

Robust Predictions for DSGE Models with Incomplete Information*

Ryan Chahrour Robert Ulbricht
Boston College Boston College

June 10, 2021

Abstract

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.

Keywords: Business cycles, DSGE models, incomplete-information, information-robust predictions.

JEL Classification: E32, D84.

*We thank our discussants Guillaume Chevillon, Stephane Dupraz, Alexandre Kohlhas, Jianjun Miao and Anna Orlik for many valuable comments. Ulbricht gratefully acknowledges financial support from the European Union's Horizon 2020 Program under grant agreement No. 649396. *Email addresses:* ryan.chahrour@bc.edu, ulbricht@bc.edu.

1 Introduction

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. At least in part, this is 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 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 who 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* autocovariance struc-

ture 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 micro-data 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 U.S. 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 ag-

ported 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 U.S. 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 U.S. 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 non-informational residual. We find that, in principle, confidence-fluctuations can account for a large portion of the U.S. 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 \(2014\)](#). 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

2 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 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.

2.1 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 $\mathcal{I}_{i,t}^j$. Here, $j \in \{0, 1, \dots, J\}$ 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 $g_{i,t} = [g_{i,t}; g_t^a]$, where $g_{i,t}$ denotes a $n \times 1$ vector of purely atomistic endogenous variables that satisfy the adding-up constraint $\int_0^1 g_{i,t} di = 0$, and g_t^a denotes a $n_{ga} \times 1$ vector of endogenous aggregate variables (which may but are not limited to include the "mean component" of $g_{i,t}$).

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

$$0 = \int_0^1 E_{\mathcal{I}_{i,t}^j} \begin{pmatrix} A_1^j & A_2^j \\ B_1^j & B_2^j \end{pmatrix} \begin{pmatrix} g_{i,t+1} \\ f_{i,t+1} \end{pmatrix} + \begin{pmatrix} A_1^j & A_2^j \\ B_1^j & B_2^j \end{pmatrix} \begin{pmatrix} g_{i,t} \\ f_{i,t} \end{pmatrix} + l_{i,t}^j; \quad (1)$$

for all $i \in [0, 1]$ and $t = 0, 1, \dots$. Here, $f_{i,t} = [f_{i,t}; f_t^a]$ is an exogenous column vector of stochastic variables. In analogy to the endogenous vector $g_{i,t}$, we partition the exogenous vector into an atomistic component, $f_{i,t}$, and an aggregate component, f_t^a , where the atom-

³Here, 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.

⁴Notice that 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.

istic component satisfies the adding up constraint $\int_0^1 f_{i,t} di = 0$. 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, $\{f, g\}_{i,t}$

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

Assumption 1 (Information bounds). $I_{i;t}^j \leq I_{i;t}^j \leq I_t$.

Assumption 1 defines a lower and an upper bound on information. The upper bound, I_t , simply states that agents cannot learn more than what is potentially knowable under full information. The lower bound, $I_{i;t}^j$, 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 restricts equilibrium outcomes.

Assumption 2 (Recursiveness). $I_{i;t-1}^j \leq I_{i;t}^j$.

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

To state the theorem, define

$$I_{i;t}^j = E_t[A_1^j g_{i;t+1} + A_2^j f_{i;t+1} + B_1^j g_{i;t} + B_2^j f_{i;t}] + I_{i;t}^j$$

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

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

$$E[I_{i;t}^j] = 0 \text{ for all } i, j, t \text{ and } s: I_{i;t}^j \perp I_{i;t}^s \text{ for } s \neq j \quad (3)$$

hold for i, j , and t .

The theorem gives two conditions that are jointly necessary and sufficient for 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 $I_{i;t}^j$ and $I_{i;t}^s$. The necessity of this restriction is the familiar result that expectation errors must be orthogonal to all available information, including an agent's belief $I_{i;t}^j$ itself (at the very least "one knows what one knows"). The novel part of our result is the sufficiency of this condition. For any $E \in E^{\text{primal}}(F; T)$ with $E[T_t] = 0$, we can always construct an information structure that implements E as an incomplete-information equilibrium as long as it satisfies (3).

⁶Specifically, in this case, we obtain a version of our theorem, in which condition (3) is imposed only for $s = 0$.

