information, and thereby how much the corresponding aggregates vary with the underlying fundamental. The true σ instead determines how *unequal* beliefs and choices are in the cross section, but such inequality does not matter for aggregates in our class of economies. It follows that all our results, including the characterization of (ω_f, ω_b) and K_{CG}, carry over by replacing σ with $\hat{\sigma}$.

Suppose, now, that the analyst knows all parameters except $\hat{\sigma}$ and σ and wishes to quantify the equilibrium effects of the friction under consideration (as we do, for example, in Section VI). Suppose further that the analyst combines the CG coefficient with the individual-level counterpart estimated in Bordalo et al. (2020) and Broer and Kohlhas (2019). Then, the CG coefficient *alone* allows the identification of $\hat{\sigma}$ and the quantification of its effect on the actual dynamics. The individual-level counterpart allows the identification of σ, but this does not affect the aforementioned quantitative evaluation.

A similar point applies to the cross-sectional dispersion of forecasts. A large part of it is accounted by individual-specific fixed effects, which themselves correlate with life-time experiences unrelated to the current macroeconomic context (Malmendier and Nagel, 2016). This can be accommodated in the theory by letting each agent i have a different prior mean, μ_i, about ξ_t. Such prior-mean heterogeneity is then a key determinant of the cross-sectional dispersion of forecasts. But it does not matter at all for our observational equivalence result and the offered mapping from K_{CG} to (ω_f, ω_b).

This also anticipates the exercise conducted in Table 1: for our quantitative application to inflation, we test the ability of our model to capture the cross-sectional dispersion of the forecast *errors* or the forecast *revisions* precisely because these objects partial out individual fixed effects such as those associated with the aforementioned kind of heterogeneity.

---

[13] Broer and Kohlhas (2019) establish this point in a setting where a_t follows an exogenous AR(1) process, but the logic extends to our context. When agents are overconfident ($\hat{\sigma} < \sigma$), they over-react to their information relative to what a rational agent would do, so a positive forecast revision today predicts a negative forecast error in the future. And the converse is true if agents are under-confident ($\hat{\sigma} > \sigma$). Also note that, although the formulation used in Bordalo et al. (2020) has different methodological underpinnings, it works in essentially the same way as the form overconfidence considered here.

More challenging is the evidence presented in Kohlhas and Walther (2019). In direct contradiction to CG's message, these authors argue that expectations *over-react* in the sense that average forecasts errors are negatively correlated with past outcomes. They then proceed to offer a resolution based on asymmetric attention to pro-cyclical and counter-cyclical components of the forecasted variable. In Online Appendix Section **??**, we explain how our methods can be adapted to their setting. And in Angeletos, Huo and Sastry (2021), we propose an alternative resolution, one based on the combination of informational frictions and over-extrapolation. But we leave this issue out of the present paper.

C. *Public information*

So far we have have let agents observe only private signals. If we add public signals, the CG moment is no more sufficient for uniquely identifying the information structure: there are multiple combinations of the precisions of the private and public signals that generate the same value for K_{CG}. By the same token, any given value for K_{CG} maps to a set of possible values for the pair (ω_f, ω_b).

At first glance, this poses a challenge for the quantitative strategy proposed in this section. However, as explained in Online Appendix Section **??** and illustrated in our application to inflation below, this challenge is resolved by two key observations.

First, K_{CG} puts a tight upper bound on the relative precision of the public signal. Intuitively, as information gets more and more correlated, everybody's expectations converge to those of a representative agent, and K_{CG} converges to zero. A high value for K_{CG} therefore means *necessarily* either that there is little public information to start with, or that people pay little attention to it.

Second, by varying the precision of public information between zero and the aforementioned upper bound, we can span the entire range of values (ω_f, ω_b) that are consistent with any given value of K_{CG}. In Online Appendix Section **??**, we implement this strategy in our application to inflation, which is the topic of the next section, and show that the distortions reported therein under the simplifying assumption of no public information represent a lower bound on the distortions obtained when public information is added.

VI. Application to Inflation

We now apply our toolbox the context of inflation. We argue that the theory can not only rationalize existing estimates of the Hybrid NKPC with some level of noise, but also do so with a level of noise consistent with that inferred from CG's evidence on expectations. We also illustrate how our theory ties the coefficients of the Hybrid NKPC to policy and market structures.[14]

[14]Nimark (2008) foresaw the first part of our application by showing that an econometrician would estimate a Hybrid NKPC on artificial data generated by his model. Relative to that paper, we offer a sharper illustration of this possibility and, most importantly, let the evidence on expectations bear on

A.   Operationalizing the theory

Consider the incomplete-information NKPC introduced in Section II:[15]

$$(25) \qquad \pi_t = \kappa \sum_{k=0}^{\infty} (\chi\theta)^k \overline{\mathbb{E}}_t \left[\mathtt{mc}_{t+k}\right] + \chi(1-\theta) \sum_{k=0}^{\infty} (\chi\theta)^k \overline{\mathbb{E}}_t \left[\pi_{t+k+1}\right],$$

Unlike the representation obtained in Corollary 2, this equation is structural: it is invariant to the process for the real marginal cost and the specification of information. But it is also hard to implement empirically, because it requires data on the term structure of the relevant forecasts over long horizons. This is where our toolbox comes handy: using our results, we can connect the above structural equation both to existing estimates of the Hybrid NKPC and to the available evidence on expectations.

To evaluate these connections, we henceforth interpret the time period as a quarter and impose the following parameterization: $\chi = 0.99$, $\theta = 0.6$, and $\rho = 0.95$. The value of χ requires no discussion. The value of θ is in line with micro data and textbook treatments of the NKPC. The value of ρ is obtained by estimating an AR(1) process on the labor share, the empirical proxy for the real marginal cost used in, inter alia, Galí and Gertler (1999) and Galí, Gertler and Lopez-Salido (2005).[16] Finally, the value of κ is left undetermined: because this parameter scales up and down the inflation dynamics equally under any information structure, it is irrelevant for the conclusions drawn below.[17]

B.   Connecting to Existing Estimates of the Hybrid NKPC

While an unrestricted estimation of the Hybrid NKPC allows ω_f and ω_b to be free, our theory ties them together: a higher ω_b can be obtained only if the noise is larger, which in turns requires ω_f to be smaller. A quick test of the theory is therefore whether the existing estimates of the Hybrid NKPC happen to satisfy

---

the theory. Such a connection to the expectations evidence is also absent from Woodford (2003), Mankiw and Reis (2002), Reis (2006), Kiley (2007), Maćkowiak and Wiederholt (2009, 2015) and Matejka (2016). Melosi (2016) utilizes expectations data but studies a different issue, the signaling role of monetary policy. Finally, the literature on adaptive learning (Sargent, 1993; Evans and Honkapohja, 2012) also allows for the anchoring of current outcomes to past outcomes; see in particular Carvalho et al. (2017) for an application to inflation. But the anchoring found in our paper has three distinct qualities: it is consistent with rational expectations; it is tied to the strength of the GE feedback; and it is directly comparable to that found in the DSGE literature.

