American Economic Journal: Macroeconomics 2023, 15(1): 173–208
https://doi.org/10.1257/mac.20200053

Robust Predictions for DSGE Models with Incomplete Information†

By Ryan Chahrour and Robert Ulbricht*

We provide predictions for DSGE models with incomplete information that are robust across information structures. Our approach maps an incomplete-information model into a full-information economy with time-varying expectation wedges and provides conditions that ensure the wedges are rationalizable by some information structure. Using our approach, we quantify the potential importance of information as a source of business cycle fluctuations in an otherwise frictionless model. Our approach uncovers a central role for firm-specific demand shocks in supporting aggregate confidence fluctuations. Only if firms face unobserved local demand shocks can confidence fluctuations account for a significant portion of the US business cycle. (*JEL* D82, D83, E13, E31, E32)

What are the sources of aggregate fluctuations? One common view is that business cycles are caused by shocks to the confidence of consumers and firms. The literature on business cycles has formalized this view in several ways, including modeling confidence fluctuations as a consequence of incomplete information (e.g., Lorenzoni 2009; Angeletos and La'O 2013; Benhabib, Wang, and Wen 2015). Yet, relatively few of these information-based models have been investigated quantitatively. This is at least in part because the private information structures governing people's beliefs are hard to observe in the data or—as argued by Sims (2003) and Woodford (2003)—may have no observable counterpart.

In this paper we quantify the potential importance of confidence-driven business cycles using a novel approach that bypasses the challenge of postulating ad hoc information structures. The approach takes the *economic environment* (technology, preferences, market structure) as given, but does not require a complete specification of the *information structure* that governs people's beliefs. Instead, we provide

* Chahrour: Boston College, Department of Economics (email: ryan.chahrour@bc.edu); Ulbricht: Boston College, Department of Economics (email: ulbricht@bc.edu). Simon Gilchrist was coeditor for this article. We thank three anonymous referees and our discussants Guillaume Chevillon, Stéphane Dupraz, Alexandre Kohlhas, Jianjun Miao, and Anna Orlik for valuable comments. Ulbricht gratefully acknowledges financial support from the European Union's Horizon 2020 Program under grant agreement no. 649396.
† Go to https://doi.org/10.1257/mac.20200053 to visit the article page for additional materials and author disclosure statements or to comment in the online discussion forum.

an "information-robust" characterization of all equilibria that are possible within a given economic environment.

Methodological Contribution.—We develop our methodology for a canonical class of models with dispersed or incomplete information, without any restriction on the set of signals governing people's beliefs regarding their own idiosyncratic shocks, the aggregate state of the economy, what other agents believe, and so on. Notably, our general framework encompasses virtually all linear rational expectations DSGE models explored in the literature. We show how to map these models into a "primal" economy in which all agents have full information and where deviations from full information are summarized by exogenous wedges in agents' equilibrium expectations. We then develop necessary and sufficient conditions for the existence of an information structure that is consistent with the expectation errors captured by these wedges. Subject to these conditions, the primal economy is isomorphic to the incomplete-information economy.

Exploiting this equivalence, we derive a complete characterization of all information equilibria within a given economic environment. Specifically, our characterization allows the researcher to specify a (possibly empty) *minimal information set* reflecting their prior of what constitutes a lower bound on agents' information. Our main theorem then states that an equilibrium of the primal economy corresponds to an equilibrium of the information economy if and only if the expectation errors captured by the exogenous wedges are orthogonal to the corresponding agent's actions and each element of that agent's minimal information set. In our applications we show how to use this characterization to draw concrete economic conclusions about equilibrium in the incomplete information model, without ever completely specifying the information available to agents.

Applied Contribution.—To demonstrate the usefulness of our approach, we use it to ask: *under what conditions* can changes in confidence generate sizable fluctuations in aggregate economic activity? As an illustration, we first examine this question in the context of a simple price-setting model similar to the one in Woodford (2003). The model describes the problem of price-setting firms that face exogenous aggregate demand and downward-sloping individual demand functions. Applied to this model, our methodology can be used to analytically bound the variances of endogenous variables, to sign cross-covariances among them, and to limit their autocorrelations. Among our results, we find that any information structure that allows firms to contemporaneously observe their own sales implies that aggregate inflation must be procyclical. Moreover, if either idiosyncratic or aggregate demand is observed (or constant), then aggregate output does not fluctuate.

After demonstrating our approach in this simple context, we then use it to explore the potential for confidence-driven business cycles quantitatively. Our quantitative model is a flexible price business cycle model without capital, in which households and firms live on informationally disparate "islands." The inclusion of households introduces the potential for additional aggregate demand channels that act through incomplete information. Like the price-setting example, firms on each island experience fluctuations in local demand. In addition, we allow for exogenous fluctuations in aggregate productivity, as well as temporary and persistent changes in firm-level productivity.

Whether the model generates aggregate fluctuations beyond those induced by aggregate productivity shocks depends on its ability to generate expectation errors that are correlated in the cross section. There are two potential sources of such correlation. First, agents can be jointly optimistic or pessimistic regarding the aggregate state of productivity, as in Lorenzoni (2009) or Angeletos and La'O (2010). Second, agents can be jointly optimistic or pessimistic about their own idiosyncratic conditions, as in Angeletos and La'O (2013) or Benhabib, Wang, and Wen (2015), possibly accentuated by strategic uncertainty. Both channels are disciplined by the properties of the fundamental shocks to productivity and demand. Our approach allows us to provide a general characterization of these restrictions that does not hinge on specific structural assumptions about people's information.

For reference, we first establish a novel theoretical benchmark for the case in which the stochastic process governing idiosyncratic shocks is *unrestricted* by data. For this case, we show that confidence-driven fluctuations can in principle generate *any* autocovarince structure for output and inflation, bypassing all cross-equation restrictions that obtain under full information, provided that agents do not perfectly observe demand for their local goods when making production choices. This result extends findings of Angeletos and La'O (2013) and Benhabib, Wang, and Wen (2015) that correlated information shocks can generate arbitrary macroeconomic volatility if idiosyncratic shocks are sufficiently volatile.

In light of this benchmark, we next ask: how much expectations-driven volatility can one generate for a realistic calibration of idiosyncratic shocks? We explore this question by calibrating the processes for idiosyncratic productivity and demand using existing microdata estimates (Foster, Haltiwanger, and Syverson 2008). We then compute global upper bounds on confidence-induced output fluctuations, their persistence, and the contemporaneous correlation with inflation.

For an empirically plausible calibration, we find that the volatility frontier for confidence-induced output fluctuations is hump-shaped in aggregate persistence and is decreasing in the contemporaneous correlation with inflation. For an aggregate persistence and inflation cyclicality consistent with US data, the maximal one-step-ahead volatility of confidence-induced fluctuations in output is 0.011 (approximately 90 percent of its empirical counterpart). We demonstrate that the ability to generate sizable macro volatility through confidence fluctuations hinges critically on the volatility of micro shocks to firm demand. By contrast, micro shocks to productivity play a somewhat dispensable role for generating aggregate volatility.

Why does idiosyncratic product demand play such an important role in supporting aggregate fluctuations? The answer has two key components. First, informed by the empirical evidence of Foster, Haltiwanger, and Syverson (2008), firm-specific demand fluctuations in our calibration are large, in particular relative to idiosyncratic productivity.[1] Idiosyncratic demand realizations therefore drive large fluctuations in

[1] See de Loecker (2011); Demidova, Kee, and Krishna (2012); Roberts et al. (2018); and Foster, Haltiwanger, and Syverson (2016) for further evidence that demand shocks are much larger than productivity shocks.

the payoffs about which firms or households can potentially be mistaken. Second, our baseline specification of minimal information allows both households and firms to see their island's own productivity. This still allows agents to be uncertain with respect to productivity components (temporary versus persistent and idiosyncratic versus aggregate), but it rules out expectation errors regarding firms' own *contemporaneous* productivity.

We contrast this "homogenous information" baseline with a specification in which households and firms do not share information. In this case, household uncertainty about local productivity can drive somewhat larger fluctuations. Still, the fluctuations that can be supported by uncertainty about productivity in this case are not nearly as large as those that can be generated by uncertainty regarding local demand. Across the cases we investigate, local demand uncertainty remains the most important prerequisite for large information-driven fluctuations.

Finally, we explore the degree to which confidence-driven fluctuations are consistent with US business cycle data. To this end we estimate a prototype wedge economy similar to the one in Chari, Kehoe, and McGrattan (2007), which captures the auto-covariance structure of the US business cycle by construction. We then use our theoretical results to partition the estimated wedges into an informational component, which can be microfounded through incomplete information, and a noninformational residual. We find that, in principle, confidence fluctuations can account for a large portion of the US business cycle that remains unexplained after conditioning on productivity shocks.

Again, a prerequisite for such confidence fluctuations to be sizable is that firms do not know their idiosyncratic product demands while making their production plans: If local demand is perfectly observed, at most 3 percent of observed output fluctuations can be accounted for by any type of confidence (regardless of what else firms observe). By contrast, if local demand is not observed but aggregate productivity is, up to 51 percent of output fluctuations can be explained by correlated confidence regarding local conditions, leading us to conclude that local demand shocks are crucial for the model to support aggregate sentiment fluctuations.

Related Literature.—The methodology developed in this paper is related to Bergemann and Morris (2013, 2016) and Bergemann, Heumann, and Morris (2015). These papers demonstrate the equivalence between *Bayes equilibria* in games with incomplete information and *Bayes correlated equilibria*. The approach developed in this paper is similar in that it also demonstrates the equivalence between a class of incomplete-information models with another class of full-information models. Our approach is significantly more general, however, because it is not limited to static game environments, but also applies to dynamic market economies, which is crucial for the application to business cycles. Closely related to our application to dynamic macroeconomic models, Passadore and Xandri (2021) develop robust predictions in dynamic policy games with an application to sovereign debt.

On the applied side, our analysis relates to a recent literature on confidence-driven business cycles. While the literature is mostly theoretical, there are now a few studies with a quantitative focus. In particular, Huo and Takayama (2015) quantify a version of Angeletos and La'O (2013); and Blanchard, L'Huillier, and Lorenzoni

(2013) estimate a version of Lorenzoni (2009).[2] Our approach is distinguished by our general formulation of incomplete information that does not require an ex ante stand on which agents are affected by information frictions, how information is shared in the cross section of agents, or any other parametric properties of the information structure.

The objective of this paper is also closely related to Angeletos, Collard, and Dellas (2018). Departing from the assumption of rational expectations, those authors develop a tractable framework in which agents' expectations regarding the beliefs of other agents are subject to reduced-form "confidence shocks." They show that confidence shocks can account for a significant portion of the US business cycle, but abstract from the question of whether those shocks can by microfounded by some information structure. Our approach is complimentary in that we characterize the restrictions on confidence-driven fluctuations imposed by rational expectations.

Our approach is also useful for reducing the computational burden of solving (and estimating) business cycle models with incomplete information. While the incomplete-information version of our economy is hard to solve, the corresponding primal economy permits a simple aggregate representation, in which aggregate wedges capture the average deviations from incomplete information in the cross section of agents. Conditional on these wedges, which are constrained by the restrictions characterized in our theorem, the primal economy can be solved using standard tools developed for full-information models. In this ability to reduce the computational burden of solving (and estimating) incomplete information models, our paper also relates to Rondina and Walker (2021); Acharya (2013); Huo and Takayama (2021); Acharya, Benhabib, and Huo (2018); Han, Tan, and Wu (2021); and Adams (2022), who use frequency-domain techniques to obtain analytical solutions in certain models, and Nimark (2017) who explores the asymptotic accuracy of a finite-state approximation approach to a class of dispersed information models.

Layout.—The rest of the paper is organized as follows. Section I develops our information-robust characterization approach and applies it to the simple price-setting model. Section II sets up the quantitative model. Sections III derives information-robust predictions for the quantitative model. Section IV contains the application to US business cycles. Section V concludes. Replication data are available in Chahrour and Ulbricht (2023).

I. Information-Robust Characterization

We present our main result in the context of a general linear rational expectations model with incomplete information. The framework encompasses virtually all linearized DSGE models used in the literature as well as the class of coordination games studied by Morris and Shin (2002) and others. After stating our main

[2] See also Melosi (2014, 2017) for an estimation of a variant of Woodford (2003); Maćkowiak and Wiederholt (2015) for plausible calibration of a particular DSGE model with rational inattention; and Ilut and Saijo (2021) for a quantitative DSGE model with time-varying ambiguity aversion. In these works information frictions alter the propagation of fundamental shocks (productivity, monetary), but there are no confidence-driven fluctuations.

characterization theorem, we demonstrate its application in a simple model of price setting. In the subsequent sections, we apply our methodology to a quantitative business cycle model, and use it to explore the potential importance of confidence-driven business cycles in the United States.

A. *Main Theorem*

Framework.—Consider a linear economy characterized by a system of expectational difference equations, in which date-t expectations are formed conditional on a collection of information sets $\left\{\mathcal{I}_{i,t}^{j}\right\}$. Here, $j \in \{0,1,\ldots,J\}$ indexes a collection of ex ante heterogeneous information classes that may differ arbitrarily. Within each class j, there is a continuum of ex ante symmetric information sets, indexed by $i \in [0,1]$, which may only differ in their ex post realization of shocks.[3] We normalize $j = 0$ to refer to the full information set, $\mathcal{I}_t^*$, which is defined by the history of all variables that are realized at date t.[4]

Let $g_{i,t} \equiv [\Delta g_{i,t}; g_t^a]$, where $\Delta g_{i,t}$ denotes a $n_{\Delta g} \times 1$ vector of idiosyncratic endogenous variables that satisfy the adding-up constraint $\int_0^1 \Delta g_{i,t} \mathrm{d}i = 0$, and g_t^a denotes a $n_{g^a} \times 1$ vector of endogenous aggregate variables (which may but are not limited to include the "mean component" of $\{\Delta g_{i,t}\}$).

We suppose that $g_{i,t}$ satisfies the following system of expectational difference equations:

$$0 = \sum_{j=0}^{J} \mathbb{E}\left\{ \begin{bmatrix}\mathbf{A}_1^j & \mathbf{A}_2^j\end{bmatrix} \begin{bmatrix} g_{i,t+1} \\ f_{i,t+1} \end{bmatrix} + \begin{bmatrix}\mathbf{B}_1^j & \mathbf{B}_2^j\end{bmatrix} \begin{bmatrix} g_{i,t} \\ f_{i,t} \end{bmatrix} | \mathcal{I}_{i,t}^j \right\}, \tag{1}$$

for all $i \in [0,1]$ and $t = 0,1,\ldots$. Here, $f_{i,t} \equiv [\Delta f_{i,t}; f_t^a]$ is a column vector of exogenous stochastic variables. In analogy to the endogenous vector $g_{i,t}$, we partition the exogenous vector into an idiosyncratic component, $\Delta f_{i,t}$, and an aggregate component, f_t^a, where the idioryncratic component satisfies the adding up constraint $\int_0^1 \Delta f_{i,t} \mathrm{d}i = 0$. We assume that $f_{i,t}$ follows a stationary Gaussian process and is ex ante symmetric across i.[5]

Throughout, we maintain the assumption of rational expectations, so that conditional on an information set, all expectations are formed using Bayes' law. An equilibrium is defined as a joint process for all the endogenous variables, $\{\Delta g_{i,t}\}_{i \in [0,1]} \cup g_t^a$, that solves (1) given processes for the exogenous fundamentals $\mathcal{F}_t \equiv \{\Delta f_{i,t}\}_{i\in[0,1]} \cup f_t^a$ and for information $\mathcal{I}_t \equiv \left\{\mathcal{I}_{i,t}^j\right\}_{i,j \in [0,1]\times\{1,2,\ldots,J\}}$. We use $\mathcal{E}(\mathcal{F},\mathcal{I})$ to denote the set of stationary equilibria satisfying (1). We note that nothing stated here requires equilibrium to be unique or even to exist.