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

$$p_{i;t} = E[p_t + \eta y_t + \zeta z_{i;t} | I_{i;t}]; \quad (8)$$

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

$$y_{i;t} = \frac{1}{\sigma} (\frac{p_{i;t}}{p_t})^{1-\sigma} y_t + \zeta z_{i;t}; \quad (9)$$

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

$$q_t = y_t + p_t; \quad (10)$$

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

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

$$p_{i;t} = p_t + \eta y_t + \zeta z_{i;t} + \epsilon_{i;t} \quad (11)$$

along with equations (9) and (10).

Given a process $\zeta z_{i;t}$, the equilibrium of the primal economy is straightforward to find. Defining $\epsilon_t = \int_0^1 \epsilon_{i;t} di$, aggregates in the economy are given by

$$p_t = q_t + \epsilon_t; \quad y_t = \frac{1}{\sigma} \epsilon_t; \quad (12)$$

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, $y_{i,t} \geq z_{i,t}$, 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 co-movement between output, inflation and money growth.

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

$$\rho_{\text{Var}[y_t]} \geq (1 - \alpha) \frac{\text{Corr}[y_t; \pi_t]}{1 - \text{Corr}[y_t; y_{t-1}]} \rho_{\text{Var}[\pi_t]}$$

and

$$\rho_{\text{Var}[y_t]} \geq \frac{(1 - \alpha)}{(1 - \alpha) + \alpha} \frac{\text{Corr}[y_t; dq_t]}{1 - \text{Corr}[y_t; y_{t-1}]} \rho_{\text{Var}[dq_t]}.$$

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 π_t , and differencing the resulting conditions, we have

$$\text{Cov}[y_t; dy_t] = (1 - \alpha) \text{Cov}[y_t; \pi_t] = \alpha \text{Cov}[z_{i,t}; d z_{i,t}].$$

Noting that $\text{Cov}[z_{i,t}; d z_{i,t}] = (1 - \text{Corr}[z_{i,t}; z_{i,t-1}]) \text{Var}[z_{i,t}] \geq 0$ completes the proof. \square

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 $z_{i,t}$.

Proposition 3. *Suppose $f(z_{i,t}; y_{i,t}) \geq z_{i,t}$. Then aggregate output is constant.*

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

$$\text{Cov}[z_{i,t}; y_{t-s} + p_{t-s}] = \text{Cov}[z_{i,t}; z_{i,t-s}]. \quad (18)$$

From $\sum_{i,t} Z_{i,t} di = 0$, it follows that $\text{Cov}[Z_{i,t}; Z_{i,t-s}] = \text{Cov}[Z_{i,t}; Z_{i,t-s}]$. Applying Theorem 1, it then must hold that $\text{Cov}[Z_{i,t}; Z_{i,t-s}] = 0$. Evaluating (18) at $s = 0$ and $s = 1$, using (12) to substitute for $Z_{i,t}$

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

$$E_{i,t} \sum_{s=0}^{\infty} \beta^s U(C_{i;t+s}; N_{i;t+s}) | I_{i,t}^h ;$$

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

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

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

$$P_t C_{i;t} + Q_t B_{i;t} = W_{i;t} N_{i;t} + B_{i;t-1} + D_{i;t} ;$$

where P_t is the price of the final good, Q_t

specific component,

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

where the aggregate component follows a random walk process

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

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

Final-good sector A competitive final-goods sector aggregates intermediate input goods $i \in [0; 1]$, using the technology

$$Y_t = \left(\int_0^1 Z_{i;t} Y_{i;t}^{\frac{1}{\sigma}} di \right)^{\sigma};$$

where $\sigma > 1$ 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

$$P_{i;t} = \frac{Y_{i;t}}{Y_t} P_t^{\sigma} Z_{i;t} P_t \quad (20)$$

where the aggregate price index P_t is defined by

$$P_t = \left(\int_0^1 Z_{i;t} P_{i;t}^{\frac{1}{\sigma}} di \right)^{\sigma};$$

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

$$\dot{i}_t = \rho + \eta \pi_t \quad (21)$$

^R₀¹ $\int_0^1 Y_{ij;t}^{\frac{1}{\sigma}} dj = 1$ where σ matches the elasticity of substitution across "island-varieties" specified in the

where $\alpha > 1$ and $i_t = \log(Q_t)$.⁹

Information structure Our methodology allows us to explore how a few abstract assumptions regarding $f_{i,t}^j, g_{i,j} \in [0,1]$ f, h, g