[15]Recall that π_t is the inflation rate, $\mathtt{mc}_t$ is the real marginal cost, $\chi \in (0,1)$ is the discount factor, $\theta \in (0,1)$ is the Calvo parameter, and $\kappa > 0$ is the slope of the NKPC. Online Appendix Section **??** contains a detailed derivation, a discussion of the underlying assumptions, and an explanation of a mistake in versions of this condition found in some prior work.

[16]We use seasonally adjusted business sector labor share as proxy for the real marginal cost, from 1947Q1 to 2019Q2. This yields an estimate of ρ equal to 0.97 or 0.92 depending on whether we exclude or include a linear trend.

[17]In the textbook version of the NKPC, κ is itself pinned down by χ and θ. But the literature has provided multiple rationales for why κ can differ from its textbook value (e.g., it can vary with the curvature of "Kimball aggregator"). For our purposes, this amounts to treating κ as a free parameter.

this restriction. We implement this test in Figure 2. The negatively slopped line depicts the aforementioned restriction. The crosses represent the three main estimates of the pair (ω_f, ω_b) from Galí, Gertler and Lopez-Salido (2005), and the surrounding disks give the corresponding confidence regions.[18]
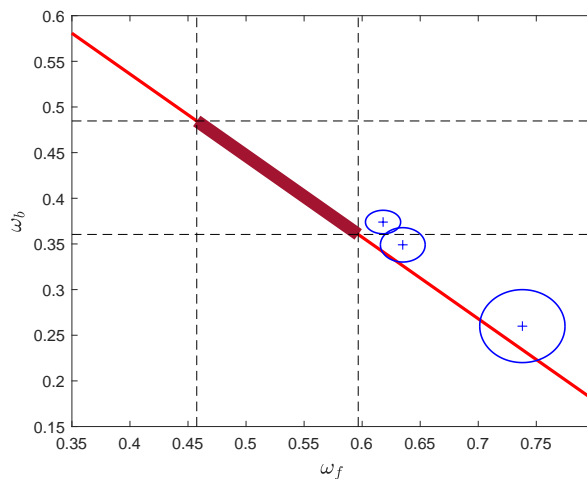

FIGURE 2. TESTING THE THEORY

Note: The straight line represents the relation between ω_f and ω_b implied by the theory. Raising the level of noise maps to moving northwest along this line. The darker, thicker segment of this line corresponds to the confidence interval of K_{CG}, the relevant moment of the inflation forecasts, as reported in column (1) of Table 1 Coibion and Gorodnichenko (2015). The three crosses represent the three estimates of the pair (ω_f, ω_b) provided in Table 1 of Galí, Gertler and Lopez-Salido (2005), and the surrounding disks give the corresponding confidence regions.

As evident in the figure, the theory passes the aforementioned test: the existing estimates of the Hybrid NKPC can be rationalized by some level of noise.[19] But is the requisite level of noise empirically plausible? We address this question next by making use of the mapping developed in Section V.

---

[18]The three estimates are taken from Table 1 of that paper. In particular, the left one of the three points shown in Figure 2 corresponds to $(\omega_f, \omega_b) = (0.618, 0.374)$ and is obtained by the GMM estimation of the closed-form solution that expresses current inflation as the sum of past inflation and all the expected future real marginal costs. The middle point corresponds to $(\omega_f, \omega_b) = (0.635, 0.349)$ and is obtained by GMM estimation of the hybrid NKPC directly. Finally, the right point corresponds to $(\omega_f, \omega_b) = (0.738, 0.260)$ and is obtained by a nonlinear instrumental variable estimation.

[19]Mavroeidis, Plagborg-Møller and Stock (2014) review the extensive literature on the empirical literature of the NKPC and questions the robustness of the estimates provided by Galí, Gertler and Lopez-Salido (2005). This debate is beyond the scope of our paper. In any event, the exercise conducted next bypasses the estimation of the Hybrid NKPC on macroeconomic data and instead infers it by calibrating our theory to survey data on expectations.

C. *Bringing in the evidence on expectations*

As already noted, CG have run regression (23) using data from the Survey of Professional Forecasters.[20] Their main OLS specification, reported in column (1) of Table 1 of that paper, yields a mean estimate for K_{CG} equal to 1.193, with a standard deviation of 0.185. Translating the 95% confidence interval through the mapping developed in Section V yields the darker, thicker segment in Figure 2. This segment thus identifies the combinations of (ω_f, ω_b) that can be rationalized with a level of noise consistent with the expectation evidence in CG.

Clearly, only the third of the three estimates provided by Galí, Gertler and Lopez-Salido (2005), that corresponding to the furthest right point in the figure, is noticeably away from this segment. This happens to be the estimate that these authors trust the least for independent, econometric, reasons. We conclude that, when the theory is disciplined by the evidence in CG, it generates distortions broadly in line with existing estimates of the Hybrid NKPC. More succinctly, the informational friction implicit in the expectations data may alone account for all the observed inertia in inflation.

D. *A decomposition*

The quantitative implications of the theory are further illustrated in Figure 3. This figure compares the impulse response function of inflation under three scenarios. The solid line corresponds to frictionless benchmark. The dashed line corresponds to the frictional case, calibrated to the mean estimate of K_{CG} reported above. The circled dotted line is explained shortly.

As evident in the figure, the quantitative bite of the informational friction is significant: the impact effect on inflation is about 60% lower than its complete-information counterpart, and the peak of the inflation response is attained 5 quarters after impact rather than on impact. But what drives this quantitative bite? The lack of information about the real marginal cost (the PE component), or the beliefs about inflation (the GE component)?

The answer to this question is provided by the circled dotted line in Figure 3. This line represents a counterfactual that shuts down the effect of the informational friction on the expectations of the behavior of others (inflation) and isolates its effect on the the expectations of the fundamental (the real marginal cost). As evident in the figure, this counterfactual is very close to the complete-information benchmark and far away from the incomplete-information case. It follows that most of documented quantitative bite is due to the GE channel, or the anchoring of the expectations of inflation.[21]

---

[20]In the present context, it would be preferable to have an estimate of K_{CG} for the average forecasts of a representative sample of US firms. Such an estimate is lacking in the literature, but the evidence in Coibion and Gorodnichenko (2012) suggests that the friction among firms and consumers is, as one would expect, larger than that among professional forecasters.