[3] Here, ex ante symmetry across i means that the unconditional distribution over $\mathcal{I}_{i,t}^j$ is identical across all i. While differences in the ex post realization of signals can also be captured by introducing additional information classes, using i to reflect these differences helps streamlining notation in models where (some) agents are ex ante identical.

[4] Notice which variables are realized at date t is definitional and, thus, something the modeler must specify. For instance, $\mathcal{I}_t^*$ could contain future innovations if they are realized at date t as in the news literature.

[5] These assumptions can be relaxed. First, in many cases, $f_{i,t}$ can be detrended along with an appropriate transformation of (1). Second, while we assume $f_{i,t}$ to be Gaussian, the assumption is not needed when one is only interested in characterizing the auto-covariance structure of $g_{i,t}$. Third, symmetry across i is without loss of generality, as one can stack an arbitrary number of shocks into $f_{i,t}$.

Primal Representation.—Our main result constitutes an isomorphism between the equilibria of the model (1) and the equilibria of a related full-information economy, which we call the "primal" representation of the model. The primal representation of model (1) is given by

$$(2)\qquad 0 = \left(\sum_{j=0}^{J}\left[\mathbf{A}_1^j \;\; \mathbf{A}_2^j\right]\right)\begin{bmatrix}\mathbb{E}_t g_{i,t+1}\\ \mathbb{E}_t f_{i,t+1}\end{bmatrix} + \left(\sum_{j=0}^{J}\left[\mathbf{B}_1^j \;\; \mathbf{B}_2^j\right]\right)\begin{bmatrix}g_{i,t}\\ f_{i,t}\end{bmatrix} + \sum_{j=1}^{J}\tau_{i,t}^j,$$

where $\mathbb{E}_t\left[\cdot\right] \equiv \mathbb{E}\left[\cdot\,|\mathcal{I}_t^*\right]$ denotes the full-information expectation operator. Compared to (1), model (2) replaces all expectation operators $\mathbb{E}\left[\cdot\,|\mathcal{I}_{i,t}^j\right]$ with $\mathbb{E}_t\left[\cdot\right] + \tau_{i,t}^j$, where $\left\{\tau_{i,t}^j\right\}$ represent the expectation errors implicit in agents' equilibrium expectations relative to full information. Notice that our notation already reflects the normalization that $j = 0$ corresponds to full information by setting $\tau_{i,t}^0 = 0$.

The key conceptional difference between the primal economy and the original one is that in the primal economy we treat agents' expectation errors as exogenous "wedges," whereas in the original economy they derive endogenously from agents' information sets. In analog to the original economy, we use $\mathcal{E}^{\text{primal}}\left(\mathcal{F},\mathcal{T}\right)$ to denote the set of stationary equilibria of the primal economy with fundamentals $\mathcal{F}$ and expectation wedges $\mathcal{T}_t \equiv \left\{\tau_{i,t}^j\right\}_{i,j\in[0,1]\times\{1,2,\ldots,J\}}$. Solving models of the form in (2) is straightforward, and the literature offers myriad strategies for obtaining $\mathcal{E}^{\text{primal}}\left(\mathcal{F},\mathcal{T}\right)$.

Characterization Theorem.—We now state our main theorem, which provides necessary and sufficient conditions on the expectation wedges in the primal representation such that they can be supported as expectation errors in an equilibrium of the original incomplete information economy.

To do so, we impose the following structure on information in the original economy.

ASSUMPTION 1 (Information Bounds): $\Theta_{i,t}^j \subseteq \mathcal{I}_{i,t}^j \subseteq \mathcal{I}_t^*$.

Assumption 1 defines a lower and an upper bound on information. The upper bound, $\mathcal{I}_t^*$, simply states that agents cannot learn more than what is potentially knowable under full information. The lower bound, $\Theta_{i,t}^j$, must be specified by the modeler. It constitutes the primary input parameter to our methodology, allowing researchers to explore how their priors regarding agents' information restrict equilibrium outcomes.

ASSUMPTION 2 (Recursiveness): $\mathcal{I}_{i,t-1} \subseteq \mathcal{I}_{i,t}$.

Assumption 2 imposes the usual rationality requirement that all agents perfectly recall past information. While perfect recall is standard, we note that our methodology easily extends to the case where agents may forget past information.[6]

[6] Specifically, in the case of no recall, we obtain a version of our theorem, in which condition (3) is imposed only for $s = 0$.

To state the theorem, define

$$\mu_{i,t}^{j} \equiv \mathbb{E}_t\left[\mathbf{A}_1^j g_{i,t+1} + \mathbf{A}_2^j f_{i,t+1} + \mathbf{B}_1^j g_{i,t} + \mathbf{B}_2^j f_{i,t}\right] + \tau_{i,t}^{j},$$

which for each (i,j,t) represents the expectation implicit in $\tau_{i,t}^{j}$. The following theorem states the implementation result.

THEOREM 1: *Fix stationary $\mathcal{F}$, $\mathcal{T}$, and $\mathcal{E} \in \mathcal{E}^{\text{primal}}(\mathcal{F},\mathcal{T})$. Then there exists an information structure $\mathcal{I}$ satisfying Assumptions 1 and 2 that implements $\mathcal{E}$ as equilibrium in the incomplete-information economy (i.e., $\mathcal{E} \in \mathcal{E}(\mathcal{F},\mathcal{I})$) if and only if (i) $\mathcal{E}\left[\tau_{i,t}^{j}\right] = 0$ and (ii)*

$$\mathbb{E}\left[\tau_{i,t}^{j}\theta\right] = 0 \textit{ for all } \theta \in \left\{\mu_{i,t-s}^{j}, \Theta_{i,t-s}^{j}\right\}_{s\geq 0} \tag{3}$$

hold for i, j, and t.

The theorem gives two conditions that are jointly necessary and sufficient for $\mathcal{T}$ to be implemented by some information structure. Condition (i) is simply a rationality requirement that an agent's beliefs cannot be perpetually biased. Condition (ii) is an orthogonality requirement between the expectation wedges and $\mu_{i,t}^{j}$ and $\Theta_{i,t}^{j}$. The necessity of this restriction is the familiar result that expectation errors must be orthogonal to all available information, including an agent's belief $\mu_{i,t}^{j}$ itself (at the very least "one knows what one knows"). The novel part of our result is the sufficiency of this condition. For any $\mathcal{E} \in \mathcal{E}^{\text{primal}}(\mathcal{F},\mathcal{T})$ with $\mathbb{E}\left[\mathcal{T}_t\right] = 0$, we can always construct an information structure that implements $\mathcal{E}$ as an incomplete-information equilibrium as long as it satisfies (3).

Sketch of Proof.—Here we illustrate the proof in a simple case. The general proof is given in Appendix Section A. Suppose equilibrium in the original economy is defined by a single condition,

$$y_t = \mathbb{E}\left[a_t | \mathcal{I}_t\right], \tag{4}$$

where $\mathbb{E}\left[a_t\right] = 0$, and let $\Theta_t = \emptyset$. Equilibrium in the primal economy is then defined by

$$y_t = a_t + \tau_t. \tag{5}$$

Let (y_t, a_t, τ_t) be a stationary Gaussian process satisfying (5). Theorem 1 states that there exists an $\mathcal{I}_t$ that supports y_t as an equilibrium in the original economy if and only if (i) $\mathbb{E}\left[\tau_t\right] = 0$ and (ii) $\mathbb{E}\left[\tau_t y_{t-s}\right] = 0$ for all $s \geq 0$. The necessity of conditions (i) and (ii) is immediate because optimal inference requires that expectation errors are unpredictable.

To see why the conditions are also sufficient, first note that by construction (y_t, a_t, τ_t) is an equilibrium in the primal economy. For (y_t, a_t, τ_t) to also solve (4), it, hence, suffices to construct an $\mathcal{I}$ such that $\mathbb{E}[a_t|\mathcal{I}_t] = a_t + \tau_t = y_t$. To do so, suppose that $\mathcal{I}_t = \{\omega_{t-s}\}_{s\geq 0}$, where $\omega_t = a_t + \tau_t$. That is, each period, the agent receives a new signal ω_t that has the same joint distribution over (ω_t, a_t) as the equilibrium "belief" y_t that we wish to implement. Projecting a_t onto $y^t \equiv \{y_{t-s}\}_{s\geq 0}$, we have

$$\mathbb{E}[a_t|\mathcal{I}_t] = \operatorname{cov}(a_t, y^t)\left[\operatorname{var}(y^t)\right]^{-1} y^t. \tag{6}$$

Notice that

$$\operatorname{cov}(y_t, y^t) = [1 \ 0 \ 0 \ \cdots]\operatorname{var}(y^t). \tag{7}$$

Further notice that (5) in combination with condition (ii) gives $\operatorname{cov}(a_t, y^t) = \operatorname{cov}(y_t - \tau_t, y^t) = \operatorname{cov}(y_t, y^t)$. We can thus use (7) to substitute out $\operatorname{cov}(a_t, y^t)$ in (6) to get

$$\mathbb{E}[a_t|\mathcal{I}_t] = [1 \ 0 \ 0 \ \cdots]\operatorname{var}(y^t)\left[\operatorname{var}(y^t)\right]^{-1} y^t = y_t.$$

We conclude that as long as conditions (i) and (ii) hold, there exists an information structure that implements τ_t and, hence, y_t. Intuitively, observing the equilibrium expectation y_t is a sufficient statistic for forming $\mathbb{E}[a_t \,|\, \mathcal{I}_t]$, giving us a simple means of implementing τ_t.

B. *Illustration: Application to Price-Setting Model*

As an example of how our approach works in practice, we present a simple price setting model and show how to derive analytical restrictions on equilibrium outcomes. The model focuses on the log-linearly approximated pricing decision of a monopolistically competitive firm, while taking aggregate demand as an exogenous process in the spirit of Woodford (2003).

Setup.—Firms in the model set their prices according to

$$p_{i,t} = \mathbb{E}[p_t + \xi y_t + \nu z_{i,t}|\mathcal{I}_{i,t}], \tag{8}$$

where $p_t \equiv \int_0^1 p_{i,t}di$ is the aggregate price index, y_t is aggregate output, $z_{i,t}$ is an idiosyncratic demand shock, and $\xi \in (0,1)$ and $\nu \in [0,1]$ are the elasticities of the target price in y_t and $z_{i,t}$. Each firm i, faces standard CES demand,

$$y_{i,t} = -\eta(p_{i,t} - p_t) + y_t + \eta z_{i,t}, \tag{9}$$

with $\eta > 1$. Finally, aggregate output and prices are related via the constant-velocity equation

$$q_t = y_t + p_t, \tag{10}$$

with q_t denoting the exogenous supply of money. We assume that $\{z_{i,t}\}$ and q_t follow independent stationary Gaussian processes, and $\int_0^1 z_{i,t}\, \mathrm{d}i = 0$.

Primal Representation.—Because only (8) contains an expectation, it is the only equation with a nontrivial expectation wedge in the primal representation of the economy. The primal representation of the economy is therefore given by

$$p_{i,t} = p_t + \xi y_t + \nu z_{i,t} + \tau_{i,t}, \tag{11}$$

along with equations (9) and (10).

Given a process $\{\tau_{i,t}\}$, the equilibrium of the primal economy is straightforward to find. Defining $\tau_t \equiv \int_0^1 \tau_{i,t}\, \mathrm{d}i$, aggregates in the economy are given by

$$p_t = q_t + \xi^{-1}\tau_t, \ y_t = -\xi^{-1}\tau_t. \tag{12}$$

Similarly, we can solve for the idiosyncratic dynamics of $\Delta p_{i,t} \equiv p_{i,t} - p_t$ and $\Delta y_{i,t} \equiv y_{i,t} - y_t$ to arrive at

$$\Delta p_{i,t} = \nu z_{i,t} + \Delta \tau_{i,t}, \ \Delta y_{i,t} = \eta(1-\nu) z_{i,t} - \eta \Delta \tau_{it}. \tag{13}$$

Notice that the equilibrium in the primal representation provides a separation of dynamics at the aggregate and idiosyncratic levels. A similar separation is always possible with appropriate definitions, and turns about to be convenient for deriving restriction on equilibrium outcomes.

Predictions.—Applying our theorem amounts to placing covariance restrictions on the outcomes captured by (12)–(13). For the purpose of this illustration, we focus on the case where firms observe their own sales; i.e., $y_{i,t} \in \Theta_{i,t}$, ruling out any information structures where $y_{i,t} \notin \mathcal{I}_{i,t}$. While this may not be entirely realistic, it provides for an instructive example to demonstrate how our methodology can be applied in practice. We note that under these assumptions, it is equivalent to assume that firms set prices or quantities (as we later assume in our quantitative exercises.)

Observe that in the notation of our general framework, $\mu_{i,t} = p_{i,t}$. Theorem 1, hence, requires $\tau_{i,t}$ to be orthogonal to $y_{i,t-s}$ and $p_{i,t-s}$ for all $s \geq 0$. Imposing these restrictions, we arrive at two key implementability conditions relating the dynamics of aggregate and idiosyncratic variables:

$$\mathrm{cov}[\tau_t, p_{t-s}] = -\mathrm{cov}[\Delta \tau_{i,t}, \Delta p_{i,t-s}] \tag{14}$$

$$\mathrm{cov}[\tau_t, y_{t-s}] = -\mathrm{cov}[\Delta \tau_{i,t}, \Delta y_{i,t-s}] \tag{15}$$

for all $s \geq 0$. Manipulating these conditions allows us to derive a series of results.

PROPOSITION 1: *The unconditional variance of output is bounded by the volatility of q_t according to*

$$\sqrt{\text{var}[y_t]} \leq \frac{(1-\nu)\eta}{(1-\nu)\eta+\nu}\sqrt{\text{var}[q_t]}.$$

PROOF:

Using (13) to substitute out $\Delta p_{i,t-s}$ and $\Delta y_{i,t-s}$ in (14) and (15), and combining conditions to eliminate $\text{cov}[\Delta \tau_{i,t}, z_{i,t-s}]$, we have

$$\nu \text{cov}[\tau_t, y_{t-s}] - \eta(1-\nu)\text{cov}[\tau_t, p_{t-s}] = \eta \text{cov}[\Delta \tau_{i,t}, \Delta \tau_{i,t-s}]. \tag{16}$$

Evaluating at $s = 0$ and using (12) to substitute out τ_t and p_t,

$$(\nu+\eta(1-\nu))\text{var}[y_t] - \eta(1-\nu)\text{cov}[y_t, q_t] = -\eta \xi^{-1}\text{var}[\Delta \tau_{i,t}]. \tag{17}$$

Noting that $\text{var}[\Delta \tau_{i,t}] \geq 0$ and, by the Cauchy-Schwartz inequality, $\text{cov}[y_t, q_t] \leq \sqrt{\text{var}[y_t]} \cdot \sqrt{\text{var}[q_t]}$, completes the proof. ∎

The proposition expresses a bound on the volatility of aggregate output relative to the volatility of nominal demand. The bound is especially stark in the simple model, necessitating an exogenous aggregate shock to generate any expectation-driven fluctuations in aggregate output. As we explore in our more general quantitative setting, this conclusion is an artifact of two simplifying assumptions: (i) the assumption that firms observe their own sales, $y_{i,t} \in \Theta_{i,t}$, which precludes firms from having uncertainty about their demand, and (ii) the absence of other firm-specific shocks affecting input prices or technology. Once we relax either of these assumptions, it will be possible to generate expectation-driven fluctuations in the absence of aggregate shocks. Before further exploring this possibility, we first demonstrate how one can use our methodology to establish related bounds on the comovement between output, inflation, and money growth.