3.2 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. Lower-case 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

$$c_{i,t} = E[c_{i,t+1} - \pi_{t+1} | \mathcal{H}_{i,t}^h]: \quad (25)$$

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

$$y_{i,t} = y_{i,t} - c_{i,t} + E[p_{i,t} | \mathcal{H}_{i,t}^f] - E[p_t | \mathcal{H}_{i,t}^h] + a_{i,t}; \quad (26)$$

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

$$p_{i,t} = \pi_t (y_t - y_{i,t}) + z_{i,t} + p_t \quad (27)$$

and

$$b_{i,t} = b_{i,t-1} + y_{i,t} - c_{i,t} + p_{i,t} - p_t; \quad (28)$$

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

$$y_t = \int_0^1 y_{i,t} di = \int_0^1 c_{i,t} di; \quad (29)$$

(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

In keeping with the literature, we do not include the real return on assets, $r_t = i_t E_t[r_{t+1}]$, or its constituents i_t , p_t and $E_t[p_{t+1}]$, in the lower bound on households' information $f|_{i;t}^h g$. 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 such that the average expected real interest $E_t[r_t]$ increases with r_t . Using our methodology, there is no need to explicitly specify the signals through which households make inference about r_t . Instead, requiring market clearing in the primal representation of the economy yields by construction a "market clearing expectation" $E_t[r_t]$ that adapts to clear the goods market in all states of the world.¹²

To see this, consider the simplified case where aggregate demand is given by $c_t = E[r_t^j | I_t]$ and aggregate supply, y_t , follows an exogenous random process. In this case, market clearing ($c_t = y_t$) requires

$$E[r_t^j | I_t] = y_t \quad (30)$$

which in conjunction with i_t pins down r_t : In the primal representation, $E[r_t^j | I_t] = r_t + \varepsilon_t$, and market clearing requires

$$r_t + \varepsilon_t = y_t \quad (31)$$

The key difference between (30) and (31) is that the expectation error, ε_t , is a primitive of the primal economy. Because ε_t is exogenous in the primal economy, the solution $r_t = y_t - \varepsilon_t$ always imposes that the implied $E[r_t^j | I_t]$ responds one-for-one to a decline in y_t , inducing precisely the sensitivity of households' expectations to economic conditions that is necessary for r_t to clear the goods market. Hence, by imposing market clearing in the primal economy, we implicitly require that agents have enough information

3.3 Primal Representation

There are two equilibrium conditions with non-trivial expectation operators. Replacing equations (25) and (26) with their primal analog, we arrive at¹³

$$c_{i,t} = E$$

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

$$E_t \left[f d c_{i,t}; d y_{i,t}; d b_{i,t}; d p_{i,t} g_{i,t} \right] \in E(F; T):$$

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

$$E \left[\begin{matrix} c_{i,t} \\ p_{i,t}^h \end{matrix} \right] = E \left[\begin{matrix} p_{i,t}^h \\ p_{i,t}^f \end{matrix} \right] = 0 \quad \text{for all } \sum_{s=0}^h g_{s=0}^h; \text{ and}$$

$$E \left[\begin{matrix} p_{i,t}^f \end{matrix} \right] = 0 \quad \text{for all } \sum_{s=0}^h g_{s=0}^h$$

hold for all i and t .