[21]The decomposition offered in Figure 3 mirrors the decomposition of PE and GE effects introduced in Section III.A. See Online Appendix Section **??** for the detailed construction.
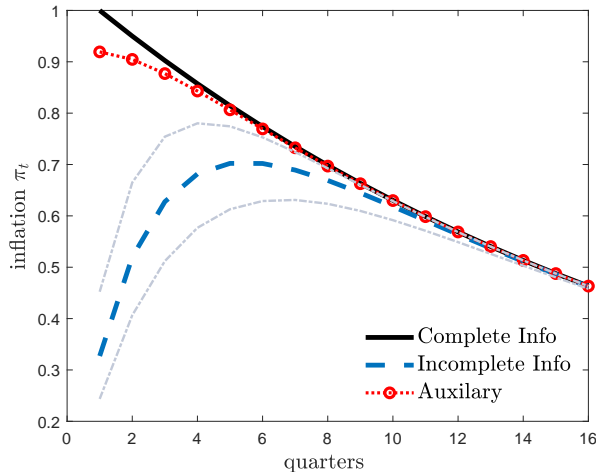
FIGURE 3. RESPONSE OF INFLATION TO HIGHER REAL MARGINAL COST

E. *Cross-sectional moments*

Thus far we have disregarded the individual-level evidence of Bordalo et al. (2020) and Broer and Kohlhas (2019). For the reasons explained in Section V, this evidence can be matched by letting agents be over- or under-confident, without influencing any of the preceding findings. This, however, does not mean that such evidence has no bite on the quantitative performance of the model. If we use the CG moment in combination with the individual-level counterpart estimated in the aforementioned papers, we can jointly identify $\hat{\sigma}$ and σ, the perceived and the actual level of noise. We can then further test the model by looking at its predictions for other, non-targeted moments, such as the cross-sectional dispersion of the individual forecast errors or that of the individual forecast revisions.

We implement this test in Table 1. We continue to denote with K_{CG} the coefficient of regression (23), and we denote with $K_{\text{BGMS/BK}}$ the individual-level counterpart. We then consider three sets of estimates for these coefficients. The first corresponds to Coibion and Gorodnichenko (2015) and to the exercise conducted above. The second and the third sets are from Bordalo et al. (2020) and Broer and Kohlhas (2019), respectively.[22] For each set, we report the identified belief parameters, the implied degrees of myopia and anchoring, and the model's predictions about the aforementioned cross-sectional moments. We finally compare the latter to their empirical counterparts.

As explained in the legend of the table, we consider two possible normalizations of the cross-sectional moments. Some normalization is needed because the analysis so far has been silent about the scale of the fluctuations in inflation. In one,

---

[22]Though both papers confirm that the original CG findings that K_{CG} is positive, they disagree on the sign of K_{BGMS}. This reflects differences in the treatment of outliers and other implementation details.

we normalize by the unconditional volatility of the quarter-to-quarter change in inflation. In the other, we normalize by the unconditional volatility of the level of inflation. We a priori prefer the first normalization, because our model is not supposed to capture low-frequency phenomena (e.g., great moderation) that may be "polluting" the second measure. But the model does a good job in both cases.

TABLE 1—MOMENTS ON AVERAGE AND INDIVIDUAL INFLATION FORECASTS

	K_{CG}	$K_{\mathrm{BGMS/BK}}$	$\hat{\sigma}$	σ	ω_f	ω_b
CG	1.19	0.00	1.76	1.76	0.52	0.43
BGMS	1.41	0.18	2.04	1.61	0.48	0.46
BK	1.27	-0.19	1.86	2.61	0.51	0.44

	Forecast error dispersion				Forecast revision dispersion			
	data^1	model^1	data^2	model^2	data^1	model^1	data^2	model^2
CG	2.07	2.03	0.40	0.24	1.94	1.63	0.38	0.19
BGMS	2.07	1.80	0.40	0.20	1.94	1.32	0.38	0.14
BK	2.07	2.98	0.40	0.34	1.94	2.31	0.38	0.26

Note: The three rows correspond to different estimates for K_{CG}, the coefficient of regression (23), and $K_{\mathrm{BGMS/BK}}$, the individual-level counterpart. In the first row, K_{CG} is taken form Panel B, Table 1 of Coibion and Gorodnichenko (2015), and $K_{\mathrm{BGMS/BK}}$ is fixed to zero. In the second row, both K_{CG} and $K_{\mathrm{BGMS/BK}}$ are taken from Table 3 of Bordalo et al. (2020). And in the third row, they are taken from Table 1 of Broer and Kohlhas (2019). The columns under forecast error dispersion correspond to the standard deviation of the cross-sectional forecast errors normalized by the standard deviation of either the quarter-to-quarter change in inflation (columns with superscript 1) or the level of inflation (with superscript 2). The columns under forecast revision dispersion correspond to the standard deviation of the cross-sectional forecast revisions with the same normalizations. We collect the forecast data from the survey of professional forecasters run by the Federal Reserve Bank of Philadelphia (Survey of Professional Forecasters, 1968-2018), and the real-time GDP deflator data from the Philadelphia Fed's Real-Time Dataset for Macroeconomists (Real-Time Data Set for Macroeconomists, 1968-2018).

F. Food for thought

We wrap up our application to inflation with a few additional insights about the possible determinants of the Hybrid NKPC implied by our analysis.

We start by studying the role of market concentration.[23] To this goal, we modify the micro-foundations as follows. There is now a continuum of markets, in each of which there is a finite number, $N \geq 2$, of competitors. We index the markets by $m \in [0, 1]$ and the firms in a given market by $i \in \{1, ..., N\}$. We let

---

[23]We thank a referee for suggesting this direction.

consumers have nested-CES preferences, so that the demand faced by firm i in market m is given by

$$Y_{i,m,t} = \left(\frac{P_{i,m,t}}{P_{m,t}}\right)^{-\psi} \left(\frac{P_{m,t}}{P_t}\right)^{-\varepsilon} Y_t,$$

where $P_{i,m,t}$ is the price of that firm, $P_{m,t}$ is the price index of the market that firm operates in, P_t is the aggregate price level, Y_t is aggregate income, $\psi > 1$ is the within-market elasticity of substitution and $\epsilon \in (0, \psi)$ is the cross-market counterpart. We finally assume that each firm has complete information about its own market but incomplete information about the entire economy.[24]

PROPOSITION 8: *In the economy described above, Corollary 2 continues to hold, modulo the following modification: both distortions decrease with market concentration (i.e., they increase with N).*

The intuition behind this result is that a higher degree of market concentration increases the strategic complementarity *within* markets and decreases it *across* markets. To the extent that firms know more about their own market than about the entire economy, this amounts to a lower bite of higher-order uncertainty, and therefore less myopia and less anchoring in the aggregate inflation dynamics.

This result links two empirical trends: the increase in market concentration (De Loecker, Eeckhout and Unger, 2020; Autor et al., 2020) and the reduction in inflation persistence (Cogley, Primiceri and Sargent, 2010; Fuhrer, 2010). Of course, this correlation does not establish causality. Still, the result illustrates how our analysis sheds new light on the possible determinants of inflation persistence.