PROPOSITION 2: *Inflation $\pi_t \equiv p_t - p_{t-1}$ and money growth $dq_t = q_t - q_{t-1}$ must be weakly procyclical. Specifically, the correlation with output is bounded below as follows:*

$$\nu\sqrt{\text{var}[y_t]} \leq (1-\nu)\eta \cdot \frac{\text{corr}[y_t, \pi_t]}{1-\text{corr}[y_t, y_{t-1}]}\sqrt{\text{var}[\pi_t]}$$

and

$$\sqrt{\text{var}[y_t]} \leq \frac{(1-\nu)\eta}{(1-\nu)\eta+\nu} \cdot \frac{\text{corr}[y_t, dq_t]}{1-\text{corr}[y_t, y_{t-1}]}\sqrt{\text{var}[dq_t]}.$$

PROOF:

As both bounds are derived following completely analogous steps, we only show the proof for inflation. Evaluating (16) for $s = 0$ and $s = 1$, using (12) to substitute for τ_t, and differencing the resulting conditions, we have

$$\nu \text{cov}[y_t, \text{d}y_t] - \eta(1-\nu)\text{cov}[y_t, \pi] = -\eta\xi^{-1}\text{cov}[\Delta\tau_{i,t}, \text{d}\Delta\tau_{i,t}].$$

Noting that $\text{cov}[\Delta\tau_{i,t}, \text{d}\Delta\tau_{i,t}] = \left(1 - \text{corr}[\Delta\tau_{i,t}, \Delta\tau_{i,t-1}]\right) \cdot \text{var}[\Delta\tau_{i,t}] \geq 0$ completes the proof. ∎

The proposition establishes that, when uncertainty originates exclusively from demand shocks, expectations-driven fluctuations must exhibit exactly the same cyclical properties as demand shocks themselves. Again, the restriction is especially stark given the assumptions of our simple model, and the restriction that inflation and money growth must be procyclical is relaxed once we allow for other sources of uncertainty.

We conclude our illustration by exploring two refinements of $\Theta_{i,t}$.

PROPOSITION 3: *Suppose* $\{z_{i,t}, y_{i,t}\} \in \Theta_{i,t}$. *Then aggregate output is constant.*

PROOF:

Using (9) to substitute out $\Delta y_{i,t}$ in (15), and combining with (14) to eliminate $\Delta p_{i,t}$, we obtain

$$\text{cov}[\tau_t, y_{t-s} + \eta p_{t-s}] = -\eta \text{cov}[\Delta\tau_{i,t}, z_{i,t-s}]. \tag{18}$$

From $\int z_{i,t}\text{d}i = 0$, it follows that $\text{cov}[\Delta\tau_{i,t}, z_{i,t-s}] = \text{cov}[\tau_{i,t}, z_{i,t-s}]$. Applying Theorem 1, it then must hold that $\text{cov}[\Delta\tau_{i,t}, z_{i,t-s}] = 0$. Evaluating (18) at $s = 0$ and $s = 1$, using (12) to substitute for τ_t, and differencing the resulting conditions, we therefore obtain

$$\sqrt{\text{var}[y_t]} = -\eta\frac{\text{corr}[y_t, \pi_t]}{1 - \text{corr}[y_t, y_{t-1}]}\sqrt{\text{var}[\pi_t]}.$$

The result then follows, because $\text{corr}[y_t, \pi_t] \geq 0$ by Proposition 2. ∎

The proposition complements the finding in Proposition 1 that expectation-driven fluctuations can only be caused by uncertainty about aggregate demand. Proposition 3 further demonstrates that even though uncertainty about $z_{i,t}$ cannot support any systematic aggregate fluctuations (when $y_{i,t} \in \Theta_{i,t}$), it is nevertheless necessary for supporting aggregate fluctuations caused by uncertainty about q_t. This is intuitive because, without uncertainty about $z_{i,t}$, firms can simply invert their idiosyncratic demand to back out the aggregate state of the economy, resolving any uncertainty about q_t.

Finally, we highlight a natural bound on the autocorrelation of endogenous fluctuations when information is revealed with some lag.

PROPOSITION 4: *Suppose that* $\{y_{t-h}, \pi_{t-\bar{h}}\} \in \Theta_{i,t}$. *Then* $\operatorname{cov}[y_t, y_{t-s}] = \operatorname{cov}[y_t, \pi_{t-s}] = \operatorname{cov}[y_t, q_{t-s}] = 0$ *for all* $s \geq \bar{h}$.

PROOF:

The result follows immediately from Theorem 1, after using (12) to substitute for $\tau_t = -\xi y_t$. ■

The proposition establishes that if the aggregate state is revealed at some lag $\bar{h}$, then this limits the autocorrelation of any expectation-driven fluctuations to within a horizon of $\bar{h}$ periods. The result echoes a similar, but more special, result in Acharya, Benhabib, and Huo (2018), which bounds the persistence of a specific type of sentiment shocks. Proposition 4, by contrast, reveals that the lagged revelation of aggregate information always eliminates subsequent autocovariances, regardless of other details of the information structure.

II. Quantitative Application

A. *Setup*

We now turn to our quantitative application. The model is a "RBC economy without capital," augmented with imperfect information. Households and firms are located on a continuum of islands, indexed by $i \in [0,1]$. On each island, a representative household interacts with a representative firm in a local labor market. Firms use the labor provided by households to produce differentiated intermediate goods, which are aggregated by a competitive final goods sector located on the mainland. There are no subperiods; all markets at date t operate simultaneously.

Households.—Preferences on island i are given by

$$\mathbb{E}\left\{\sum_{\tau=0}^{\infty} \beta^{\tau} U(C_{i,t+\tau}, N_{i,t+\tau}) | \mathcal{I}_{i,t}^{h}\right\},$$

where $\beta \in (0,1)$ is the discount factor, $N_{i,t}$ is hours worked, $C_{i,t}$ is final good consumption, and $\mathcal{I}_{i,t}^{h}$ denotes the information available to the household on island i at time t. The utility flow U is given by

$$U(C,N) = \log C - \frac{1}{1+\zeta} N^{1+\zeta},$$

where $\zeta \geq 0$ is the inverse of the Frisch elasticity of labor supply. The household's budget constraint is

$$P_t C_{i,t} + Q_t B_{i,t} \leq W_{i,t} N_{i,t} + B_{i,t-1} + D_{i,t},$$

where P_t is the price of the final good, Q_t is the nominal price of a riskless one-period bond, $B_{i,t}$ are local bond holdings, $W_{i,t}$ are local wage rates, and $D_{i,t}$ are profits of the local firm.[7] Bonds are in zero net supply, so market clearing requires $\int_0^1 B_{i,t}\,\mathrm{d}i = 0$. No other financial assets can be traded across islands, which implies that households are exposed to idiosyncratic income risk.

Intermediate-Goods Producers.—Each good i is produced by a monopolistically competitive firm with access to a linear production technology,

$$Y_{i,t} = A_{i,t} N_{i,t}. \tag{19}$$

Firms choose $N_{i,t}$ to maximize expected profits, $\mathbb{E}\left[P_{i,t} Y_{i,t} - W_{i,t} N_{i,t} | \mathcal{I}^f_{i,t}\right]$, subject to an inverse demand curve specified below. Here, $\mathcal{I}^f_{i,t}$ denotes the date-t information available to the firm on island i, which may differ from households' information. The wage rate $W_{i,t}$ is determined competitively.[8] The productivity $A_{i,t}$ consists of an aggregate and an island-specific component,

$$\log A_{i,t} = \log A_t + \Delta a_{i,t},$$

where the aggregate component follows a random walk process

$$\log A_t = \log A_{t-1} + \epsilon_t.$$

The innovation ϵ_t is i.i.d. across time with zero mean and constant variance. The island-specific component $\Delta a_{i,t}$ follows a time-invariant, stationary random process that is i.i.d. across islands and is normalized so that $\int_0^1 \Delta a_{i,t}\,\mathrm{d}i = 0$.

Final-Good Sector.—A competitive final-goods sector aggregates intermediate input goods $i \in [0,1]$, using the technology

$$Y_t = \left(\int_0^1 Z_{i,t} Y_{i,t}^{\frac{\eta-1}{\eta}} di\right)^{\frac{\eta}{\eta-1}},$$

where $\eta > 1$ is the elasticity of substitution among input varieties, $Y_{i,t}$ denotes the input of intermediate good i at time t, and $Z_{i,t}$ is an island-specific demand shifter following a time-invariant, stationary process that is i.i.d. across islands and satisfies $\int_0^1 \log(Z_{i,t})\, di = 0$. Profit maximization yields the inverse input demands, given by

$$P_{i,t} = \left(\frac{Y_{i,t}}{Y_t}\right)^{-1/\eta} Z_{i,t} P_t, \tag{20}$$

[7] Following Maćkowiak and Wiederholt (2015), we assume that bond positions adjust to clear the budget constraint independently of the information available to households.

[8] Formally, firm i is representative of a continuum of firms, $j \in [0,1]$, competing in the local labor market. Each of these firms produces a separate variety (i,j) that is aggregated to $Y_{i,t}$ using the technology $Y_{i,t} = \left(\int_0^1 Y_{ij,t}^{1-1/\eta}\mathrm{d}j\right)^{\eta/(\eta-1)}$, where η matches the elasticity of substitution across "island-varieties" specified in the final good technology below. Clearly, the setting collapses to the one in the main text, where $Y_{i,t}$ is produced by a representative firm i that is competitive in the local labor market and faces isoelastic demand from the final good sector with elasticity $-\eta$.

where the aggregate price index P_t is defined by

$$P_t = \left(\int_0^1 Z_{i,t}^{\eta} P_{i,t}^{1-\eta} \mathrm{d}i\right)^{\frac{1}{1-\eta}}.$$

Monetary Policy.—We close the model by specifying a simple interest rate rule, pinning down the equilibrium rate of inflation, $\pi_t \equiv \log(P_t/P_{t-1})$. Specifically, we assume that the central bank sets nominal bond prices such that

$$i_t = \phi \pi_t, \tag{21}$$

where $\phi > 1$ and $i_t = -\log(Q_t)$.[9]

Information Structure.—Our methodology allows us to explore how a few abstract assumptions regarding $\{\mathcal{I}_{i,t}^j\}_{i,j \in [0,1] \times \{f,h\}}$ restrict equilibrium behavior, without the need to fully specify a parametric information structure. As a baseline, we consider the case where firms and households share the same information within islands ($\mathcal{I}_{i,t} \equiv \mathcal{I}_{i,t}^f = \mathcal{I}_{i,t}^h$) and where the joint information set $\mathcal{I}_{i,t}$ is bounded below by

$$\Theta_{i,t}^{\text{sym}} = \{A_{i,t}, C_{i,t}, N_{i,t}, Y_{i,t}, W_{i,t}, \mathcal{I}_{t-\bar{h}}^*\} \cup \Theta_{i,t-1}^{\text{sym}}. \tag{22}$$

Under this baseline, households and firms observe local output (and hence productivities) in addition to the local consumption, employment, and wages. Moreover, all agents eventually learn the truth at some horizon $\bar{h} \geq 0$.[10] The assumption of finite revelation is not required by our theorem, but is useful in our application because it ensures that observing a history of growth rates of a variable is equivalent to observing its level.

As an alternative to this baseline, we also explore the case in which firms and households have access to different information. In our most general (i.e., least restrictive) specification, information is bounded below by

$$\Theta_{i,t}^{h} = \{C_{i,t}, N_{i,t}, W_{i,t}, \mathcal{I}_{t-\bar{h}}^*\} \cup \Theta_{i,t-1}^{h} \tag{23}$$

$$\Theta_{i,t}^{f} = \{A_{i,t}, N_{i,t}, Y_{i,t}, W_{i,t}, \mathcal{I}_{t-\bar{h}}^*\} \cup \Theta_{i,t-1}^{f}. \tag{24}$$

Because different agent classes on a single island have different information under this specification, we refer to this case as one of "heterogenous information."

Throughout, we assume that the full information set contains any variables dated t or earlier. Hence, we rule out "news," because future innovations to A_t, $\{\Delta a_{i,t}\}$ and $\{Z_{i,t}\}$ are not part of $\mathcal{I}_t^*$.

[9] The rule also contains a constant intercept ensuring consistency with the natural rate at the zero-inflation steady state. The term is omitted since it drops out after we log-linearize the model below.

[10] Here we specify $\Theta_{i,t}^{\text{sym}}$ recursively to emphasize that households have access to past information but note that doing so is redundant given Assumption 2. Similarly, because $\mu_{i,t}$ (in the notation of the general framework) is a monotone transformation of $C_{i,t}$ and $N_{i,t}$ and (together with $Y_{i,t}$) can be used to infer $A_{i,t}$ and $W_{i,t}$, one could without loss of generality omit $\{C_{i,t}, N_{i,t}, A_{i,t}, W_{i,t}\}$ from $\Theta_{i,t}^{\text{sym}}$. The smallest set $\Theta_{i,t}$ yielding identical equilibrium restrictions as $\Theta_{i,t}^{\text{sym}}$ is therefore $\Theta_{i,t} = \{Y_{i,t}, \mathcal{I}_{t-\bar{h}}^*\}$.

B. *Equilibrium Conditions*

We work with a log-linear approximation to the model around the balanced growth path of the economy with no heterogeneity and full information. Lowercase letters denote log-deviations of a variable from this path, in which $y_{i,t} = a_t$ for all i and $\pi_t = 0$.

The households' Euler equation is given by

$$c_{i,t} = \mathbb{E}\left[c_{i,t+1} - \phi \pi_t + \pi_{t+1} \,|\, \mathcal{I}^h_{i,t}\right]. \tag{25}$$

Combining firms' demand for labor with households' supply, local labor market clearing requires

$$y_{i,t} = \xi\left(y_{i,t} - c_{i,t} + \mathbb{E}\left[p_{i,t} \,|\, \mathcal{I}^f_{i,t}\right] - \mathbb{E}\left[p_t \,|\, \mathcal{I}^h_{i,t}\right]\right) + a_{i,t}, \tag{26}$$

where $\xi \equiv 1/(\zeta + 1)$. The linearized price index p_t is given by $p_t = \int_0^1 p_{i,t} \mathrm{d}i$. The linearized demand relation and budget constraint take the form

$$p_{i,t} = \eta^{-1}(y_t - y_{i,t}) + z_{i,t} + p_t \tag{27}$$

and

$$\beta b_{i,t} = b_{i,t-1} + y_{i,t} - c_{i,t} + p_{i,t} - p_t, \tag{28}$$

where $b_{i,t} \equiv B_{i,t}/(P_t C_{i,t})$ is in levels rather than logs because $B_{i,t}$ can take negative values. Given a process for fundamentals and information $\{a_{i,t}, z_{i,t}, \mathcal{I}^f_{i,t}, \mathcal{I}^h_{i,t}\}$, an equilibrium of the model is a set of processes $\{c_{i,t}, y_{i,t}, b_{i,t}, p_{i,t}\}$ and $\{y_t, \pi_t\}$ that are consistent with (25)–(28), with Bayesian updating, and with market clearing for goods,

$$y_t = \int_0^1 y_{i,t} \mathrm{d}i = \int_0^1 c_{i,t} \mathrm{d}i. \tag{29}$$

(As usual, market clearing for bonds is implied by (28) and (29).)