Proposition 5 is an immediate corollary to Theorem 1. Here, the restriction to finite MA processes arises because /

captured by $\begin{pmatrix} p \\ t \end{pmatrix} = \begin{pmatrix} p:f \\ t \end{pmatrix} + \begin{pmatrix} p:h \\ t \end{pmatrix}$, which corresponds to the labor wedge in our economy that is composed of a household and a firm component. The aggregate "wedges" $\begin{pmatrix} c \\ t \end{pmatrix}$ and $\begin{pmatrix} p \\ t \end{pmatrix}$ are the sole drivers of the output gap and inflation. If all agents had full information ($\begin{pmatrix} c \\ t \end{pmatrix} = \begin{pmatrix} p \\ t \end{pmatrix} = 0$), 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 $\begin{pmatrix} c \\ t; p \\ t \end{pmatrix}^\theta$ can be obtained using standard numerical tools. In our case, a closed-form solution is also available. Substituting for \hat{y}_t in (34), $\begin{pmatrix} c \\ t \end{pmatrix}$ is characterized by the prediction formula

$$\begin{pmatrix} c \\ t \end{pmatrix} = \mathbb{E}_t[\begin{pmatrix} p \\ t+1 \end{pmatrix} + \begin{pmatrix} c \\ t+1 \end{pmatrix} + \begin{pmatrix} p \\ t+1 \end{pmatrix}]. \quad (36)$$

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

Lemma 1. *Let $\begin{pmatrix} c \\ t \end{pmatrix} = A(L)u_t$, where $A(L)$ is a square-summable lag polynomial in non-negative powers of L and the innovations u_t are orthogonal white noise. Then there exists a unique stationary equilibrium process for $(\hat{y}_t; \begin{pmatrix} c \\ t \end{pmatrix})$, given by*

$$\hat{y}_t = \sum_{i=0}^{\infty} A(L)u_t \quad (37)$$

and

$$\begin{pmatrix} c \\ t \end{pmatrix} = \sum_{i=0}^{\infty} \frac{(1-L)^{-1}A(L)(1-L)^{-1}A(L)^{-1}}{L^{-1}}u_t. \quad (38)$$

4 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 4.1 maps the restrictions stated in Proposition 5 into restrictions on the dynamics of the "macro" wedges determining the behavior of the aggregate economy. Sections 4.2 and 4.3 then use these restrictions on the macro wedges to characterize feasible volatility and co-movement patterns of output and inflation under varying assumptions on information and fundamentals.

4.1 Feasible Dynamics of Aggregate Wedges

We begin by mapping the orthogonality restrictions in Proposition 5 into restrictions on the macro wedges $\begin{pmatrix} c \\ t \end{pmatrix}$ and $\begin{pmatrix} p \\ t \end{pmatrix}$. To streamline the exposition, we only detail the derivation

for the baseline case $\text{sym}_{i;t}$ depicted in (22), in which firms and households have symmetric information.

To begin, observe that for $\text{sym}_{i;t}, f_{i;t}, s_{i;t}, \text{sym}_{i;t}, g_{s=0}$ satisfies Assumption 3 with

$$S_{i;t} = f_{i;t} dc_{i;t} + dy_{i;t} + da_{i;t} g_{s=0}$$

Here we have used that (i) $n_{i;t}$ and $w_{i;t}$ are linear combinations of $(c_{i;t}, y_{i;t}, a_{i;t})$ and are therefore informationally redundant; and (ii) that for any finite horizon h , observing the sequence of differences $f_{i;t} S_{i;t} g_{s=0}^h$ in addition to l_{t+h} contains the same information as the corresponding sequence of levels.

To proceed, define $\epsilon_{i;t} = (c_{i;t}, p_{i;t}^h, p_{i;t}^h)^\theta$ and let $\epsilon_{i;t} = \epsilon_{i;t} + \epsilon_{i;t}$ denote the idiosyncratic portion of the expectation wedges. Similarly, let $(c_{i;t}, y_{i;t})$ denote the idiosyncratic deviations from aggregate output. By construction the Δ -component of any variable is

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) $MA(h)$ process for the idiosyncratic and aggregate components of $\tilde{y}_{i,t}$, 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 $\tilde{y}_t^c: \tilde{y}_t^p$ for which (39) can be satisfied with some processes for the idiosyncratic components $\tilde{y}_{i,t}^c: \tilde{y}_{i,t}^p$. 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.)