We conclude with two additional ideas along these lines. The first one regards the conduct of monetary policy. Under the lens of our approach, a more hawkish monetary policy, such as that followed in the post-Volker era, is predicted to contribute to lower inflation persistence by reducing the effective degree of strategic complementarity in the firms' pricing decisions.

The second idea regards the economy's input-output structure. Rubbo (2020) has recently argued, in a setting abstracting from informational frictions, that changes in the input-output structure help explain the flattening of the NKPC. Our own analysis suggests that, in the presence of informational frictions, such changes may have also influenced the endogenous persistence in inflation, or the backward-looking component of the Hybrid NKPC.[25]

The exploration of these ideas is left for future work. But Section VIII paves the way for them by extending our tools to multi-variate systems and networks.

---

[24]The logic for the offered result requires only that information is more correlated within a market than across markets, or that firms face less higher-order uncertainty about their immediate links in the market network than about their remote links. The sharper assumption that firms face no higher-order uncertainty about their immediate links only simplifies the exposition.

[25]La'O and Tahbaz-Salehi (2020) make a similar point as Rubbo (2020) in a setting where nominal rigidity originates in incomplete information, but abstract from forward-looking behavior and learning, which are the forces highlighted here.

VII. Application to Consumption and Bridge to HANK

Now we turn to the effects of incomplete information on aggregate demand. As already shown in Corollary 3, the Euler equation is modified as if there is additional discounting together with habit persistence. In this section, we illustrate the quantitative potential of this idea. We also build a bridge to the HANK literature by showing that the habit-like sluggishness generated by the informational friction is amplified when the agents with the highest MPC are also the ones with the most cyclical income (Patterson, 2019; Flynn, Patterson and Sturm, 2019).

A. A HANK-like extension

We consider a perpetual-youth, overlapping-generations version of the New Keynesian model, along the lines of Piergallini (2007), Del Negro, Giannoni and Patterson (2015), and Farhi and Werning (2019). As in those papers, finite horizons (mortality risk) serve as convenient proxies for liquidity constraints, self-control problems, and other micro-level frictions that help explain why most estimates of the MPC in microeconomic data are almost an order of magnitude larger than that predicted by the textbook infinite-horizon model. We take this basic insight a step further by letting heterogeneity in mortality risk capture heterogeneity in the MPC. We couple this with heterogeneity in cyclical exposure. And, crucially, we let information be incomplete.

There are n types, or groups, of consumers, indexed by $g \in \{1, ..., n\}$, with respective mass π_g. In each period, a consumer in group g remains alive with probability $\varpi_g \in (0, 1]$; with the remaining probability, she dies and gets replaced by a new consumer of the same type. Consumers can trade actuarially fair annuities, so the return to saving, conditional on survival, is R_t/ϖ_g. This makes sure that the mortality risk does not distort intertemporal smoothing. Still, heterogeneity in ω_g matters because it maps to heterogeneity in MPCs. On top of that, different groups can have different exposures to the business cycle: the (log) income of group g is $y_{g,t} = \phi_g y_t$, where $\phi_g \geq 0$ is the elasticity of that group's income with respect to aggregate income and $\sum_g \pi_g \phi_g = 1$.

These assumptions allow us to study how the propagation mechanism under consideration, namely that related to incomplete information and higher-order beliefs, depends on heterogeneity in MPCs and business-cycle exposures. But they also open the door to a *separate* propagation mechanism: the dynamics of wealth inequality and the associated role of fiscal policy. To isolate the effects of interest, to nest the present application to the abstract analysis of Section VIII, and to obtain a sharp theoretical result (Proposition 9 below), we neutralize the second mechanism by letting appropriate fiscal transfers undo any wealth inequality triggered by interest-rate shocks.[26]

---

[26] An earlier draft had not clarified this assumption, without which the wealth distribution becomes a relevant state variable for the aggregate dynamics. We thank Dmitriy Sergeyev for pointing out this. See Online Appendix Section **??** for details.

As shown in Online Appendix Section **??**, the group-level spending can be expressed as follows:

$$(26) \qquad c_{g,t} = m_g\phi_g \sum_{k=0}^{\infty}(1-m_g)^k\overline{\mathbb{E}}_t^g[c_{t+k}] - (1-m_g)\sum_{k=0}^{\infty}(1-m_g)^k\overline{\mathbb{E}}_t^g[r_{t+k}],$$

where $m_g \equiv 1 - \chi\varpi_g$, χ is the subjective discount rate, and $\overline{\mathbb{E}}_t^g$ is the average expectation. For each g, equation (26) follows from aggregating the consumption functions of the individuals within group g and replacing their income in terms of aggregate consumption. The collection of these equations across g recasts the demand block of the economy as a dynamic network among the various groups of consumers. This echoes Auclert, Rognlie and Straub (2019), which develops similar network representations for more general HANK economies.

Inspection of (26) reveals, first, that m_g identifies the MPC of group g and, second, that the strategic complementarity, or the Keynesian cross, depends on how the product $m_g\phi_g$ varies across groups, or whether a higher MPC is positively correlated with a higher business-cycle exposure. Patterson (2019) provides evidence of such a positive correlation and shows how it translates to a steeper Keynesian cross in a static, complete-information context. In the light of our insight of how the as-if distortions introduced by informational frictions depend on GE feedback mechanisms, one may expect such a positive correlation to translate also to more myopia and habit-like persistence in the aggregate consumption dynamics. We verify this intuition in part (iii) below, at least under the simplifying assumption of two groups.

PROPOSITION 9 (HANK): *(i) Under complete information, there exists a scalar $\varsigma > 0$ such that aggregate consumption obeys a textbook Euler condition of the following form:*

$$c_t = -\varsigma r_t + \mathbb{E}_t[c_{t+1}].$$

(ii) Under incomplete information, there exist scalars $\omega_f < 1$ and $\omega_b > 0$ such that aggregate consumption obeys a hybrid Euler condition of the form:

$$c_t = -\varsigma r_t + \omega_f\mathbb{E}_t[c_{t+1}] + \omega_b c_{t-1},$$

where the scalar $\varsigma > 0$ is the same as that under complete information and the scalars $\omega_f < 1$ and $\omega_b > 0$ are functions of σ and $(\pi_g, m_g, \phi_g)_{g\in\{1,...,n\}}$.

(iii) Suppose there are two groups, with $m_1 > m_2$. An increase in ϕ_1, the business-cycle exposure of high-MPC group, maps to a lower ω_f and a higher ω_b, that is, more as-if myopia and anchoring in the aggregate dynamics.

Part (i) mirrors an irrelevance result from Werning (2015). With complete information, the DIS curve of our HANK economy is the same as a representative agent's Euler condition. There is neither extra discounting of the future nor habit-like persistence. Heterogeneity matters at most for ς, the elasticity of aggregate

consumption with respect to the real interest rate.

Part (ii) extends Corollary 3 to heterogeneity in MPC and business-cycle exposure. Once again, incomplete information amounts to adding myopia and habit-like persistence in the DIS curve. But now heterogeneity interacts with information in shaping the magnitude of these distortions.