Comment on Prices, Information, and Market Clearing.—In many general equilibrium models with incomplete information, it is relatively simple for agents to infer the value of the economy's aggregate fundamentals from observing aggregate prices. As argued by Lorenzoni (2009), this is largely an artifact of the simplicity of models, whereas, in practice, the ability of agents to learn about the economy's fundamentals is likely impaired by a large number of shocks, model misspecification, and the possible presence of structural breaks. To capture these effects within simple models like ours, the literature has therefore utilized various ways of introducing noise into price systems.[11]

[11] Common approaches include the addition of noise traders (e.g., Grossman and Stiglitz 1980; Hellwig 1980), the decentralization of markets (e.g., Lorenzoni 2009; Angeletos and La'O 2013), and the use of rational inattention that introduces noise directly into information sets (e.g., Maćkowiak and Wiederholt 2015; Mondria, Vives, and Yang 2022).

In keeping with the literature, we do not include the real return on assets, $r_t \equiv i_t - \mathbb{E}_t[\pi_{t+1}]$, or its constituents i_t, p_t, and $\mathbb{E}_t[p_{t+1}]$, in the lower bound on households' information $\{\mathcal{I}^h_{i,t}\}$. However, we note that by imposing market clearing on the aggregate goods market, we *implicitly* require that households observe some noisy version of r_t so that the average expected real interest, $\bar{\mathbb{E}}^h_t[r_t] \equiv \int \mathbb{E}[r_t | \mathcal{I}^h_{i,t}] \mathrm{d}i$, increases with r_t. As long as this is the case, market clearing is guaranteed by some interest rate r_t (which generally differs from the one that would emerge under full information). Using our methodology, there is no need to explicitly specify the signals through which households make inference about r_t. Simply imposing market clearing in the primal economy entails that $\bar{\mathbb{E}}^h_t[r_t]$ indeed has nonzero elasticity with respect to r_t, ensuring that the goods market clears in all states of the world.[12]

To see this, consider the simplified case where aggregate demand is given by $c_t = -\bar{\mathbb{E}}^h_t[r_t]$ and aggregate supply, y_t, follows an exogenous random process. In this case market clearing ($c_t = y_t$) requires

$$\bar{\mathbb{E}}^h_t[r_t] = -y_t, \tag{30}$$

which in conjunction with $\{\mathcal{I}^h_{i,t}\}$ pins down r_t. In the primal representation, we have $\bar{\mathbb{E}}^h_t[r_t] = r_t + \tau_t$ with $\tau_t = \int \tau_{i,t} \mathrm{d}i$, and market clearing requires

$$r_t + \tau_t = -y_t. \tag{31}$$

The key difference between (30) and (31) is that the average expectation error, τ_t, is a primitive of the primal economy. Accordingly, we can always ensure market clearing by setting $r_t = -y_t - \tau_t$, which entails that the equilibrium expectation $\bar{\mathbb{E}}^h_t[r_t] = r_t + \tau_t$ indeed adjusts to economic conditions. As further detailed in online Appendix Section A.1, this precisely rules out the case where households have no information at all about r_t (which entails $\tau_t = -r_t$). By imposing market clearing in the primal representation, we are effectively ruling out this no-information case without having to specify the underlying price signals parametrically.

C. *Primal Representation*

There are two equilibrium conditions with nontrivial expectation operators. Replacing equations (25) and (26) with their primal analog, we arrive at[13]

$$c_{i,t} = \mathbb{E}_t\left[\left(c_{i,t+1} - \tau^c_{i,t+1}\right) - \phi \pi_t + \pi_{t+1}\right] + \tau^c_{i,t} \tag{32}$$

$$y_{i,t} = \xi\left(y_{i,t} - c_{i,t} + p_{i,t} - p_t + \tau^{p,f}_{i,t} - \tau^{p,h}_{i,t}\right) + a_{i,t}. \tag{33}$$

[12] All other markets clear as usual, without the need for any information beyond the lower bounds given by (23)–(24): labor markets clear through $\{W_{i,t}\}$, which is observed by firms and households within each island; the market for input goods clears through $\{P_{i,t}\}$, which is observed by the final goods sector; and the market for bonds clears by Walras' law.

[13] Here $\tau^c_{i,t}$ is specified after rewriting (25) in its nonrecursive form. With this normalization, $\tau^c_{i,t}$ defines the gap relative to the optimal level of consumption that household i would choose if it had full information at t and all future dates.

Here, $\tau^c_{i,t}$ and $\tau^{p,h}_{i,t}$ have the interpretation of households' prediction errors, relative to full information, regarding their consumption target and the aggregate price index. On the firms' side, $\tau^{p,f}_{i,t}$ has the interpretation of firms' prediction error regarding their inverse product demand, $p_{i,t}$. Note that all wedges are defined relative to the full-information targets that obtain taking as given the behavior of the rest of the economy (given expectation errors made by other households and firms).

One unique feature of our environment is that we allow for non-stationarity in aggregate productivity, whereas most of the incomplete-information literature requires stationary fundamentals. Stationarizing the representation in (32)–(33) is straightforward, but invoking Theorem 1 requires us to find a stationary representation of $\left\{\mu^j_{i,t}, \Theta^j_{i,t}\right\}_{i,j\in[0,1]\times\{f,h\}}$ that contains the same information as the minimal information set in the original representation. A convenient way to do this is to assume that all past information is revealed at a finite horizon $\bar{h} \geq 0$ as in (22). In this case, we can replace any nonstationary sequences in $\left\{\mu^j_{i,t}, \Theta^j_{i,t}\right\}_{i,j\in[0,1]\times\{f,h\}}$ by their first-differences, which, in combination with $\mathcal{I}^*_{t-\bar{h}}$, contain the same information as the corresponding sequences in levels. The following assumption formalizes this requirement.

ASSUMPTION 3 (Finite Revelation): *For each* $(i,j) \in [0,1] \times \{f,h\}$, *there exists a stationary information process* $\mathcal{S}^j_{i,t}$ *such that* $\left\{\mu^j_{i,t-s}, \Theta^j_{i,t-s}\right\}_{s\geq 0}$ *is informationally equivalent to* $\left\{\mathcal{S}^j_{i,t-s}\right\}_{s=0}^{\bar{h}-1} \cup \mathcal{I}^*_{t-h}$ *for some* $\bar{h} \geq 0$.

Applying Theorem 1 then yields the following result.

PROPOSITION 5: *Fix stationary processes for*

$$\mathcal{F}_t \equiv \left\{\Delta a_{i,t}, z_{i,t}\right\}_{i\in[0,1]} \cup \left\{da_t\right\}$$

$$\mathcal{T}_t \equiv \left\{\tau^c_{i,t}, \tau^{p,h}_{i,t}, \tau^{p,f}_{i,t}\right\}_{i\in[0,1]},$$

and, using $d(\cdot)$ *to denote the first difference of a variable, fix*

$$\mathcal{E}_t \equiv \left\{dc_{i,t}, dy_{i,t}, db_{i,t}, dp_{i,t}\right\}_{i\in[0,1]} \cup \left\{dy_t, \pi_t\right\} \in \mathcal{E}(\mathcal{F}, \mathcal{T}).$$

Then there exists an information structure $\mathcal{I}$ *satisfying Assumptions 1–3 that implements* $\mathcal{E}$ *as equilibrium in the incomplete-information economy if and only if (i)* $\left(\tau^c_{i,t}, \tau^{p,h}_{i,t}, \tau^{p,f}_{i,t}\right)$ *follows a MA*(h) *process of order* $h < \bar{h}$, *(ii)* $\mathbb{E}\left[\left(\tau^c_{i,t}, \tau^{p,h}_{i,t}, \tau^{p,f}_{i,t}\right)\right] = 0$, *and (iii)*

$$\mathbb{E}\left[\tau^c_{i,t}\theta\right] = \mathbb{E}\left[\tau^{p,h}_{i,t}\theta\right] = 0 \quad \textit{for all} \quad \theta \in \left\{\mathcal{S}^h_{i,t-s}\right\}_{s=0}^{\bar{h}-1}, \textit{ and}$$

$$\mathbb{E}\left[\tau^{p,f}_{i,t}\theta\right] = 0 \quad \textit{for all} \quad \theta \in \left\{\mathcal{S}^f_{i,t-s}\right\}_{s=0}^{\bar{h}-1}$$

hold for all i and t.

Proposition 5 is an immediate corollary to Theorem 1. Here, the restriction to finite MA processes arises because $\mathcal{I}^*_{t-\bar{h}} \in \Theta^j_{i,t}$ under Assumption 3; because all innovations to $\left(\tau^c_{i,t-\bar{h}}, \tau^{p,h}_{i,t-\bar{h}}, \tau^{p,f}_{i,t-\bar{h}}\right)$ are part of $\mathcal{I}^*_{t-\bar{h}}$, the orthogonality requirement of Theorem 1 implies that $\left(\tau^c_{i,t}, \tau^{p,h}_{i,t}, \tau^{p,f}_{i,t}\right)$ has a finite MA representation of order $\bar{h} - 1$. In Sections III and IV, we use Proposition 5 to analyze the feasible dynamics of rational expectation errors in a calibrated version of our model.

D. *Aggregation and Equilibrium in the Primal Economy*

Before exploring how Proposition 5 restricts the equilibrium dynamics in this economy, we conclude this section with an explicit characterization of equilibrium in the aggregate primal economy. Unlike the solution to the incomplete-information economy, which requires keeping track of the cross-sectional distribution of beliefs, the primal economy permits a simple aggregate representation of equilibrium.

Let $\tau^c_t = \int_0^1 \tau^c_{i,t}\,\mathrm{d}i$, $\tau^{p,h}_t = \int_0^1 \tau^{p,h}_{i,t}\,\mathrm{d}i$, and $\tau^{p,f}_t = \int_0^1 \tau^{p,f}_{i,t}\,\mathrm{d}i$ denote the "macro" wedges. Integrating over (32) and (33), we have

$$\hat{y}_t = \mathbb{E}_t\left[\hat{y}_{t+1} - \tau^c_{t+1} - \phi\pi_t + \pi_{t+1}\right] + \tau^c_t \tag{34}$$

$$\hat{y}_t = \xi\left(\tau^{p,f}_t - \tau^{p,h}_t\right), \tag{35}$$

where $\hat{y}_t \equiv y_t - a_t$ is the level of output relative to its (full-information) potential.

Equations (34) and (35) define the aggregate dynamics in the primal economy. Common prediction errors in the Euler equation, captured by τ^c_t, show up as an Euler equation wedge. Similarly, common prediction errors regarding each islands' terms-of-trade, $p_{i,t} - p_t$, are captured by $\tau^p_t \equiv \tau^{p,f}_t - \tau^{p,h}_t$, which corresponds to the labor wedge in our economy that is composed of a household and a firm component. The aggregate "wedges" τ^c_t and τ^p_t are the sole drivers of the output gap and inflation. If all agents had full information ($\tau^c_t = \tau^p_t = 0$), the aggregate economy would be in its first-best equilibrium in which output reaches its potential in every period ($y_t = a_t$) and inflation is always zero.

In general a solution for endogenous variables as a function of the joint process $\tau_t \equiv \left(\tau^c_t, \tau^p_t\right)'$ can be obtained using standard numerical tools. In our case a closed-form solution is also available. Substituting for $\hat{y}_t$ in (34), π_t is characterized by the prediction formula

$$\pi_t = \phi^{-1}\mathbb{E}_t\left[\xi d\tau^p_{t+1} - d\tau^c_{t+1} + \pi_{t+1}\right]. \tag{36}$$

Following Hansen and Sargent (1980, 1981), we obtain an explicit solution for inflation, stated in the following.

LEMMA 1: *Let $\tau_t = A(L)\,u_t$, where $A(L)$ is a square-summable lag polynomial in nonnegative powers of L and the innovations u_t are orthogonal white noise. Then there exists a unique stationary equilibrium process for $(\hat{y}_t, \pi_t)$, given by*

$$\hat{y}_t = \left[0 \;\; \xi\right]A(L)\,u_t \tag{37}$$

and

$$\pi_t = \left[-1 \;\; \xi\right] \frac{(1-L)A(L) - \left(1-\phi^{-1}\right)A\left(\phi^{-1}\right)}{\phi L - 1} u_t. \tag{38}$$

III. Inference about the Aggregate Economy

In this section we explore how the theoretical restrictions of Proposition 5 translate into restrictions on the behavior of the aggregate economy. In a first step, Section IIIA maps the restrictions stated in Proposition 5 into restrictions on the dynamics of the "macro" wedges determining the behavior of the aggregate economy. Sections IIIB and IIIC then use these restrictions on the macro wedges to characterize feasible volatility and comovement patterns of output and inflation under varying assumptions on information and fundamentals.

A. *Feasible Dynamics of Aggregate Wedges*

We begin by mapping the orthogonality restrictions in Proposition 5 into restrictions on the macro wedges τ_t^c and τ_t^p. To streamline the exposition, we only detail the derivation for the baseline case $\Theta_{i,t}^{\text{sym}}$ depicted in (22), in which firms and households have symmetric information.

To begin, observe that for $\Theta_{i,t}^{\text{sym}}$, $\left\{\mu_{i,t-s}, \Theta_{i,t-s}^{\text{sym}}\right\}_{s\geq 0}$ satisfies Assumption 3 with

$$\mathcal{S}_{i,t} = \left\{dc_{i,t}, dy_{i,t}, da_{i,t}\right\}.$$

Here we have used that (i) $n_{i,t}$ and $w_{i,t}$ are linear combinations of $\left(c_{i,t}, y_{i,t}, a_{i,t}\right)$ and are therefore informationally redundant; and (ii) that for any finite horizon $\bar{h}$, observing the sequence of *differences* $\left\{\mathcal{S}_{i,t-s}\right\}_{s=0}^{\bar{h}-1}$ in addition to $\mathcal{I}_{t-\bar{h}}^*$ contains the same information as the corresponding sequence of *levels*.

To proceed, define $\tau_{i,t} \equiv \left(\tau_{i,t}^c, \tau_{i,t}^{p,f}, \tau_{i,t}^{p,h}\right)'$ and let $\Delta\tau_{i,t} \equiv \tau_{i,t} - \tau_t$ denote the idiosyncratic portion of the expectation wedges. Similarly, let $\left(\Delta c_{i,t}, \Delta y_{i,t}\right)$ denote the idiosyncratic deviations from aggregate consumption and output. By construction the "Delta"-component of any variable is orthogonal to any aggregate variable. Hence, for any two variables $x_{i,t}$ and $y_{i,t}$, we have $\text{cov}\left[x_{i,t}, y_{i,t}\right] = \text{cov}\left[x_t, y_t\right] + \text{cov}\left[\Delta x_{i,t}, \Delta y_{i,t}\right]$. The orthogonality requirement between $\tau_{i,t}$ and $\mathcal{S}_{i,t}$ can then be written as

$$\text{cov}\left[\tau_t, \left(dy_{t-s}, dy_{t-s}, \epsilon_{t-s}\right)\right] = -\text{cov}\left[\Delta\tau_{i,t}, \left(\Delta dc_{i,t-s}, \Delta dy_{i,t-s}, \Delta da_{i,t-s}\right)\right] \tag{39}$$

$$\text{for all } s \geq 0.$$

Condition (39) requires that any aggregate comovement on the left-hand side is exactly offset by corresponding "Delta" comovements on the right-hand side. It is the analogue to conditions (14) and (15) in the simple price-setting application.