4.2 Unrestricted Micro-Shock Benchmark

Before proceeding to our quantitative results, we provide a theoretical benchmark for the case where we treat the idiosyncratic fundamentals, $\tilde{f}_{i,t} = (a_{i,t}; z_{i,t})$, 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 \tilde{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 $\tilde{f}_{i,t}$. By construction, the chosen process for $\tilde{f}_{i,t}$ 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) $MA(h)$ process for $(\tilde{y}_t^c; \tilde{y}_t^p)$ and set $\tilde{y}_{i,t}^{\text{sym}}$ as in (22). Then for any aggregate productivity process, a , there exist idiosyncratic processes $\tilde{y}_{i,t}^c$ and $\tilde{y}_{i,t}^p$, such that \tilde{y}_t can be implemented in the incomplete information economy.*

Proposition 6 provides a striking benchmark: Absent micro-data that disciplines $\tilde{f}_{i,t}$, correlated optimism and pessimism (across islands), can be used to generate *any* joint process in $(\tilde{y}_t^c; \tilde{y}_t^p)$. 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 \tilde{y}_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.

4.3 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 $\hat{f}_{i,t}$ at an empirically plausible calibration? We explore this question numerically, calibrating $\hat{f}_{i,t}$ to existing micro-data.

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

Next, we set the incomplete information horizon to $h - 1 = 14$ quarters. While we do not have strong priors on $\hat{f}_{i,t}$,

$z_{i,t}$ and $a_{i,t}$ 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 Loecker 2011; Demidova, Kee and Krishna 2012; Roberts et al. 2017; 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 σ_y as a function of its persistence and the cyclical nature of inflation σ_π that our model can generate in the absence of aggregate shocks to fundamentals ($\text{Var}[\epsilon_t] = 0$).


Formally, define $\sigma_y(\rho)$ as the one-step-ahead volatility of output induced by ϵ_t . Similarly, define $\sigma_y(\rho) = \text{Corr}[\hat{y}_t, \hat{y}_{t-1}]$ as the first-order autocorrelation of \hat{y}_t , and define $\sigma_y(\rho) = \text{Corr}[\hat{y}_t, \hat{\pi}_t]$ as the contemporaneous correlation with inflation. We use Lemma 2 to numerically trace out the *volatility frontier* for output as a function of its autocorrelation σ_y and its contemporaneous correlation with inflation σ_π :

$$\max_{\sigma_y}(\sigma_\pi; \sigma_y) = \max_{\rho} f(\sigma_y(\rho))g(\sigma_\pi(\rho))$$

subject to

$$\begin{aligned} \sigma_y(\rho) &= \sigma_y \\ \sigma_\pi(\rho) &= \sigma_\pi \end{aligned}$$

and the implementability condition (39). Here σ_y and σ_π are independent (zero-mean) MA(h) processes.



gure.

The sensitivity is strongest in α_z and α_z , indicating that correlated expectation errors about the demand shocks $fz_{i,t}g$ are of critical importance for supporting fluctuations in aggregate confidence. In particular, a reduction in α_z from its baseline value of 0.2504 to 0.01, reduces \bar{y}^{\max} by a factor of three to 0.37 percent; an increase in α_z to 1.00, increases \bar{y}^{\max} to 3.39 percent. Those comparative statics reflect the naturally increasing shape of \bar{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 ($\text{Var}[E\{fz_{i,t}j|I_{i,t}g] = \text{Var}[z_{i,t}]$). In the extreme case where $\alpha_z \neq 0$, rationality requires that $E[z_{i,t}|I_{i,t}] = 0$ for all t , even if $I_{i,t}$ contains no information about $z_{i,t}$.

Similarly to α_z , variations in the persistence of $z_{i,t}$ also have a significant impact on \bar{y}^{\max} : a reduction of α_z from its baseline value of 0.976 to 0.5, reduces \bar{y}^{\max} to 0.35 percent. An increase in the persistence of $z_{i,t}$ to 0.99, increases \bar{y}^{\max} to 3.18. The role of α_z for supporting expectation errors is two-fold. First, $\text{Var}[z_{i,t}]$ is increasing in α_z , again increasing the potential for expectation errors. Second, persistence in $z_{i,t}$ (or in $a_{i,t}$), 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 cyclicity of inflation ($\bar{y} = 0.3$).¹⁹

By contrast, variations in the parameters of $fa_{i,t}g$ result in only moderate variations in \bar{y}^{\max} . In particular, reducing α_x or α_l to 0.01, implies only marginally smaller values of \bar{y}^{\max} , suggesting that the idiosyncratic productivity shocks $fa_{i,t}g$ 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 $a_{i,t} \geq i_{i,t}$, firms and households always know their current productivity, limiting productivity-related confusion to uncertainty about the composition of $a_{i,t}$, 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