Part (iii) completes the picture by showing how exactly heterogeneity matters. An increase in the business-cycle exposure of the high-MPC group (and a corresponding reduction in the business-cycle exposure of the low-MPC group) translates to both more myopia and more habit-like persistence.

The basic logic behind this result was anticipated above. Its proof utilizes the techniques developed in Section VIII. In the remainder of this section, we use a numerical example to illustrate our findings.

B. *Numerical Example*

Figure 4 compares four economies. The first one corresponds to the textbook, representative-agent benchmark. We refer to this benchmark as "Complete Information" in the figure. The second economy is a variant of the first one that adds habit persistence, of the type and magnitude found in the DSGE literature.[27] We refer to this economy as "Complete Info + Habit." The remaining two economies remove habit but add incomplete information. Both of them feature an average MPC equal to $\overline{m} = .30$, which is roughly consistent with the relevant evidence. The one referred to as "Incomplete Info" in the figure, abstracts from heterogeneity; this is the economy described in Corollary 3. The other one, which is referred to as "Incomplete Info + HANK" in the figure, adds heterogeneity: there are two groups of consumers, with $m_1 = .55$, $m_2 = .05$, $\phi_1 = 2$, and $\phi_2 = 0$.[28]

Let us first compare "Incomplete Info" to "Complete Info + Habit." This extends the lesson of the previous section from the inflation context to the consumption context: the informational friction alone generates a similar degree of sluggishness as that generated by habit persistence in the DSGE literature. Importantly, whereas the degree of habit assumed in that literature is far larger than that supported by micro-economic evidence (Havranek, Rusnak and Sokolova, 2017), the informational friction assumed here is broadly consistent with survey evidence. This illustrates how our approach help merge the gap between the micro and macro estimates of habit.

---

[27]In particular, we assume external habit and specify the per-period utility as $\log(C_t - b\bar{C}_t)$, where C_t and $\bar{C}_t$ denote, respectively, own consumption and aggregate consumption. In equilibrium, $\bar{C}_t = C_t$ and the log-linearized Euler condition reduces to the following low of motion of consumption:

$$c_t = -\frac{1-b}{1+b}r_t + \frac{1}{1+b}\mathbb{E}_t[c_{t+1}] + \frac{b}{1+b}c_{t-1}.$$

We finally set $b = .7$, which is in the middle of the macro-level estimates reported in the meta-analysis by Havranek, Rusnak and Sokolova (2017).

[28]For the incomplete-information economies, we target $K_{\mathrm{CG}} = 0.9$. This is in the middle of the range of values Angeletos, Huo and Sastry (2021) estimate when they repeat the CG regression on forecasts of unemployment, with the rationale being that unemployment is a proxy for the output gap in the model.
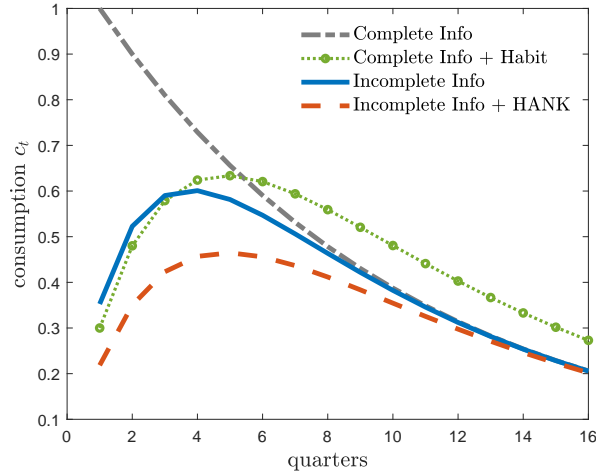
FIGURE 4. RESPONSE OF CONSUMPTION TO LOWER INTEREST RATES

Relatedly, if we consider an extension with transitory idiosyncratic income shocks along the lines of Online Appendix Section **??**, our economy can feature simultaneously two properties: a large and front-loaded response to such shocks at the micro level, in line with the relevant microeconomic evidence; and a dampened and sluggish response to monetary policy at the macro level, in line with the relevant macroeconomic evidence. By contrast, if there was true habit persistence in consumption of the kind and level assumed in the DSGE literature, the micro-level responses would also be dampened and sluggish, contradicting the relevant microeconomic evidence. This idea is pushed further, and is more carefully quantified, in a recent paper by Auclert, Rognlie and Straub (2020).

Finally, let us inspect the economy "Incomplete Info + HANK." Needless to say, this economy is not meant to capture a realistic degree of heterogeneity: our two-group specification is only a gross approximation to the kind of heterogeneity captured in the quantitative HANK literature (e.g., Kaplan and Violante, 2014; Kaplan, Moll and Violante, 2018)). Nevertheless, this economy helps illustrate how such heterogeneity, and in particular the kind of positive cross-sectional correlation between MPCs and income cyclicality documented in Patterson (2019), can reinforce both the habit-like sluggishness and the myopia-like dampening generated by incomplete information.

C.   *Informational friction plus wealth dynamics*

In the preceding analysis we used appropriate fiscal transfers to make sure that the wealth distribution is not a state variable for the aggregate dynamics and to nest the exercise into the analysis of Section VIII. We now shut down these transfers and study how the endogenous dynamics of wealth matter both in isolation and in combination with our mechanisms.

Consider first the case with complete information and suppose again that there are two groups, with only the high-MPC group being exposed to the business cycle ($\omega_1 < \omega_2$ and $\phi_1 > 0 = \phi_2$), and consider a negative innovation in η_t. This causes, in equilibrium, an expansion. But because only the first group's income is exposed to it, and because the income increase is less than permanent, this group will try to save some of this increase, while the second group has no such incentive. Along with the fact that the total saving of the two groups has to be zero, this explains why the first group responds to the shock by saving and accumulating wealth, whereas the second group responds by borrowing and accumulating debt. But since the first group has a larger MPC, the accumulation of wealth by this group helps increase aggregate spending in the future. This suggests that, even with complete information, the wealth dynamics add persistence to the response of aggregate demand to interest-rate shocks.
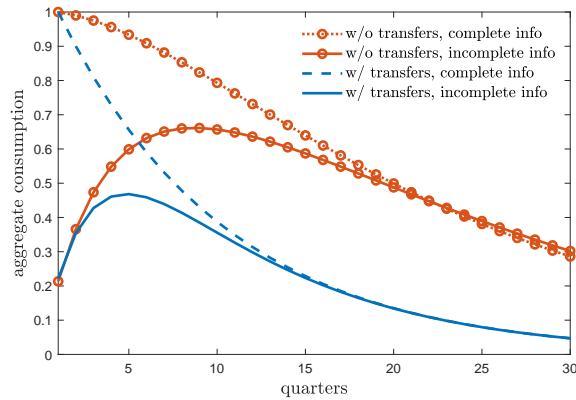

FIGURE 5. SHUTTING DOWN THE FISCAL TRANSFERS

We verify this intuition in Figure 5 and proceed to show how this source of persistence extends to the case of incomplete information, without however upsetting, and indeed only reinforcing, our own message. This figure compares the response of consumption to a negative (expansionary) interest rate shock under four scenarios. Two of them replicate the complete-information and the incomplete-information HANK cases from Figure 4. The remaining two show how the results change when fiscal transfers are switched off and, equivalently, the aforementioned wealth channel is switched on. Regardless of the information structure, this channel adds persistence.[29] The effect of the informational friction, which is our own focal point, is qualitatively the same whether the wealth channel is present or not. Perhaps more interestingly, the two mechanisms rein-

[29]This channel also adds amplification. To focus on the persistence effects, in the figure we renormalize the magnitude of the shock as we change the fiscal rule so that the complete-information response of consumption on impact remains 1.

force each other, yielding a much more pronounced hump-shaped response than each mechanism alone.