The main complication compared to the price-setting application is that the endogenous "Delta"-variables on the right-hand side, $\Delta dc_{i,t}$ and $\Delta dy_{i,t}$, can no longer be expressed as static functions of $\Delta\tau_{i,t}$ and fundamentals. Instead, $\Delta dc_{i,t}$ and

$\Delta dy_{i,t}$ are themselves a solution to a system of expectational difference equations. Specifically, subtracting y_t from both sides of (32) and (33), we obtain

$$\Delta c_{i,t} = \mathbb{E}_t\big[\Delta c_{i,t+1} - \Delta \tau^c_{i,t+1}\big] + \Delta \tau^c_{i,t} \tag{40}$$

$$\Delta y_{i,t} = \xi\big(\Delta y_{i,t} - \Delta c_{i,t} + \Delta p_{i,t} + \Delta \tau^p_{i,t}\big) + \Delta a_{i,t} \tag{41}$$

for $\Delta \tau^p_{i,t} = \Delta \tau^{p,f}_{i,t} - \Delta \tau^{p,h}_{i,t}$. Together with (27) and (28), conditions (40) and (41) define a (fictitious) small open economy, which can be solved independently from the economy's aggregates. While the endogenous nature of $\Delta dc_{i,t}$ and $\Delta dy_{i,t}$ impedes further analytical progress that parallels Propositions 1–4, condition (39) similarly entails restrictions on aggregate volatility and the (auto-)covariance structure of the economy, which can be characterized numerically.

For our numerical analysis below, we exploit that for any (zero mean) $\mathrm{MA}\big(\bar{h}\big)$ process for the idiosyncratic and aggregate components of $\tau_{i,t}$, condition (39) is both necessary and sufficient for the implementation of these wedges by some information structure. The set of feasible aggregate fluctuations is thus characterized by the set of aggregate processes $\{\tau^c_t, \tau^p_t\}$ for which (39) can be satisfied with some processes for the idiosyncratic components $\{\Delta \tau^c_{i,t}, \Delta \tau^p_{i,t}\}$. In general, one can obtain this characterization by numerically solving for the map from wedges to covariances, which entails finding equilibrium in the "Delta"-economy. In our case, we are able to simplify the search by solving the "Delta-economy" in closed form, which allows for a more efficient numerical implementation (see the derivation following Lemma 2 in the online Appendix for details.)

B. *Unrestricted Micro-Shock Benchmark*

Before proceeding to our quantitative results, we provide a theoretical benchmark for the case where we treat the idiosyncratic fundamentals, $\Delta f_{i,t} = \big(\Delta a_{i,t}, z_{i,t}\big)$, as unrestricted. Previous literature has shown that if idiosyncratic fundamentals are sufficiently volatile, then confusion about these shocks can be used to support aggregate fluctuations in $\hat{y}_t$, even if there are no aggregate shocks to fundamentals. This is because expectation errors regarding *local* shocks can be correlated across islands even though the underlying fundamentals are purely idiosyncratic (e.g., Angeletos and La'O 2013; Benhabib, Wang, and Wen 2015).

In the spirit of this literature, the following benchmark uses our methodology to characterize what dynamics are possible if we place no restrictions on $\Delta f_{i,t}$. By construction, the chosen process for $\Delta f_{i,t}$ has no direct impact on the aggregate economy. Its only role is to provide a source of uncertainty, which can be used to support aggregate fluctuations when information is incomplete.

PROPOSITION 6: *Fix a (zero mean)* $\mathrm{MA}\big(\bar{h}\big)$ *process* τ *for* $\big(\tau^c_t, \tau^p_t\big)$ *and set* $\Theta^{\mathrm{sym}}_{i,t}$ *as in* (22)*. Then for any aggregate productivity process*, a, *there exist idiosyncratic processes* $\Delta\tau$ *and* Δf, *such that* τ *can be implemented in the incomplete information economy.*

Proposition 6 provides a striking benchmark: absent microdata that disciplines $\Delta f_{i,t}$, correlated optimism and pessimism (across islands), can be used to generate *any* joint process in $(\hat{y}_t, \pi_t)$. Going beyond the results in Angeletos and La'O (2013) and Benhabib, Wang, and Wen (2015) on *volatility*, the benchmark shows that "sentiment" fluctuations can implement arbitrary *processes* for τ_t and, by implication, arbitrary autocorrelation structures among the aggregate variables, potentially bypassing all cross-equation restrictions that emerge under full information.[14] Intuitively, expectation errors can plausibly be correlated, both because information can be correlated between households and firms and because expectation errors by households generally affect both their consumption and labor supply.

C. *Quantitative Results*

In light of the "everything goes" result in Proposition 6, a natural question to ask is: what are the restrictions on aggregate dynamics once we fix $\Delta f_{i,t}$ at an empirically plausible calibration? We explore this question numerically, calibrating $\Delta f_{i,t}$ to existing microdata.

Parametrization.—We interpret one period as a quarter and set the discount factor β to 0.99. The inverse Frisch elasticity ζ is set to 0.5, the elasticity of substitution between input varieties η is set to 7.5, and the elasticity of the interest rate ϕ is set to 1.5. These values are within the range typically used by the literature.

Next, we set the incomplete information horizon to $\bar{h} - 1 = 14$ quarters. While we do not have strong priors regarding $\bar{h}$, our choice is consistent with the horizon at which Coibion and Gorodnichenko (2015) find a significant response in professional forecasters' expectation errors to various fundamental and nonfundamental shocks. Below, we explore the sensitivity of our results to $\bar{h}$, and show that once the horizon $\bar{h} - 1$ exceeds six periods, it has little impact on results.

It remains to choose processes for the island-specific productivities and demand. We separate local productivities into a persistent component, $x_{i,t}$, and a purely transient component, $\omega_{i,t}$,

$$\Delta a_{i,t} = x_{i,t} + \omega_{i,t},$$

where $\omega_{i,t}$ is i.i.d. with zero mean and variance σ_ω^2. The separation ensures that agents can be potentially confused about the precise state of $\Delta a_{i,t}$, even if there are no aggregate productivity shocks. The persistent components $\{x_{i,t}\}$ as well as the local demand shocks $\{z_{i,t}\}$ follow independent AR(1) processes with auto-correlations (ρ_x, ρ_z) and one-step-ahead variances (σ_x^2, σ_z^2). The variance and persistence parameters are set based on Foster, Haltiwanger, and Syverson (2008), who use plants' price data to disentangle demand from physical productivity shocks at the

[14] In three related contributions, Huo and Takayama (2015); Angeletos, Collard, and Dellas (2018); and Ilut and Saijo (2021) all provide examples of how learning may introduce nonzero correlation in wedges. However, in contrast to the result in Proposition 6, these comovement patterns are restricted by the specifics of the information-structures considered in these papers, translating into nontrivial cross-equation restrictions.

plant level. Specifically, we set $\rho_x = \rho_z = 0.976$, $\sigma_x = 0.0552$, $\sigma_\omega = 0.0478$, and $\sigma_z = 0.2504$, which imply within-product dispersions and quarterly autocorrelations of $z_{i,t}$ and $\Delta a_{i,t}$ that match the corresponding statistics in Foster, Haltiwanger, and Syverson (2008).[15]

It is worth noting that, in line with popular views, the data of Foster, Haltiwanger, and Syverson (2008) imply that demand shocks are much larger than productivity shocks (see also de Loecker 2011; Demidova, Kee, and Krishna 2012; Roberts et al. 2018; Foster, Haltiwanger, and Syverson 2016 for similar results). Intuitively, this is consistent with the idea that fluctuations in demand reflect both demand and supply shocks upstream in the production chain, which amplifies demand uncertainty relative to the uncertainty about within-firm technology. We explore the robustness of our results with respect to the scale of idiosyncratic shocks, considering a variety of calibrations in the exercises that follow.

Volatility Frontier (Definition).—We compute the maximal output volatility—as a function of its persistence and the cyclicality of inflation—that our model can generate in the absence of aggregate shocks to fundamentals ($\mathrm{var}[\epsilon_t] = 0$).

Formally, define $\sigma_{\hat{y}}(\tau) \equiv \sqrt{\mathrm{var}[\hat{y}_t | \mathcal{I}^*_{t-1}]}$ as the one-step-ahead volatility of output induced by τ. Similarly, define $\rho_{\hat{y}}(\tau) \equiv \mathrm{corr}[\hat{y}_t, \hat{y}_{t-1}]$ as the first-order autocorrelation of $\hat{y}_t$, and define $\gamma_{\hat{y}\pi}(\tau) \equiv \mathrm{corr}[\hat{y}_t, \pi_t]$ as the contemporaneous correlation with inflation. We numerically trace out the *volatility frontier* for output as a function of its autocorrelation $\rho_{\hat{y}}$ and its contemporaneous correlation with inflation $\gamma_{\hat{y}\pi}$:

$$\sigma_{\hat{y}}^{\max}(\bar{\rho}_{\hat{y}}, \bar{\gamma}_{\hat{y}\pi}) \equiv \max_{\tau, \Delta\tau} \{\sigma_{\hat{y}}(\tau)\}$$

subject to

$$\rho_{\hat{y}}(\tau) = \bar{\rho}_{\hat{y}}$$

$$\gamma_{\hat{y}\pi}(\tau) = \bar{\gamma}_{\hat{y}\pi}$$

and the implementability condition (39). Here τ and $\Delta\tau$ are independent (zero-mean) $\mathrm{MA}(\bar{h})$ processes.[16] The process for the idiosyncratic fundamentals $\Delta f = (\Delta a, z)$ is given by our calibration.

Baseline Case.—Figure 1 presents the volatility frontier for the baseline where firms and households have symmetric information within islands and $\Theta_{i,t}^{\mathrm{sym}}$ is given by (22). Here $\sigma_{\hat{y}}^{\max}$ is denominated in percentage deviations from the balanced growth path. The most striking feature is the discrepancy at $\gamma_{\hat{y}\pi} = 0$. When inflation is procyclical ($\gamma_{\hat{y}\pi} > 0$), incomplete information can explain an output volatility up to 1.8 percent. Evaluating $\sigma_{\hat{y}}^{\max}$ at values consistent with US data, $\gamma_{\hat{y}\pi} = 0.3$ and $\rho_{\hat{y}} = 0.9$, the maximal volatility amounts to 1.1 percent, which is about 9/10

[15] The underlying calibration targets are 0.976 and 0.943 for the quarterly persistence rates of $z_{i,t}$ and $\Delta a_{i,t}$, respectively, and 1.16 and 0.26 for the (unconditional) within-product dispersions.

[16] Without loss of generality, we restrict τ to load on at most three innovations. Similarly, we restrict $\Delta\tau$ to load on at most three innovations in addition to the fundamental shocks that drive Δf.

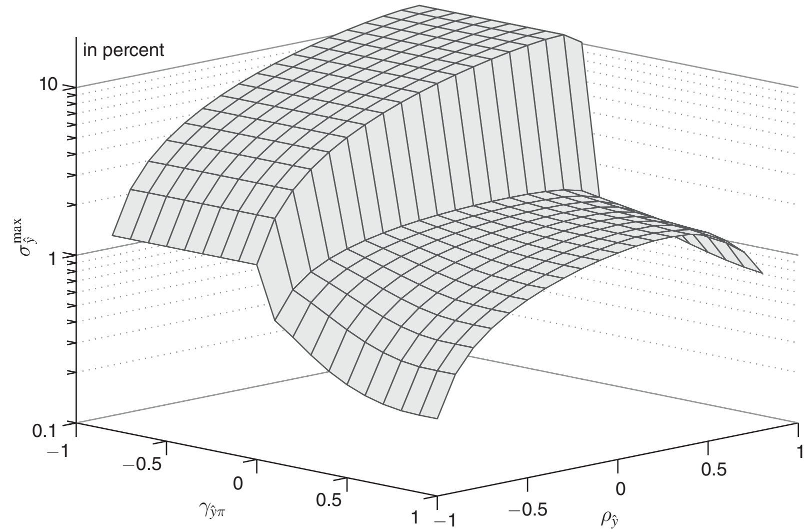

Figure 1. Feasibility Frontier

Notes: The graph shows the maximal output volatility (denominated in percentage deviations from the balanced growth path) that can be generated by incomplete information as a function of aggregate persistence $\rho_{\hat{y}}$ and the contemporaneous correlation with inflation $\gamma_{\hat{y}\pi}$.

of the corresponding volatility in the United States. By contrast, when inflation is countercyclical $(\gamma_{\hat{y}\pi} < 0)$, the maximal volatility is increased by about one order of magnitude.

The reason for the discrepancy is a fundamental difference in the channels through which the model generates procyclical and countercyclical inflation dynamics. As suggested by Proposition 2 (see online Appendix Section A.4 for a variant of the proposition applying to the quantitative model), countercyclical inflation dynamics are intrinsically tied to expectation errors regarding local demand, which can be quite large for the calibrated process for $z_{i,t}$. By contrast, procyclical inflation dynamics (typically) require some nominal misconception,[17] which is disciplined by the volatility of *aggregate* prices.

Micro Shocks and Macro Volatility.—How do changes in the specification of $\{\Delta a_{i,t}\}$ and $\{z_{i,t}\}$ affect the volatility frontier $\sigma_{\hat{y}}^{\max}$? To explore the link from micro-shocks to macro-volatility, we conduct comparative statics exercises in σ_x, σ_ω, σ_z, ρ_x, and ρ_z. Here we focus on the case where the macro-correlations $\gamma_{\hat{y}\pi}$ and $\rho_{\hat{y}}$ are respectively fixed at 0.3 and 0.9, consistent with US data. In online Appendix C, we extend the analysis to the case where inflation is countercyclical, finding qualitatively similar results to the ones below.

[17] Perceived fluctuations in *local* demand cannot induce procyclical inflation dynamics because of consumption smoothing. Under standard preferences, consumption (typically) goes up by less than output in response to a temporary increase in local demand. (This is true as long as $z_{i,t}$ is not too persistent; in our calibration it holds for $\rho_z \leq 0.997$.) The Taylor principle $(\phi > 1)$ then implies that expansions caused by correlated errors regarding $\{z_{i,t}\}$ must be accompanied by a *drop* in inflation so that consumption and output are equilibrated through the expected decline in the real interest rate.