¹⁹In order to generate pro-cyclical 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 $E[z_{i,t}|I_{i,t}]$ for determining labor demand.

a particular choice of production. Formally, information is now bounded by

$$i:t = \bar{p}_{i:t} \text{sgs}_0 \left[\begin{array}{c} \text{sym} \\ i:t \end{array} \right]$$

with $\begin{array}{c} \text{sym} \\ i:t \end{array}$ given by (22). Because $\begin{array}{c} p:f \\ i:t \end{array}$ measures firms' expectation error regarding $p_{i,t}$, an immediate consequence of including $p_{i,t}$ in $\begin{array}{c} f \\ i:t \end{array}$ is that $\begin{array}{c} p:f \\ i:t \end{array} = 0$ 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, $w_{i,t}$ and $a_{i,t}$, 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 $\begin{array}{c} f \\ i:t \end{array} z_{i,t}g$, shutting down $\begin{array}{c} p:f \\ i:t \end{array}$ reduces \bar{y}^{\max} to 0.41, suggesting that uncertainty about demand is key to generating sizable aggregate fluctuations. Moreover, compared to the case where $\begin{array}{c} \text{sym} \\ i:t \end{array}$ is given by (22), the sensitivity of \bar{y}^{\max} in the parameters of $\begin{array}{c} f \\ i:t \end{array} z_{i,t}g$ is reduced, whereas the sensitivity in the parameters of $\begin{array}{c} f \\ i:t \end{array} a_{i,t}g$ 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 $z_{i,t} + p_t^{-1}y_t$ from (20), reducing the scope to generate waves of optimism and pessimism via $z_{i,t}$ and, by implication, increasing the model's reliance on $a_{i,t}$ for supporting aggregate fluctuations in confidence.²⁰

Heterogeneous information We next relax the assumption that households and firms share the same information set, setting $\begin{array}{c} h \\ i:t \end{array}$ and $\begin{array}{c} f \\ i:t \end{array}$ as in (23) and (24). The resulting volatility frontier is depicted by the red lines in Figure 2. For the baseline calibration, this increases \bar{y}^{\max} to 4.49 percent. This reflects the additional flexibility in $\begin{array}{c} f \\ i:t \end{array}$ and $\begin{array}{c} h \\ i:t \end{array}$, due to households not being required to perfectly know the local firm's productivity (i.e., $a_{i,t}; y_{i,t} \not\subseteq \begin{array}{c} h \\ i:t \end{array}$) and firms not being required to perfectly know households' consumption ($c_{i,t} \not\subseteq \begin{array}{c} f \\ i:t \end{array}$). Specifically, this enables waves of optimism and pessimism among households about income fluctuations caused by $a_{i,t}$ and $z_{i,t}$, translating to aggregate demand fluctuations| even if $a_{i,t}$ and $z_{i,t}$ are observed by firms. The stark increase in \bar{y}^{\max} 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 ($\begin{array}{c} f \\ i:t \end{array}$ includes $\bar{p}_{i:t} \text{sgs}_0$ in addition to (24)). The results are depicted by the blue lines in Figure 2). Compared to the symmetric-information case without demand uncertainty, \bar{y}^{\max} is slightly increased to 0.49. However, the difference between symmetric

²⁰The 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 $z \neq 0$, which affects local wealth and households' consumption choice.

Table 1: Summary of estimated U.S. wedges

Standard deviation	First-order autocorr.	Contemporaneous correlation	
		with \hat{c}_t^c	with \hat{c}_t^p