VIII. Multivariate Systems, or Networks

We close the paper with the extension of our analytical results to multi-variate systems, or networks. We already made implicit use of this extension in our HANK application. Here, we fill in the details and develop tools that could aid analytical and quantitative evaluations of how informational frictions and network structures interact in a variety of applications.

The economy consists of n groups, each containing a continuum of agents. Groups are indexed by $g \in \{1, ..., n\}$, agents by (i, g) where $i \in [0, 1]$ is an agent's name and g her group affiliation (e.g., consumer or firm). The best response of agent i in group g is specified as follows:

$$(27) \qquad a_{i,g,t} = \varphi_g \mathbb{E}_{i,g,t}[\xi_t] + \beta_g \mathbb{E}_{i,g,t}[a_{i,g,t+1}] + \sum_{j=0}^{n} \gamma_{gj} \mathbb{E}_{i,g,t}[a_{j,t+1}].$$

The parameter φ_g captures the direct, contemporaneous exposure of an agent in group g to the exogenous shock, holding constant her expectations of both her own future actions and the actions of others. The parameter $\{\beta_g\}$ captures the additional, forward-looking, PE effect that obtains because of the consideration of own future actions. Finally, the parameter $\{\gamma_{g,j}\}$ captures the dependence of the optimal action of an agent in group g to her expectation of the average action of group j. This allows for rich strategic of GE interactions both within groups (when $j = g$) and across groups (when $j \neq g$).[30]

Turning now to the information structure, this is specified as a collection of private Gaussian signals, one per agent and per period. The period-t signal received by agent i in group t is given by

$$(28) \qquad x_{i,g,t} = \xi_t + u_{i,g,t}, \quad u_{i,g,t} \sim \mathcal{N}(0, \sigma_g^2).$$

where $\sigma_g \geq 0$ parameterizes the noise of group g. Notice that, by allowing σ_g to differ across g, we can accommodate information heterogeneity in addition to payoff and strategic heterogeneity. For instance, firms could be more informed than consumers on average, and "sophisticated" consumers could be more informed than "unsophisticated" ones.

Let $\boldsymbol{a}_t = (a_{g,t})$ be a column vector collecting the aggregate actions of all the groups (e.g., the vector of aggregate consumption and aggregate inflation). Let

---

[30]Like our baseline framework, the extension considered here rules out the dependence of an agent's best response on the *concurrent* choices of others. This, however, is without serious loss of generality for two reasons. First, in all applications of interest, this dependence vanishes as the length of the time period goes to zero. Second, if we incorporate a general form of such dependence by adding the term $\sum_j \alpha_{g,i} \mathbb{E}_{i,j,t}[a_{i,j,t}]$ in equation (27), the results stated below, namely Propositions 10 and 11, continue to hold, modulo a minor adjustment in the cubic that appears in condition (29).

$\boldsymbol{\varphi} = (\varphi_g)$ be a column vector containing the value of φ_g across the groups. Let $\boldsymbol{\beta} = \mathbf{diag}\{\beta_g\}$ be a $n \times n$ diagonal matrix whose off-diagonal elements are zero and whose diagonal elements are the values of β_g across groups. Finally, let $\boldsymbol{\gamma} = (\gamma_{gk})$ be an $n \times n$ matrix collecting the interaction parameters, γ_{gj}, and let $\boldsymbol{\delta} \equiv \boldsymbol{\beta} + \boldsymbol{\gamma}$. Similarly to Section II, we impose that $\beta_g \in (0, 1)$ and the spectral radius of $(\mathbf{I} - \boldsymbol{\beta})^{-1}\boldsymbol{\gamma}$ is less than 1. The following extensions of Propositions 2 and 3 are then possible.

PROPOSITION 10 (Solution): *There exists a unique equilibrium, and the aggregate outcome $a_{g,t}$ of each group g is given by*

$$(29) \qquad a_{g,t} = \sum_{j=1}^{n} \psi_{g,j} \left\{ \frac{1 - \frac{\vartheta_j}{\rho}}{1 - \vartheta_j L} \xi_t \right\},$$

*where $\{\psi_{g,j}\}$ are fixed scalars, characterized in Online Appendix Section **??**, and $\{\vartheta_g\}$ are the inverse of the outside roots of the following polynomial:*
(30)
$$C(z) = \det\left((\boldsymbol{\delta} - \boldsymbol{\gamma} - \mathbf{I}z) \, \mathbf{diag}\left\{ z^2 - \left(\rho + \frac{1}{\rho} + \frac{1}{\rho \sigma_g^2} \right) z + 1 \right\} - z \, \mathbf{diag}\left\{ \frac{1}{\rho \sigma_g^2} \right\} \boldsymbol{\gamma} \right).$$

PROPOSITION 11 (Observational Equivalence): *There exist matrices $\boldsymbol{\omega}_f$ and $\boldsymbol{\omega}_b$ such that the incomplete-information economy is observationally equivalent to the following complete-information economy:*

$$(31) \qquad \boldsymbol{a}_t = \boldsymbol{\varphi}\xi_t + \boldsymbol{\omega}_f \boldsymbol{\delta}\mathbb{E}_t[\boldsymbol{a}_{t+1}] + \boldsymbol{\omega}_b \boldsymbol{a}_{t-1}.$$

One subtlety with representation (31) is that it is not unique: there are multiple values of the matrices $\boldsymbol{\omega_f}$ and $\boldsymbol{\omega_b}$ that replicate the incomplete-information equilibrium. Intuitively, it is possible to make agents myopic vis-a-vis the future by letting them discount enough either only their own group's future actions, or the future actions of other groups too.[31] This complicates the interpretation and the comparative statics of the provided representation but is of little substantial consequence: although the representation in terms of condition (31) is not unique, the equilibrium itself is determinate, and so are its observable properties, which can be directly obtained from Proposition 10.