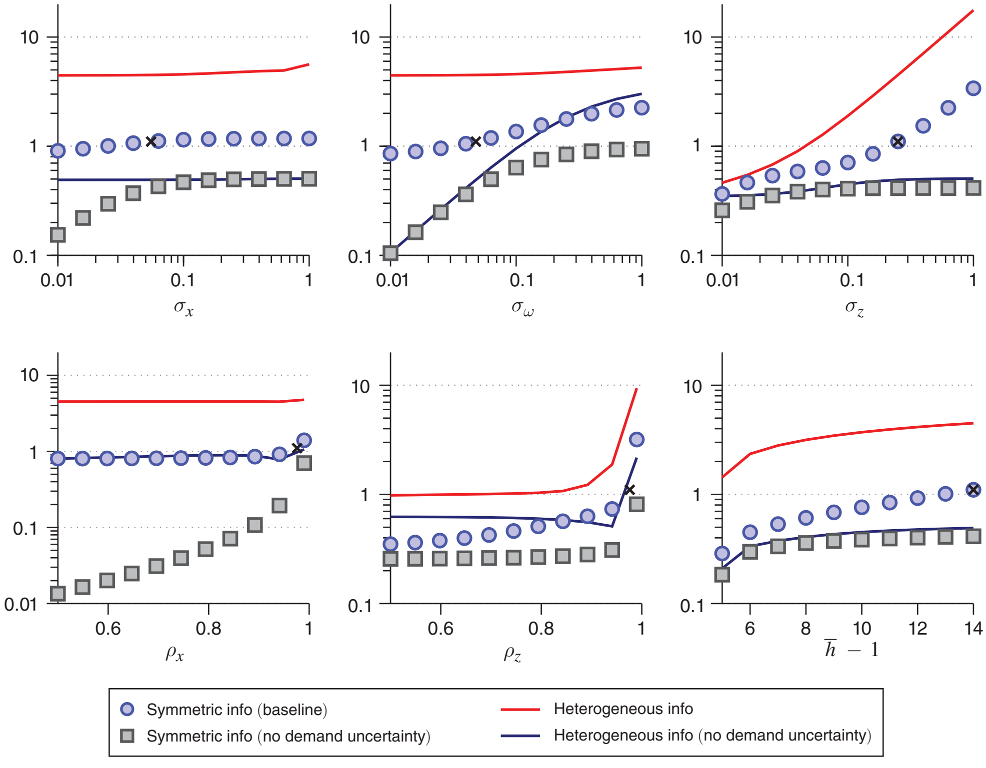

Figure 2. Comparative Statics

Notes: Feasibility frontier for alternate specifications of the micro-shocks $\{\Delta a_{i,t}, z_{i,t}\}$ and for alternate information-bounds $\{\Theta_{i,t}\}$. The graphs show the maximal output volatility $\sigma_{\hat{y}}^{\max}$ (denominated in percentage deviations from the balanced growth path) that can be generated by incomplete information for the case where $\rho_{\hat{y}} = 0.9$ and $\gamma_{\hat{y}\pi} = 0.3$. The "$\times$"-marks indicate the case where both the micro-shocks and $\Theta_{i,t}$ are fixed at their baseline values shown in Figure 1.

The results are presented in Figure 2. The blue dots in panels 1–5 correspond to the case where households have symmetric information as assumed above. For comparison, the baseline calibration, for which $\sigma_{\hat{y}}^{\max} \approx 1.1$ percent, is indicated by the "$\times$"-marks in the figure.

The sensitivity is strongest in σ_z and ρ_z, indicating that correlated expectation errors about the demand shocks $\{z_{i,t}\}$ are of critical importance for supporting fluctuations in aggregate confidence. In particular, a reduction in σ_z from its baseline value of 0.2504 to 0.01 reduces $\sigma_{\hat{y}}^{\max}$ by a factor of three to 0.37 percent; an increase in σ_z to 1.00 increases $\sigma_{\hat{y}}^{\max}$ to 3.39 percent. Those comparative statics reflect the naturally increasing shape of $\sigma_{\hat{y}}^{\max}$ in any fundamental volatility. Intuitively, the more volatile $z_{i,t}$ (and $a_{i,t}$), the larger the potential for agents to make expectation errors, which is a direct consequence of the law of total variance ($\mathrm{var}\left[\mathbb{E}\{z_{i,t}|\mathcal{I}_{i,t}\}\right] \leq \mathrm{var}[z_{i,t}]$). In the extreme case where $\sigma_z \to 0$, rationality requires that $\mathbb{E}\left[z_{i,t}|\mathcal{I}_{i,t}\right] = 0$ for all t, even if $\mathcal{I}_{i,t}$ contains no information about $z_{i,t}$.

Similarly to σ_z, variations in the persistence of $z_{i,t}$ also have a significant impact on $\sigma_{\hat{y}}^{\max}$: a reduction of ρ_z from its baseline value of 0.976 to 0.5, reduces $\sigma_{\hat{y}}^{\max}$ to

0.35 percent. An increase in the persistence of $z_{i,t}$ to 0.99, increases $\sigma_{\hat{y}}^{\max}$ to 3.18. The role of ρ_z for supporting expectation errors is two-fold. First, $\mathrm{var}[z_{i,t}]$ is increasing in ρ_z, again increasing the potential for expectation errors. Second, persistence in $z_{i,t}$ (or in $\Delta a_{i,t}$), enables optimism and pessimism regarding the wealth of the local household, independently from the direct effects on contemporaneous labor supply and demand. As fluctuations in *perceived* wealth translate into fluctuations in desired consumption, they can be used to induce pro-cyclical inflation dynamics as in Lorenzoni (2009), which is instrumental for generating the targeted cyclicality of inflation $(\gamma_{\hat{y}\pi} = 0.3)$.[18]

By contrast, variations in the parameters of $\{a_{i,t}\}$ result in only moderate variations in $\sigma_{\hat{y}}^{\max}$. In particular, reducing σ_x or σ_ω to 0.01, implies only marginally smaller values of $\sigma_{\hat{y}}^{\max}$, suggesting that the idiosyncratic productivity shocks $\{\Delta a_{i,t}\}$ play a somewhat dispensable role in our calibration. This reflects two factors. First, given our calibration, productivity is less volatile than demand, implying that there is less scope for productivity-related confusion in the first place. Second, because $a_{i,t} \in \Theta_{i,t}$, firms and households always know their current productivity, limiting productivity-related confusion to uncertainty about the composition of $\Delta a_{i,t}$, whose relevance in turn is determined by the persistence of $x_{i,t}$.

No Demand Uncertainty.—So far, we have not taken a stand whether or not agents know the inverse demand for the local good, $p_{i,t}$. As an alternative, we now consider the case where $p_{i,t}$ is perfectly observed, so that there is no uncertainty about the revenues associated with a particular choice of production. Formally, information is now bounded by

$$\Theta_{i,t} = \{p_{i,t-s}\}_{s\geq 0} \cup \Theta_{i,t}^{\mathrm{sym}}$$

with $\Theta_{i,t}^{\mathrm{sym}}$ given by (22). Because $\tau_{i,t}^{p,f}$ measures firms' expectation error regarding $p_{i,t}$, an immediate consequence of including $p_{i,t}$ in $\Theta_{i,t}^{f}$ is that $\tau_{i,t}^{p,f} = 0$ for all i and t, so that fluctuations in aggregate output can only be driven by the households' component of the labor wedge. Intuitively, firms only need to know their marginal costs, $w_{i,t} - a_{i,t}$, and their local demand, $p_{i,t}$, to behave *as if* they have full information (see also Hellwig and Venkateswaran 2014).

For the baseline parametrization of $\{\Delta a_{i,t}, z_{i,t}\}$, shutting down $\tau_t^{p,f}$ reduces $\sigma_{\hat{y}}^{\max}$ to 0.41, suggesting that uncertainty about demand is key to generating sizable aggregate fluctuations. Moreover, compared to the case where $\Theta_{i,t}^{\mathrm{sym}}$ is given by (22), the sensitivity of $\sigma_{\hat{y}}^{\max}$ in the parameters of $\{z_{i,t}\}$ is reduced, whereas the sensitivity in the parameters of $\{a_{i,t}\}$ is heightened (illustrated by the gray squares in Figure 2). This is because when $p_{i,t}$ is known, agents can back out the state of $z_{i,t} + p_t - \eta^{-1} y_t$ from (20), reducing the scope to generate waves of optimism and pessimism via $z_{i,t}$

[18] In order to generate procyclical inflation dynamics through optimism and pessimism about $z_{i,t}$, the information structure must mute the direct substitution effect on labor demand. This can be achieved, for instance, by making agents (sufficiently) informed about $p_{i,t}$ (coupled with some nominal misconception as in Lucas (1972, 1973), so that $p_{i,t}$ does not fully reveal $z_{i,t}$), which is a sufficient statistic about $\mathbb{E}[z_{i,t}|\mathcal{I}_{i,t}]$ for determining labor demand.

and, by implication, increasing the model's reliance on $\Delta a_{i,t}$ for supporting aggregate fluctuations in confidence.[19]

Heterogeneous Information.—We next relax the assumption that households and firms share the same information set, setting $\Theta^h_{i,t}$ and $\Theta^f_{i,t}$ as in (23) and (24). The resulting volatility frontier is depicted by the red lines in Figure 2. For the baseline calibration, this increases $\sigma^{\max}_{\hat{y}}$ to 4.49 percent. This reflects the additional flexibility in $\mathcal{I}^f_{i,t}$ and $\mathcal{I}^h_{i,t}$, due to households not being required to perfectly know the local firm's productivity (i.e., $a_{i,t}, y_{i,t} \notin \Theta^h_{i,t}$) and firms not being required to perfectly know households' consumption ($c_{i,t} \notin \Theta^f_{i,t}$). Specifically, this enables waves of optimism and pessimism among households about income-fluctuations caused by $\Delta a_{i,t}$ and $z_{i,t}$, translating to aggregate demand fluctuations—even if $\Delta a_{i,t}$ and $z_{i,t}$ are observed by firms. The stark increase in $\sigma^{\max}_{\hat{y}}$ suggests that the usual assumption of symmetric information may in fact be quite restrictive.

Finally, we explore a variant of the heterogeneous information setting where firms face no demand uncertainty ($\Theta^f_{i,t}$ includes $\{p_{i,t-s}\}_{s\geq 0}$ in addition to (24)). The results are depicted by the blue lines in Figure 2). Compared to the symmetric-information case without demand uncertainty, $\sigma^{\max}_{\hat{y}}$ is slightly increased to 0.49. However, the difference between symmetric and heterogeneous information is now much less pronounced, suggesting that imposing informational symmetry is somewhat less restrictive when firms know their demand while making their production choices.

Effects of Incomplete-Information Horizon.—As a final comparative static, we evaluate the sensitivity of $\sigma^{\max}_{\hat{y}}$ in the incomplete information horizon $\bar{h}$. Because the autocorrelation of any MA$\left(\bar{h} - 1 \leq 4\right)$ process is bounded above by less than the targeted autocorrelation ($\rho_{\hat{y}} = 0.9$), we have $\sigma^{\max}_{\hat{y}} = 0$ for all $\bar{h} - 1 \leq 4$. Conditional on $\bar{h} - 1 \geq 5$, the impact of $\bar{h}$ is moderate, especially for the cases without demand uncertainty. For the baseline symmetric information case, the impact is somewhat more pronounced, reducing $\sigma^{\max}_{\hat{y}}$ to 0.76 when $\bar{h} - 1$ is reduced to 10 quarters.

IV. Application to US Business Cycles

We now explore the degree to which US business cycle data is consistent with a theory of incomplete information. To this end, we first estimate an unrestricted wedge process $\hat{\tau}_t \equiv \left(\hat{\tau}^c_t, \hat{\tau}^p_t\right)$ that in the tradition of Chari, Kehoe, and McGrattan (2007) best describes the data. We then partition $\hat{\tau}_t$ into an informational component τ^{info}_t (restricted by our theoretical characterization) and an unrestricted residual component τ^{resid}_t, and maximize the contribution of the informational component τ^{info}_t under varying assumptions on $\left\{\Delta a_{i,t}, z_{i,t}\right\}$ and $\left\{\Theta_{i,t}\right\}$.

[19] The sensitivity in $z_{i,t}$ is not reduced to zero for two reasons. First, $z_{i,t}$ serves as noise about the aggregate state. Second, despite there being no uncertainty about *current* $p_{i,t}$, expectation errors about $z_{i,t}$ continue to translate into optimism and pessimism about *future* prices whenever $\rho_z \neq 0$, which affects local wealth and households' consumption choice.

Table 1—Summary of Estimated US Wedges

	Standard deviation	First-order autocorr.	Contemporaneous correlation with $\hat{\tau}_t^c$	with $\hat{\tau}_t^p$	with $\hat{\epsilon}_t$
$\hat{\tau}_t^c$	0.051	0.91	1.00	·	·
$\hat{\tau}_t^p$	0.044	0.91	0.99	1.00	·
$\hat{\epsilon}_t$	0.010	—	−0.27	−0.27	1.00

A. *Methodology*

Here we briefly describe the initial estimation step and then formalize our approach to partitioning the estimated wedge process into an informational and residual component. A detailed description of the initial estimation can be found in online Appendix B. Throughout the model is calibrated as in Section IIIC.

Summary of Estimation Step.—We use the generalized method of moments (GMM) to estimate the process $\hat{\tau}_t$ that best matches the auto-covariance structure of quarterly US data on real per capita output, inflation, nominal interest rates, and per capita hours, targeting all auto-covariances between zero and eight quarters. All moments are computed at business cycle frequencies, applying a high-pass filter with a cutoff of 32 quarters to the model and the data. We model $\hat{\tau}$ as MA(14) processes, which loads on two intrinsic innovations, denoted by $\hat{u}_t$, in addition to the productivity shock $\hat{\epsilon}_t$.

Despite targeting more data series than there are shocks, the estimated process $\hat{\tau}_t$ fits the data quite well: the model replicates the US auto-covariance structure within the confidence bands of the data (see Figure 4 in the online Appendix). The productivity shock $\hat{\epsilon}_t$ explains about 36 percent of the filtered variance in $\hat{y}_t$ and about 11 percent to the filtered variance of y_t.[20] The remaining fluctuations are explained by intrinsic innovations in the estimated wedges $\hat{\tau}_t^c$ and $\hat{\tau}_t^p$.

Table 1 summarizes key moments of the estimated wedges $\left(\hat{\tau}_t^c, \hat{\tau}_t^p\right)$ and the estimated productivity shock $\hat{\epsilon}_t$.

Most noticeable is the strong positive correlation between the Euler wedge and the labor wedge $(\text{corr}\left[\hat{\tau}_t^c, \hat{\tau}_t^p\right] = 0.99)$ and both wedges' negative correlation with productivity growth $(\text{corr}\left[\hat{\tau}_t, \hat{\epsilon}_t\right] = -0.27)$. The high correlation of $\hat{\tau}_p$ and $\hat{\tau}_t$ may be somewhat surprising in light of previous results in the wedge accounting literature. We note this finding is *not* a consequence of abstracting from capital per se, but rather the fact that we measure the Euler wedge directly from data on i_t and π_t, whereas the business cycle literature typically infers real interest rates indirectly by using the time series of investment to infer the marginal product of capital through the lens of a model.[21] Using our approach, the real rate fed into the Euler

[20] The contribution of a_t to $\hat{y}_t$ exceeds the one to y_t, due to a negative correlation between a_t and $\hat{y}_t$, reflecting a slow adjustment in response to productivity shocks.

[21] Our wedges also differ from those measured in standard RBC models due to our imposition of $C_t = Y_t$. Because consumption is highly correlated with output, this difference is minor. Abstracting from differences in measurement, the wedges implied by our model are identical to those implied by standard RBC models. Specifically, given households' preferences and data for $\left\{C_t, r_t\right\}$, the Euler wedge is trivially identical. Moreover, with preferences

equation moves very little, reflecting the low volatility of both inflation and the nominal rate. Accordingly, to match both empirical consumption and inflation dynamics, the model requires the Euler and labor wedge to be highly correlated. (For intuition, notice that Lemma 1 implies that, as $\mathrm{var}[\pi_t] \to 0$, the two wedges are perfectly correlated.)