5.1.2 Partitioning of the estimated wedges

We partition the estimated wedge process \hat{a}_t into an informational component \hat{a}_t^{info} and a residual component \hat{a}_t^{resid} ,

$$\hat{a}_t = \hat{a}_t^{\text{info}} + \hat{a}_t^{\text{resid}}. \quad (42)$$

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

$$\begin{aligned} \hat{a}_t^{\text{info}} &= \text{info}(L) \hat{a}_t^{\text{info}} + u_t^{\text{info}}(L) u_t^{\text{info}} \\ \hat{a}_t^{\text{resid}} &= \text{resid}(L) \hat{a}_t^{\text{resid}} + u_t^{\text{resid}}(L) u_t^{\text{resid}}, \end{aligned}$$

where info , u^{info} , resid and u^{resid} are square-summable lag polynomials in non-negative powers of L . The innovations, \hat{a}_t^{info} , \hat{a}_t^{resid} , u_t^{info} and u_t^{resid} , are mutually orthogonal white noise. In particular, \hat{a}_t^{info} and \hat{a}_t^{resid} are innovations to aggregate productivity, satisfying

$$\hat{a}_t = \hat{a}_t^{\text{info}} + \hat{a}_t^{\text{resid}}, \quad (43)$$

with standard deviations σ^{info} and σ^{resid} . The corresponding lag-polynomial info captures how incomplete information regarding a_t influences the propagation of productivity shocks.²³ The innovations u_t^{info} and u_t^{resid} , each two-dimensional, are intrinsic shocks to \hat{a}_t^{info} and \hat{a}_t^{resid} . Accordingly, the lag-polynomial u^{info} defines intrinsic fluctuations in \hat{a}_t^{info} , driven by expectation errors, whereas u^{resid} defines intrinsic fluctuations in the residual wedges \hat{a}_t^{resid} .

The defining difference between \hat{a}_t^{info} and \hat{a}_t^{resid} is that we impose the conditions of Theorem 1 on \hat{a}_t^{info} , whereas \hat{a}_t^{resid} remains unrestricted. We gauge the potential role of incomplete information for explaining the U.S. business cycle by maximizing the contribution of expectation errors u_t^{info} to the filtered variance of \hat{y}_t . Let $\hat{y}_t^{\text{tfp}} = E[\hat{y}_t | (\hat{a}_t^{\text{info}}, \hat{a}_t^{\text{resid}})_s = 0]$, $\hat{y}_t^{\text{info}} = E[\hat{y}_t | (u_t^{\text{info}})_s = 0]$, and $\hat{y}_t^{\text{resid}} = E[\hat{y}_t | (u_t^{\text{resid}})_s = 0]$ denote the projection of the output gap on aggregate productivity, expectation errors, and residual shocks, respectively. Independence of the innovations implies $\text{Var}[\hat{y}_t] = \text{Var}[\hat{y}_t^{\text{tfp}}] + \text{Var}[\hat{y}_t^{\text{info}}] + \text{Var}[\hat{y}_t^{\text{resid}}]$. Then the maximal contribution of u_t^{info} is given by:

$$\max_{\text{info}, \text{resid}, \text{info}, \text{resid}} \text{Var}[\hat{y}_t^{\text{info}}] = \text{Var}[\hat{y}_t] \quad (44)$$

²³Conversely, 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 \hat{a}_t^{resid} . If we instead let \hat{a}_t^{info} and \hat{a}_t^{resid} load jointly on the combined productivity shock \hat{a}_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 \hat{a}_t^{resid} . Below we also consider the case where agents perfectly observe aggregate productivity, in which case both settings give identical results.

standard deviations, $(\sigma_x; \sigma_l; \sigma_z)$, by up to 1 order of magnitude relative to the baseline calibration.²⁴ With the exception of the symmetric information benchmark, all specifications allow households and firms to have access to potentially heterogeneous information.

5.2.1 Benchmarks

As benchmark, we first consider the symmetric information case where $\sigma_{i,t}^{\text{sym}}$ is set as in (22) and the heterogeneous information case where $h_{i,t}$ and $f_{i,t}$ 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 U.S. business cycle fluctuations unexplained by the productivity shock ($\text{Var}[y_t^{\text{info}}] = \text{Var}[y_t | f_{a_t}, g_{s_0}] = 1$), provided that $(\sigma_x; \sigma_l; \sigma_z)$ are at least as volatile as in our baseline calibration (scale = 1).²⁵ 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).

5.2.2 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 U.S. 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 $f_{i,t}$ and $h_{i,t}$ as in (23) and (24), augmented by f_{a_t}, g_{s_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.