Proposition 10 is indeed quite telling. It shows that the equilibrium outcome can now be expressed as a linear combination of n terms, each of which is an AR(2) process that has a similar structure as in our baseline analysis. The one root of these processes is the same across g and is given, naturally, by that of the fundamental. The other root, denoted above by ϑ_g, encodes how the information

---

[31] Indeed, both of the following two choices are possible: let $\boldsymbol{\omega}_f$ have unit off-diagonal elements, meaning that a distortion is applied only to expectations of own-group future outcomes; or let the elements of each row of $\boldsymbol{\omega}_f$ be the same, meaning that the same distortion is applied to all expectations. If one of these choices is made, there is no residual indeterminacy.

friction faced by group g interacts with the network structure of the economy.

In the knife-edge case in which $\boldsymbol{\gamma}$ is diagonal, meaning that the behavior of each group is independent of that of other groups, each ϑ_g is pinned down by the characteristics of group g alone and the outcome of that group is given by the corresponding AR(2) process alone ($\psi_{g,j} = 0$ for $j \neq g$). For generic $\boldsymbol{\gamma}$, instead, each ϑ_g depends on the entire $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ matrices, that is, on all the PE and GE parameters, as well as on all the information parameters. Furthermore, the outcome of a group depends on all the n different AR(2) processes.

To illustate how the network structure matters, let $\boldsymbol{\beta} = 0$ and $\sigma_g = \sigma$ for all g. In this case, we show in Online Appendix Section **??** that the polynomial given in condition (30) reduces to the product of n quadratics, one for each ϑ_g. Furthermore, each ϑ_g is determined in the same manner as in our baseline analysis, namely as the reciprocal of the largest solution of cubic (17), with the g-th eigenvalue of the matrix $\boldsymbol{\gamma}$ in place of the scalar γ. Because the eigenvalues of $\boldsymbol{\gamma}$ encode the GE feedback both within and across groups, we have that an increase in either kind of feedback maps to a higher ϑ_g and, thereby, to both less amplitude and more volatility. The essence of our baseline analysis is thus fully preserved.

Finally, note that the results presented here not only offer a robustness of our main insights to multi-variate systems and networks, but also a straightforward numerical algorithm: one only needs to solve the polynomial in condition (30).

IX. Conclusion

We developed a toolbox for analyzing and quantifying the equilibrium effects of informational frictions and of the associated higher-order uncertainty. We represented these effects as the combination of two behavioral distortions: a form of myopia, or extra discounting of the future; and a form of habit, or anchoring of current behavior to past behavior. We further showed how these as-if distortions increase with the strength of the underlying strategic interaction or GE feedback, and how they can be disciplined with available evidence on expectations. And we used these results to argue that the friction implicit in survey evidence of expectations is large enough to generate a comparable amount of sluggishness in the dynamics of inflation and aggregate spending as that captured in the DGSE literature with more ad hoc modeling devices.

While connecting the theory to the available evidence on expectations, we clarified *which* such evidence is best suited for the purpose of quantifying the distortions of interest: it is evidence on *average* forecasts, such as that provided in Coibion and Gorodnichenko (2015), as opposed to evidence on *individual* forecasts, such as that provided in Bordalo et al. (2020) and Broer and Kohlhas (2019). Left outside this paper was a more comprehensive investigation of the lessons contained in surveys of expectations for macroeconomic theory.

We undertake this task in a follow-up paper (Angeletos, Huo and Sastry, 2021). There, we use a variety of existing evidence along with new evidence of our own

to argue that, among a large set of candidate theories, the one that best accounts for the joint dynamics of inflation, aggregate spending and forecasts thereof in the US is a theory that blends two frictions: incomplete information or rational inattention, as in the present paper and the literature we have built on; and over-extrapolation, as in Greenwood and Shleifer (2014) and Gennaioli, Ma and Shleifer (2015). This points in the opposite direction than cognitive discounting and level-k thinking, two close cousins of under-extrapolation, but leaves room for the kinds of myopia and anchoring accommodated via our approach.

Another element of our contribution was to extend our tools to multi-variate systems and networks. We illustrated the use of these extended tools within a HANK economy. Other possible applications include production networks, whether in the context of the NKPC (La'O and Tahbaz-Salehi, 2020; Rubbo, 2020) or in the context of the RBC framework (Acemoglu et al., 2012; Baqaee and Farhi, 2019; Nimark, Chahrour and Pitschner, 2019), as well as dynamic extensions of the more abstract incomplete-information networks studied in Bergemann, Heumann and Morris (2017).

REFERENCES

Acemoglu, Daron, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. 2012. "The Network Origins of Aggregate Fluctuations." *Econometrica*, 80(5): 1977–2016.

Angeletos, George-Marios, and Alessandro Pavan. 2007. "Efficient Use of Information and Social Value of Information." *Econometrica*, 75(4): 1103–1142.

Angeletos, George-Marios, and Chen Lian. 2018. "Forward Guidance without Common Knowledge." *American Economic Review*, 108(9): 2477–2512.

Angeletos, George-Marios, Zhen Huo, and Karthik Sastry. 2021. "Imperfect Macroeconomic Expectations: Evidence and Theory." In *NBER Macroeconomics Annual 2020.* Vol. 35, Chapter 1.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub. 2019. "The Intertemporal Keynesian Cross." *mimeo.*

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub. 2020. "Micro Jumps, Macro Humps: monetary policy and business cycles in an estimated HANK model." *mimeo.*

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen. 2020. "The fall of the labor share and the rise of superstar firms." *Quarterly Journal of Economics*, 135(2): 645–709.

Baqaee, David Rezza, and Emmanuel Farhi. 2019. "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem." *Econometrica*, 87(4): 1155–1203.

Bergemann, Dirk, Tibor Heumann, and Stephen Morris. 2017. "Information and Interaction." *Cowles Foundation Discussion Paper*, , (2088).

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer. 2020. "Overreaction in Macroeconomic Expectations." *American Economic Review*, 110(9): 2748–82.

Broer, Tobias, and Alexandre Kohlhas. 2019. "Forecaster (Mis-)Behavior." *IIES miméo*.

Carvalho, Carlos, Stefano Eusepi, Emanuel Moench, and Bruce Preston. 2017. "Anchored Inflation Expectations." *mimeo.*

Choi, Yongok, Giacomo Rondina, and Todd B Walker. 2020. "Information Aggregation Bias and Samuelson's Dictum." *mimeo.*

Cogley, Timothy, Giorgio E Primiceri, and Thomas J Sargent. 2010. "Inflation-gap persistence in the US." *American Economic Journal: Macroeconomics*, 2(1): 43–69.

Coibion, Olivier, and Yuriy Gorodnichenko. 2012. "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy*, 120(1): 116–159.

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 P Giannoni, and Christina Patterson. 2015. "The Forward Guidance Puzzle." *FRB of New York mimeo.*

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger. 2020. "The rise of market power and the macroeconomic implications." *Quarterly Journal of Economics*, 135(2): 561–644.