Partitioning of the Estimated Wedges.—We partition the estimated wedge process $\hat{\tau}_t$ into an informational component τ_t^{info} and a residual component τ_t^{resid},

$$\hat{\tau}_t = \tau_t^{\mathrm{info}} + \tau_t^{\mathrm{resid}}. \tag{42}$$

In parallel to $\hat{\tau}_t$, we model both components as statistically independent MA(14) processes,

$$\tau_t^{\mathrm{info}} = \Phi_\epsilon^{\mathrm{info}}(L)\,\epsilon_t^{\mathrm{info}} + \Phi_u^{\mathrm{info}}(L)\,u_t^{\mathrm{info}}$$

$$\tau_t^{\mathrm{resid}} = \Phi_\epsilon^{\mathrm{resid}}(L)\,\epsilon_t^{\mathrm{resid}} + \Phi_u^{\mathrm{resid}}(L)\,u_t^{\mathrm{resid}},$$

where $\Phi_\epsilon^{\mathrm{info}}$, Φ_u^{info}, $\Phi_\epsilon^{\mathrm{resid}}$, and Φ_u^{resid} are square-summable lag polynomials in nonnegative powers of L. The innovations, $\epsilon_t^{\mathrm{info}}$, $\epsilon_t^{\mathrm{resid}}$, u_t^{info}, and u_t^{resid}, are mutually orthogonal white noise. In particular, $\epsilon_t^{\mathrm{info}}$ and $\epsilon_t^{\mathrm{resid}}$ are innovations to aggregate productivity, satisfying

$$\hat{\epsilon}_t = \epsilon_t^{\mathrm{info}} + \epsilon_t^{\mathrm{resid}}, \tag{43}$$

with standard deviations $\sigma_\epsilon^{\mathrm{info}}$ and $\sigma_\epsilon^{\mathrm{resid}}$. The corresponding lag-polynomial $\Phi_\epsilon^{\mathrm{info}}$ captures how incomplete information regarding a_t influences the propagation of productivity shocks.[22] The innovations u_t^{info} and u_t^{resid}, each two-dimensional, are intrinsic shocks to τ_t^{info} and τ_t^{resid}. Accordingly, the lag-polynomial Φ_u^{info} defines intrinsic fluctuations in τ_t^{info}, driven by expectation errors, whereas Φ_u^{resid} defines intrinsic fluctuations in the residual wedges τ_t^{resid}.

The defining difference between τ_t^{info} and τ_t^{resid} is that we impose the conditions of Theorem 1 on τ_t^{info}, whereas τ_t^{resid} remains unrestricted. We gauge the potential role of incomplete information for explaining the US business cycle by maximizing the contribution of expectation errors u_t^{info} to the filtered variance of $\hat{y}_t$. Let $\hat{y}_t^{\mathrm{tfp}} \equiv \mathbb{E}\big[\hat{y}_t \big| \big(\epsilon_{t-s}^{\mathrm{info}}, \epsilon_{t-s}^{\mathrm{resid}}\big)_{s\geq 0}\big]$, $\hat{y}_t^{\mathrm{info}} \equiv \mathbb{E}\big[\hat{y}_t \big| \big(u_{t-s}^{\mathrm{info}}\big)_{s\geq 0}\big]$, and $\hat{y}_t^{\mathrm{resid}} \equiv \mathbb{E}\big[\hat{y}_t \big| \big(u_{t-s}^{\mathrm{resid}}\big)_{s\geq 0}\big]$ denote the projection of the output gap on aggregate productivity, expectation errors, and residual shocks, respectively. Independence of the innovations implies

unchanged, any difference in the labor wedge must be due to a change in firms' marginal product of labor. However, under Cobb Douglas production, the marginal product of labor in a model with capital share $(1-\alpha)$ is just $\alpha Y_t/N_t$, which in log-deviations from the steady state is identical to our labor wedge Y_t/N_t.

[22] Conversely, $\Phi_\epsilon^{\mathrm{resid}}$ captures the effects of other potential frictions in propagating productivity shocks. Splitting aggregate productivity into two independent innovations ensures that the volatility generated by incomplete information is independent of the residual wedges τ_t^{resid}. If we instead let τ_t^{info} and τ_t^{resid} load jointly on the combined productivity shock ϵ_t, we find that one can increase the variance contribution of u_t^{info} almost arbitrarily through incomplete information regarding a_t and its propagation through τ_t^{resid}. Below we also consider the case where agents perfectly observe aggregate productivity, in which case both settings give identical results.

$\mathrm{var}[\hat{y}_t] = \mathrm{var}[\hat{y}_t^{\mathrm{tfp}}] + \mathrm{var}[\hat{y}_t^{\mathrm{info}}] + \mathrm{var}[\hat{y}_t^{\mathrm{resid}}]$. Then the maximal contribution of u_t^{info} is given by

$$\max_{\tau^{\mathrm{info}},\tau^{\mathrm{resid}},\sigma_\epsilon^{\mathrm{info}},\sigma_\epsilon^{\mathrm{resid}}} \left\{ \mathrm{var}\left[\hat{y}_t^{\mathrm{info}}\right] / \mathrm{var}\left[\hat{y}_t\right] \right\} \tag{44}$$

subject to two constraints. First, there must exist a (zero-mean) $\mathrm{MA}(\bar{h})$ process for $\{\Delta \tau_{i,t}\}$ so that the informational component τ_t^{info} is implementable as characterized in Theorem 1. Second, we require that the auto-covariance structure for $(\hat{y}_t, \pi_t, \epsilon_t)$ induced by $(\tau_t^{\mathrm{info}}, \tau_t^{\mathrm{resid}}, \epsilon_t^{\mathrm{info}}, \epsilon_t^{\mathrm{resid}})$ is identical to the one induced by $(\hat{\tau}_t, \hat{\epsilon}_t)$. Thus, our partitioned wedges are constrained to produce output, productivity and inflation dynamics that jointly match those of the United States.

Observe that $\mathrm{var}[\hat{y}_t^{\mathrm{tfp}}]$ and $\mathrm{var}[\hat{y}_t]$ are fully pinned down by the estimated wedge process $\hat{\tau}_t$. Hence, instead of maximizing the contribution of u_t^{info} to $\mathrm{var}[\hat{y}_t]$, we can equivalently maximize the contribution of u_t^{info} to the portion of $\hat{y}_t$ that is not driven by the productivity shock, $\mathrm{var}[\hat{y}_t | \{a_{t-s}\}_{s\geq 0}] = \mathrm{var}[\hat{y}_t] - \mathrm{var}[\hat{y}_t^{\mathrm{tfp}}]$.

B. *Results*

The results are presented in Figure 3. To assess which conditions are necessary for incomplete information to generate sizable aggregate fluctuations, we consider five specifications for the lower bounds $\{\Theta_{i,t}\}$, represented by the five lines in the graph. Along the principal axis, we also consider variations in the parametrization of the micro-shocks $\{\Delta a_{i,t}, z_{i,t}\}$, scaling their standard deviations, $(\sigma_x, \sigma_\omega, \sigma_z)$, by up to ± 1 order of magnitude relative to the baseline calibration.[23] With the exception of the symmetric information benchmark, all specifications allow households and firms to have access to potentially heterogeneous information.

Benchmarks.—As benchmark, we first consider the symmetric information case where $\Theta_{i,t}^{\mathrm{sym}}$ is set as in (22) and the heterogenous information case where $\Theta_{i,t}^h$ and $\Theta_{i,t}^f$ are set as in (23) and (24). In both cases few restrictions are imposed on information beyond rational expectations. Perhaps not surprisingly in light of our theoretical benchmark in Proposition 6, confidence shocks can fully account for all US business cycle fluctuations unexplained by the productivity shock ($\mathrm{var}[\hat{y}_t^{\mathrm{info}}]/\mathrm{var}[\hat{y}_t | \{a_{t-s}\}_{s\geq 0}] \approx 1$), provided that $(\sigma_x, \sigma_\omega, \sigma_z)$ are at least as volatile as in our baseline calibration (scale ≥ 1).[24] For the asymmetric information case (red line), the result is also robust to a downward-scaling of the micro-shocks by up to a factor of three. For the symmetric information case (blue dotted line), a reduction in the micro-volatilities by a factor of two (three), reduces the maximal contribution to 90 percent (67 percent).

[23] The scaling is applied to all three micro-shocks proportionately to their respective baseline values; i.e., the scaled standard deviations are given by $(\sigma_x, \sigma_\omega, \sigma_z) \times$ scale.

[24] Note that this also implies a perfect account of all inflation-dynamics that are unexplained by the productivity shock, since the partitioning of the wedges is constrained to implement the empirical covariance structure for $(\hat{y}_t, \pi_t, \epsilon_t)$.

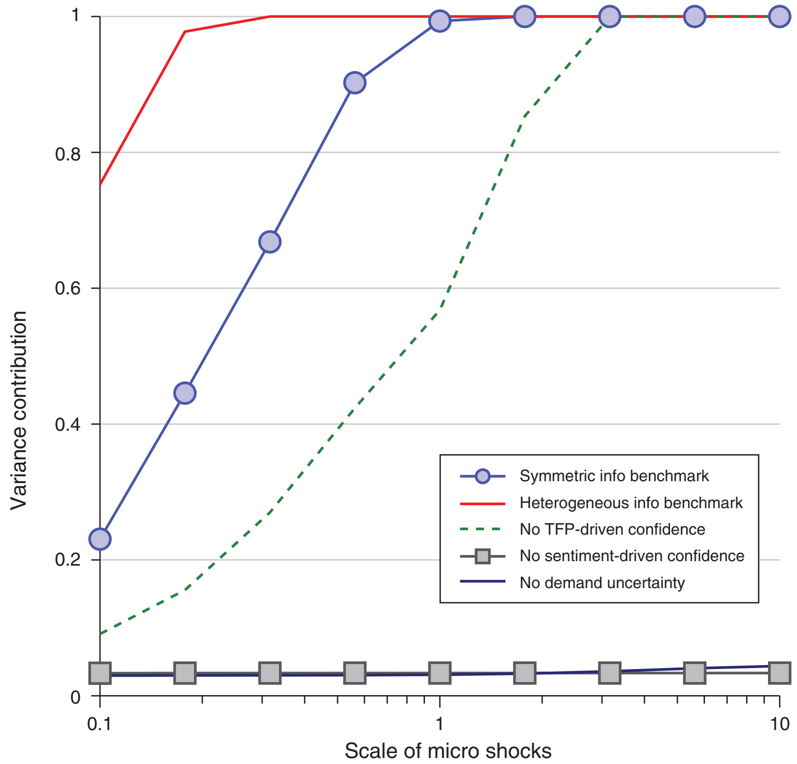

Figure 3. Maximal Contribution to US Business Cycle Volatility

Notes: The graph shows the maximal variance contribution of u_t^{info} to the portion of the US output gap not driven by productivity, $\text{var}\left[\hat{y}_t \middle| \{a_{t-s}\}_{s\geq 0}\right]$, computed at business cycle frequencies. The lines correspond to different assumptions on the lower bound of information $\{\Theta_{i,t}\}$. The variation on the principal axis considers alternative values for $(\sigma_x, \sigma_\omega, \sigma_z)$, which are scaled by up to ± 1 order of magnitude relative to the baseline calibration (scale $=1$).

Sentiments versus Noisy Learning about Aggregate Shocks.—The benchmarks show that, in combination with productivity shocks, rational fluctuations in confidence have the potential to fully account for the US business cycle. We now take a closer look at which type of confidence fluctuations are necessary to achieve this. Specifically, we differentiate between two types of confidence: (i) correlated confidence about *idiosyncratic* business conditions (aka "sentiment shocks"), and (ii) correlated confidence about *aggregate* productivity as in Angeletos and La'O (2010) or about future average productivity as in Lorenzoni (2009).

First, consider the case of sentiment shocks. We isolate their potential contribution by imposing perfect knowledge about the history of aggregate productivity by setting $\Theta_{i,t}^f$ and $\Theta_{i,t}^h$ as in (23) and (24), augmented by $\{a_{t-s}\}_{s\geq 0}$, eliminating any scope for TFP-driven fluctuations in confidence. Comparing the resulting contribution (dashed green line) with the benchmark reveals that for *small* scales of the micro shocks, confidence about aggregate productivity is indeed key for explaining the data. On the other hand, when there is sufficient idiosyncratic volatility (scale ≥ 3), sentiment shocks alone can do as well as the benchmark. For the baseline calibration (scale $= 1$), sentiment shocks can account for 57 percent of non-productivity fluctuations in US output.

Next, consider the case without sentiment shocks. To eliminate them, we set $\Theta^f_{i,t}$ and $\Theta^h_{i,t}$ as in (23) and (24), augmented by $\{x_{i,t-s}, z_{i,t-s}\}_{s\geq 0}$. Here we do not include the i.i.d.-productivities, $\omega_{i,t}$, in $\Theta^f_{i,t}$ or $\Theta^h_{i,t}$ as this would allow firms to fully back out a_t from observing $a_{i,t}$. However, because $\omega_{i,t}$ is serially uncorrelated and firms know $a_{i,t}$, expectation errors about $\omega_{i,t}$ have *no* direct effect on their actions, so that all fluctuations in confidence indeed reflect imperfect information about the aggregate productivity state. The quantitative results are shown by the gray squared lines in Figure 3. Under the baseline calibration of the micro-shocks (scale $= 1$)[25], TFP-driven fluctuations in confidence can explain at most 3.4 percent of the empirical output volatility, indicating that sentiment-driven fluctuations in confidence are indispensable for explaining the US business cycle with information frictions. This is because aggregate productivity shocks have only a limited importance by themselves, which in turn limits the potential for optimism regarding them to drive the business cycle.[26]

Interestingly, however, the two cases without sentiment- and productivity-driven confidence add up to less than the benchmark, indicating a complementarity between sentiments and confidence about aggregate productivity. Such complementarity may arise, because confidence fluctuations of one type may serve as noise in endogenous signals regarding the other type of fundamental shock.[27] Confidence about aggregate productivity shocks may therefore induce additional confidence about local conditions, and visa versa.

No Demand Uncertainty.—The final specification explores the case where firms know their demand when making their production choices, where $\Theta^f_{i,t}$ as in (24) is augmented by $\{p_{i,t-s}\}_{s\geq 0}$ (solid blue line). In this case, the maximal contribution to the empirical business cycle volatility amounts to 4.1 percent, which is almost as low as when fully shutting down all sentiment fluctuations. The result reinforces our earlier finding that demand uncertainties are key for generating sizable sentiment fluctuations and, more generally, sizable confidence fluctuations of any kind.

Implied Variance Contribution to US Output.—The results in Figure 3 show the business cycle contributions to output volatility that is unexplained by productivity, $\text{var}\left[\tilde{y}_t \middle| \{a_{t-s}\}_{s\geq 0}\right]$ (equivalently $\text{var}\left[y_t \middle| \{a_{t-s}\}_{s\geq 0}\right]$). Table 2 computes the implied contribution to the overall volatility in y_t and $\hat{y}_t$.

The discrepancy between the three columns reflects the contribution of the productivity shock to y_t and $\tilde{y}_t$. Looking at the contribution to y_t, sentiment-driven

[25] Here we recalibrate the local productivity shocks to attribute all productivity dispersion to $\omega_{i,t}$. This ensures that the inclusion of $x_{i,t}$ in $\Theta^f_{i,t}$ and $\Theta^h_{i,t}$ does not mechanically reduce the idiosyncratic noise that prevents firms from learning a_t from observing $a_{i,t} - x_{i,t} = a_t + \omega_{i,t}$.

[26] See Angeletos, Collard, and Dellas (2020) for independent evidence that productivity shocks play a small role in the business cycle. Indeed, Cochrane (1994) argues that all directly measurable aggregate shocks play a small role in driving business cycle fluctuations.

[27] See also Chahrour and Gaballo (2021).