Evans, George W, and Seppo Honkapohja. 2012. *Learning and expectations in macroeconomics.* Princeton University Press.

Farhi, Emmanuel, and Iván Werning. 2019. "Monetary Policy, Bounded Rationality, and Incomplete Markets." *American Economic Review*, 109(11): 3887–3928.

Flynn, Joel, Christina Patterson, and John Sturm. 2019. "Shock Propagation and the Fiscal Multiplier: The Role of Heterogeneity." *MIT and Northwestern University mimeo.*

Fuhrer, Jeffrey C. 2010. "Inflation persistence." In *Handbook of monetary economics.* Vol. 3, 423–486. Elsevier.

Gabaix, Xavier. 2020. "A Behavioral New Keynesian Model." *American Economic Review*, 110(8): 2271–2327.

Galí, Jordi, and Mark Gertler. 1999. "Inflation dynamics: A structural econometric analysis." *Journal of Monetary Economics*, 44(2): 195–222.

Galí, Jordi, Mark Gertler, and J David Lopez-Salido. 2005. "Robustness of the estimates of the hybrid New Keynesian Phillips curve." *Journal of Monetary Economics*, 52(6): 1107–1118.

Garcıa-Schmidt, Mariana, and Michael Woodford. 2019. "Are Low Interest Rates Deflationary? A Paradox of Perfect-Foresight Analysis." *American Economic Review*, 109(1): 86–120.

Gennaioli, Nicola, Yueran Ma, and Andrei Shleifer. 2015. "Expectations and investment." In *NBER Macroeconomics Annual 2015, Volume 30.* University of Chicago Press.

Golub, Ben, and Stephen Morris. 2019. "Expectations, Networks and Conventions." *mimeo.*

Greenwood, Robin, and Andrei Shleifer. 2014. "Expectations of returns and expected returns." *Review of Financial Studies*, 27(3): 714–746.

Havranek, Tomas, Marek Rusnak, and Anna Sokolova. 2017. "Habit Formation in Consumption: A Meta-Analysis." *European Economic Review*, 95: 142–167.

Huo, Zhen, and Marcelo Zouain Pedroni. 2020. "A Single-Judge Solution to Beauty Contests." *American Economic Review*, 110(2): 526–68.

Huo, Zhen, and Naoki Takayama. 2018. "Rational Expectations Models with Higher Order Beliefs." Yale University miméo.

Iovino, Luigi, and Dmitriy Sergeyev. 2017. "Quantitative Easing without Rational Expectations." *Work in progress.*

Jung, Jeeman, and Robert J. Shiller. 2005. "Samuelson's Dictum and the Stock Market." *Economic Inquiry*, 43(2): 221–228.

Kaplan, Greg, and Giovanni L. Violante. 2014. "A Model of the Consumption Response to Fiscal Stimulus Payments." *Econometrica*, 82(4): 1199–1239.

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante. 2018. "Monetary Policy According to HANK." *American Economic Review*, 108(3): 697–743.

Kiley, Michael T. 2007. "A Quantitative Comparison of Sticky-Price and Sticky-Information Models of Price Setting." *Journal of Money, Credit and Banking*, 39(s1): 101–125.

Kohlhas, Alexandre, and Ansgar Walther. 2019. "Asymmetric Attention." *IIES miméo.*

La'O, Jennifer, and Alireza Tahbaz-Salehi. 2020. "Optimal Monetary Policy in Production Networks." *mimeo.*

Lucas, Robert E. Jr. 1972. "Expectations and the Neutrality of Money." *Journal of Economic Theory*, 4(2): 103–124.

Maćkowiak, Bartosz, and Mirko Wiederholt. 2009. "Optimal Sticky Prices under Rational Inattention." *American Economic Review*, 99(3): 769–803.

Maćkowiak, Bartosz, and Mirko Wiederholt. 2015. "Business Cycle Dynamics under Rational Inattention." *Review of Economic Studies*, 82(4): 1502–1532.

Malmendier, Ulrike, and Stefan Nagel. 2016. "Learning from inflation experiences." *Quarterly Journal of Economics*, 131(1): 53–87.

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): 1295–1328.

Matejka, Filip. 2016. "Rationally Inattentive Seller: Sales and Discrete Pricing." *Review of Economic Studies*, 83(3): 1125–1155.

Mavroeidis, Sophocles, Mikkel Plagborg-Møller, and James H Stock. 2014. "Empirical evidence on inflation expectations in the New Keynesian Phillips Curve." *Journal of Economic Literature*, 52(1): 124–88.

Melosi, Leonardo. 2016. "Signalling effects of monetary policy." *Review of Economic Studies*, 84(2): 853–884.

Morris, Stephen, and Hyun Song Shin. 1998. "Unique Equilibrium in a Model of Self-fulfilling Currency Attacks." *American Economic Review*, 587–597.

Morris, Stephen, and Hyun Song Shin. 2002. "Social Value of Public Information." *American Economic Review*, 92(5): 1521–1534.

Morris, Stephen, and Hyun Song Shin. 2006. "Inertia of Forward-looking Expectations." *The American Economic Review*, 152–157.

Nimark, Kristoffer. 2008. "Dynamic Pricing and Imperfect Common Knowledge." *Journal of Monetary Economics*, 55(2): 365–382.

Nimark, Kristoffer, Ryan Chahrour, and Stefan Pitschner. 2019. "Sectoral Media Focus and Aggregate Fluctuations." *mimeo.*

Patterson, Christina. 2019. "The Matching Multiplier and the Amplification of Recessions." *Northwestern University mimeo.*

Piergallini, Alessandro. 2007. "Real Balance Effects and Monetary Policy." *Economic Inquiry*, 44(3): 497–511.

Real-Time Data Set for Macroeconomists. 1968-2018. "Data Files - Real-Time Data Set (NOUTPUT)." https://www.philadelphiafed.org/surveys-and-data/real-time-data-research/noutput (accessed February 15, 2021).

Reis, Ricardo. 2006. "Inattentive producers." *Review of Economic Studies*, 73(3): 793–821.

Rubbo, Elissa. 2020. "Networks, Phillips Curves, and Monetary Policy." *mimeo.*

Sargent, Thomas J. 1993. *Bounded rationality in macroeconomics.* Oxford University Press.

Sims, Christopher A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics*, 50(3): 665–690.

Survey of Professional Forecasters. 1968-2018. "Individual Forecasts: Survey of Professional Forecasters." https://www.philadelphiafed.org/surveys-and-data/real-time-data-research/individual-forecasts (accessed February 15, 2021).

Tirole, Jean. 2015. "Cognitive Games and Cognitive Traps." *Toulouse School of Economicss mimeo.*

Werning, Iván. 2015. "Incomplete Markets and Aggregate Demand." *NBER Working Paper No. 21448.*

Woodford, Michael. 2003. "Imperfect Common Knowledge and the Effects of Monetary Policy." *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps.*