Table 2—Implied Variance Contribution to US Output

	Contribution to		
	$\operatorname{var}[y_t \mid \{a_{t-s}\}_{s\geq 0}]$	$\operatorname{var}[y_t]$	$\operatorname{var}[\hat{y}_t]$
Heterogeneous info benchmark	1.00	0.89	0.64
Symmetric info benchmark	0.99	0.89	0.63
No TFP-driven confidence	0.57	0.51	0.36
No sentiment-driven confidence	0.03	0.03	0.02
No demand uncertainty	0.03	0.03	0.02

Notes: The table shows the share of output that can be accounted by the intrinsic shocks to the informational component of the estimated wedges, u_t^{info}. The contribution of the productivity shock to $\operatorname{var}[y_t]$ and $\operatorname{var}[\hat{y}_t]$ is 11 and 36 percent, respectively. All variance contributions are computed at business cycle frequencies for the baseline calibration of $\{\Delta a_{i,t}\}$ and $\{z_{i,t}\}$ (i.e., scale $=1$ in Figure 3).

fluctuations in confidence can account for 51 percent of the empirical volatility. Importantly, however, for a theory of incomplete information to generate significant fluctuations in confidence, firms must face some uncertainty about their *idiosyncratic* product demands. If this is not the case, then confidence fluctuations can at most explain 3 percent of the empirical volatility in y_t.

V. Taking Stock

We have developed a method to quantify the potential of DSGE models with imperfect information without taking a fully structural stand on the private information of agents. Along the way, we established a *conditional* equivalence, which holds under the conditions of Theorem 1, between models with dispersed information and a prototype wedge economy similar to the one in Chari, Kehoe, and McGrattan (2007). The informational foundation for these wedges is distinguished from existing theories in its ability to generate arbitrary correlation patterns between these wedges (Proposition 6). Correlated wedges, in turn, are critical for the empirical viability of confidence fluctuations because the data imply a strong correlation between the aggregate labor wedge and the Euler wedge.

Expectations are a natural candidate for generating the observed correlation, both because information can be correlated between households and firms and because expectation errors by households generally affect both their consumption and labor supply. Our results indicate, however, that two features are crucial to achieve a quantitively important role for such a foundation: (i) micro-shocks must be sufficiently volatile and (ii) idiosyncratic demand must be uncertain at the time of production choices. Regarding (i), our analysis suggests that observed micro-level volatility is indeed large enough to support substantial aggregate volatility. Regarding (ii), the presence of idiosyncratic demand uncertainties has long been acknowledged in business practices (Fisher et al. 1994) and in operations research (Fisher and Raman 1996; Mula et al. 2006). Yet, given the pivotal role that these uncertainties may play in supporting aggregate fluctuations, our results suggest to us that further research is warranted regarding the degree to which firms misperceive their own demand shocks when making input choices.

Appendix A. Proof of Main Theorem

Consider any expectation wedge $\tau_{i,t}^{j} \in \mathcal{T}_t$ from the primal economy and the corresponding lower bound $\Theta_{i,t}^{j}$ on $\mathcal{I}_{i,t}^{j}$ in the incomplete information economy. Define the expectation "targets"

$$a_{i,t}^{j} \equiv \mathbf{A}_1^{j} g_{i,t+1} + \mathbf{A}_2^{j} f_{i,t+1} + \mathbf{B}_1^{j} g_{i,t} + \mathbf{B}_2^{j} f_{i,t},$$

as pinned down by the equilibrium $\mathcal{E} \in \mathcal{E}^{\text{primal}}(\mathcal{F},\mathcal{T})$ of the primal economy.

We want to show that conditions (i) and (ii) are jointly necessary and sufficient for the construction of some $\mathcal{I}_{i,t}^{j} \supseteq \mathcal{S}_{i,t}^{j} \equiv \left\{\mu_{i,t-s}^{j}, \Theta_{i,t-s}^{j}\right\}_{s\geq 0}$ such that

$$\mathbb{E}\left[a_{i,t}^{j}|\mathcal{I}_{i,t}^{j}\right] = \mathbb{E}\left[a_{i,t}^{j}|\mathcal{I}_{t}^{*}\right] + \tau_{i,t}^{j}. \tag{A1}$$

When this is true, any solution to (2) is trivially also a solution to (1).

To conserve notation, we suppress (i,j) subscripts going forward.

Necessity.—Necessity is immediate, since optimal inference requires that expectation errors are orthogonal to variables in the information set and are unpredictable. To see this, rearrange (A1) to get

$$\tau_t = \mathbb{E}\left[a_t|\mathcal{I}_t\right] - \mathbb{E}\left[a_t|\mathcal{I}_t^{*}\right]. \tag{A2}$$

Computing the unconditional expectation over (A2) yields $\mathbb{E}[\tau_t] = 0$. Similarly, postmultiplying (A2) by μ_t and $\theta_t \in \Theta_t$ gives

$$\mathbb{E}\left[\tau_t\mu_t\right] = \mathbb{E}\left[a_t\mu_t|\mathcal{I}_t\right] - \mathbb{E}\left[a_t\mu_t|\mathcal{I}_t^{*}\right]$$

$$\mathbb{E}\left[\tau_t\theta_t\right] = \mathbb{E}\left[a_t\theta_t|\mathcal{I}_t\right] - \mathbb{E}\left[a_t\theta_t|\mathcal{I}_t^{*}\right],$$

as $\theta_t \subseteq \mathcal{I}_t \subseteq \mathcal{I}_t^{*}$. Again taking the unconditional expectation over the right-hand sides, we have $\mathbb{E}[\tau_t\mu_t] = \mathbb{E}[\tau_t\theta_t] = 0$ for all $\theta_t \in \Theta_t$.

Sufficiency.—We demonstrate sufficiency by construction. Let $\hat{a}_t \equiv \mathbb{E}\left[a_t|\mathcal{I}_t^{*}\right]$, and consider the information set $\mathcal{I}_t = \mathcal{S}_t \cup \left\{s_{t-\tau}\right\}_{\tau\geq 0}$, where $s_t \equiv \hat{a}_t + \tau_t = \mu_t$ is a signal that replicates the correlation structure of the expectation we wish to implement. Notice that $\mathcal{I}_t$ inherits recursiveness from $\mathcal{S}_t$, ensuring consistency with Assumption 2.

From the law of iterated expectations, we have $\mathbb{E}[a_t|s_t] = \mathbb{E}[\hat{a}_t|s_t]$ as $s_t \subseteq \mathcal{I}_t^{*}$. Projecting $\hat{a}_t$ onto s_t, we obtain

$$\begin{aligned}\mathbb{E}\left[a_t|s_t\right] &= \operatorname{cov}\left[\hat{a}_t, s_t\right]\operatorname{var}\left[s_t\right]^{-1} s_t \\ &= \operatorname{cov}\left[s_t - \tau_t, s_t\right]\operatorname{var}\left[s_t\right]^{-1} s_t \\ &= \operatorname{var}\left[s_t\right]\operatorname{var}\left[s_t\right]^{-1} s_t \\ &= s_t,\end{aligned} \tag{A3}$$

where the second line follows from the definition of s_t and the third line follows from condition (ii) of the Theorem and the fact that $s_t = \mu_t \in \mathcal{S}_t$. Noting that by construction no other $\theta_t \in \mathcal{S}_t$ can improve the forecast about a_t,[28] we obtain

$$\mathbb{E}\left[a_t|s_t\right] = \mathbb{E}\left[a_t|\mathcal{I}_t\right] = \mathbb{E}\left[a_t|\mathcal{I}_t^*\right] + \tau_t.$$

As the argument above applies to any $\tau_{i,t}^j \in \mathcal{T}$, we have constructed exactly the information sets needed to satisfy (A1) for all (i,j,t).

REFERENCES

Acharya, Sushant. 2013. "Dispersed Beliefs and Aggregate Demand Management." Unpublished.

Acharya, Sushant, Jess Benhabib, and Zhen Huo. 2018. "The Anatomy of Sentiment-Driven Fluctuations." NBER Working Paper 23136.

Adams, Jonathan J. 2022. "Macroeconomic Models with Incomplete Information and Endogenous Signals." Unpublished.

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas. 2018. "Quantifying Confidence." *Econometrica* 86 (5): 1689–1726.

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas. 2020. "Business-Cycle Anatomy." *American Economic Review* 110 (10): 3030–70.

Angeletos, George-Marios, and Jennifer La'O. 2010. "Noisy Business Cycles." In *NBER Macroeconomics Annual*, Vol. 24, edited by Daron Aceboglu, Kenneth Rogoff, and Michael Woodford, 319–78. Chicago, IL: University of Chicago Press.

Angeletos, George-Marios, and Jennifer La'O. 2013. "Sentiments." *Econometrica* 81 (2): 739–79.

Benhabib, Jess, Pengfei Wang, and Yi Wen. 2015. "Sentiments and Aggregate Demand Fluctuations." *Econometrica* 83 (2): 549–85.

Bergemann, Dirk, Tibor Heumann, and Stephen Morris. 2015. "Information and Volatility." *Journal of Economic Theory* 158 (B): 427–65.

Bergemann, Dirk, and Stephen Morris. 2013. "Robust Predictions in Games with Incomplete Information." *Econometrica* 81 (4): 1251–1308.

Bergemann, Dirk, and Stephen Morris. 2016. "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games." *Theoretical Economics* 11 (2): 487–522.

Blanchard, Olivier J., Jean-Paul L'Huillier, and Guido Lorenzoni. 2013. "News, Noise, and Fluctuations: An Empirical Exploration." *American Economic Review* 103 (7): 3045–70.

Chahrour, Ryan, and Gaetano Gaballo. 2021. "Learning from House Prices: Amplification and Business Fluctuations." *Review of Economic Studies* 88 (4): 1720–59.

Chahrour, Ryan, and Robert Ulbricht. 2023. " Replication data for: Robust Predictions for DSGE Models with Incomplete Information." American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. https://doi.org/10.38886/E148421V1.

Chari, V.V., Patrick J. Kehoe, and Ellen R. McGrattan. 2007. "Business Cycle Accounting." *Econometrica* 75 (3): 781–836.

Cochrane, John H. 1994. "Shocks." *Carnegie-Rochester Conference Series on Public Policy* 41: 295–364.

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.

De Loecker, Jan. 2011. "Product Differentiation, Multiproduct Firms, and Estimating the Impact of Trade Liberalization on Productivity." *Econometrica* 79 (5): 1407–51.

Demidova, Svetlana, Hiau Looi Kee, and Kala Krishna. 2012. "Do Trade Policy Differences Induce Sorting? Theory and Evidence from Bangladeshi Apparel Exporters." *Journal of International Economics* 87 (2): 247–61.

[28] To see this, note that the forecast error conditional on s_t is necessarily uncorrelated with any other $\theta_t \in \mathcal{S}_t$: $\mathrm{cov}\left[a_t - \mathbb{E}\{a_t|s_t\}, \theta_t\right] = \mathrm{cov}\left[a_t - s_t, \theta_t\right] = \mathrm{cov}\left[a_t - \hat{a}_t - \tau_t, \theta\right] = \mathrm{cov}\left[-\tau_t, \theta_t\right] = 0$. Here, the first equality follows from (47); the second one follows per the definition of τ_t; the third one follows because $a_t - \hat{a}_t$ defines the forecast error under full information $\mathcal{I}_t^*$, so that any $\theta_t \in \mathcal{S}_t \subset \mathcal{I}_t^*$ must be orthogonal to it; and the last equality follows from the conditions of the theorem.

Fisher, Marshall L., Janice H. Hammond, Walter R. Obermeyer, and Ananth Raman. 1994. "Making Supply Meet Demand in an Uncertain World." *Harvard Business Review* 72 (3): 83–93.

Fisher, Marshall, and Ananth Raman. 1996. "Reducing the Cost of Demand Uncertainty through Accurate Response to Early Sales." *Operations Research* 44 (1): 87–99.

Foster, Lucia, John Haltiwanger, and Chad Syverson. 2008. "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?" *American Economic Review* 98 (1): 394–425.

Foster, Lucia, John Haltiwanger, and Chad Syverson. 2016. "The Slow Growth of New Plants: Learning about Demand?" *Economica* 83: 91–129.

Grossman, Sanford J., and Joseph E. Stiglitz. 1980. "On the Impossibility of Informationally Efficient Markets." *American Economic Review* 70 (3): 393–408.

Han, Zhao, Fei Tan, and Jieran Wu. 2021. "Analytic Policy Function Iteration." Unpublished.

Hansen, Lars Peter, and Thomas J. Sargent. 1980. "Formulating and Estimating Dynamic Linear Rational Expectations Models." *Journal of Economic Dynamics and Control* 2: 7–46.

Hansen, Lars Peter, and Thomas J. Sargent. 1981. "A Note on Wiener-Kolmogorov Prediction Formulas for Rational Expectations Models." *Economics Letters* 8 (3): 255–60.

Hellwig, Martin F. 1980. "On the Aggregation of Information in Competitive Markets." *Journal of Economic Theory* 22 (3): 477–98.

Hellwig, Christian, and Venky Venkateswaran. 2014. "Dispersed Information, Sticky Prices and Monetary Business Cycles: A Hayekian Perspective." Unpublished.

Huo, Zhen, and Naoki Takayama. 2015. "Higher Order Beliefs, Confidence, and Business Cycles." Unpublished.

Huo, Zhen, and Naoki Takayama. 2021. "Rational Expectations Models with Higher-Order Beliefs." Unpublished.

Ilut, Cosmin, and Hikaru Saijo. 2021. "Learning, Confidence, and Business Cycles." *Journal of Monetary Economics* 117: 354–76.

Lorenzoni, Guido. 2009. "A Theory of Demand Shocks." *American Economic Review* 99 (5): 2050–84.

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

Lucas, Robert E., Jr. 1973. "Some International Evidence on Output-Inflation Tradeoffs." *American Economic Review* 63 (3): 326–34.

Maćkowiak, Bartosz, and Mirko Wiederholt. 2015. "Business Cycle Dynamics under Rational Inattention." *Review of Economic Studies* 82 (4): 1502–32.

Melosi, Leonardo. 2014. "Estimating Models with Dispersed Information." *American Economic Journal: Macroeconomics* 6 (1): 1–31.

Melosi, Leonardo. 2017. "Signaling Effects of Monetary Policy." *Review of Economic Studies* 84 (2): 853–84.

Mondria, Jordi, Xavier Vives, and Liyan Yang. 2022. "Costly Interpretation of Asset Prices." *Management Science* 68 (1): 52–74.

Morris, Stephen, and Hyun Song Shin. 2002. "Social Value of Public Information." *American Economic Review* 92 (5): 1521–34.

Mula, J., R. Poler, J.P. García-Sabater, and F.C. Lario. 2006. "Models for Production Planning under Uncertainty: A Review." *International Journal of Production Economics* 103 (1): 271–85.

Nimark, Kristoffer P. 2017. "Dynamic Higher Order Expectations." Unpublished.

Passadore, Juan, and Juan Pablo Xandri. 2021. "Robust Predictions in Dynamic Policy Games." Unpublished.

Roberts, Mark J., Daniel Yi Xu, Xiaoyan Fan, and Shengxing Zhang. 2018. "The Role of Firm Factors in Demand, Cost, and Export Market Selection for Chinese Footwear Producers." *Review of Economic Studies* 85 (4): 2429–61.

Rondina, Giacomo, and Todd B. Walker. 2021. "Confounding Dynamics." *Journal of Economic Theory* 196: 105251.

Sims, Christopher A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50 (3): 665–90.

Woodford, Michael. 2003. "Imperfect Common Knowledge and the Effects of Monetary Policy." In *Knowledge, Information and Expectations in Modern Macroeconomics—In Honor of Edmund S. Phelps*, edited by Philippe Aghion, Roman Frydman, Joseph Stiglitz and Michael Woodford, 25–58. Princeton, NJ: Princeton University Press.