²⁴The scaling is applied to all three micro-shocks proportionately to their respective baseline values; i.e., the scaled standard deviations $\text{tw5sr8(e)-36n3(b)27(y)8334(a)]T$

For the baseline calibration (scale = 1), sentiment shocks can account for 57 percent of non-productivity fluctuations in U.S. output.

Next, consider the case without sentiment shocks. To eliminate them, we set $\epsilon_{i,t}^f$ and $\epsilon_{i,t}^h$ as in (23) and (24), augmented by $\epsilon_{i,t}^f = \sigma_f z_{i,t}$ and $\epsilon_{i,t}^h = \sigma_h z_{i,t}$. Here we do not include the iid-productivities, $\epsilon_{i,t}$, in $\epsilon_{i,t}^f$ or $\epsilon_{i,t}^h$ as this would allow firms to fully back out a_t from observing $a_{i,t}$. However, because $\epsilon_{i,t}$ is serially uncorrelated and firms know $a_{i,t}$, expectation errors about $\epsilon_{i,t}$ 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)²⁶, 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 U.S. 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.²⁷

Interestingly, however, the two cases without sentiment- and productivity-driven con -

Table 2: Implied variance contribution to U.S. output

	Contribution to		
	$\text{Var}[y_t f a_t, s g_s, o]$	$\text{Var}[y_t]$	$\text{Var}[\hat{y}_t]$
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	0.04	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 $\text{Var}[y_t]$ and $\text{Var}[\hat{y}_t]$ is 11 and 36 percent, respectively. All variance contributions are computed at business cycle frequencies for the baseline calibration of $f a_{i,t} g$ and $f z_{i,t} g$ (i.e., scale = 1 in Figure 3).

5.2.4 Implied variance contribution to U.S. output

The results in Figure 3 show the business-cycle contributions to output volatility that is unexplained by productivity, $\text{Var}[y_t | f a_t, s g_s, o]$ (equivalently $\text{Var}[y_t | f a_t, s g_s, o]$). 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 \hat{y}_t . Looking at the contribution to y_t , sentiment-driven 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 .

6 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 quantitatively 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.

References

Benhabib, Jess, Pengfei Wang, and Yi Wen. 2015. "Sentiments and Aggregate Demand Fluctuations." *Econometrica*, 83(2): 549{585.

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

Foster, Lucia, John Haltiwanger, and Chad Syverson. 2016. "The Slow Growth of New Plants: Learning about Demand?" *Economica*, 83(329): 91-129.

A Proof of Main Theorem

Consider any expectation wedge $J_{i;t} \geq T_t$ from the primal economy and the corresponding lower bound $J_{i;t}$ on $I_{i;t}^J$ in the incomplete information economy. Define the expectation "targets"

$$a_{i;t}^j = A_1^j g_{i;t+1} + A_2^j f_{i;t+1} + B_1^j g_{i;t} + B_2^j f_{i;t}$$

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

We want to show that conditions (i) and (ii) are jointly necessary and sufficient for the construction of some $J_{i;t}^j = S_{i;t}^j f_{i;t}^j + g_{i;t}^j$ such that

$$E[a_{i;t}^j | I_{i;t}^J] = E[a_{i;t}^j | I_t] + J_{i;t}^j \quad (45)$$

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 (45) to get

$$J_{i;t}^j = E[a_{i;t}^j | I_t] - E[a_{i;t}^j | I_{i;t}^J] \quad (46)$$

Computing the unconditional expectation over (46) yields $E[J_{i;t}^j] = 0$. Similarly, postmultiplying (46) by $J_{i;t}^j$ and $J_{i;t}^j \geq J_{i;t}^j$ gives

$$\begin{aligned} E[J_{i;t}^j] &= E[a_{i;t}^j | I_t] - E[a_{i;t}^j | I_{i;t}^J] \\ E[J_{i;t}^j] &= E[a_{i;t}^j | I_t] - E[a_{i;t}^j | I_{i;t}^J] \end{aligned}$$

as $J_{i;t}^j = I_{i;t}^j - I_t$. Again taking the unconditional expectation over the right-hand sides, we have $E[J_{i;t}^j] = E[J_{i;t}^j] = 0$ for all $J_{i;t}^j \geq J_{i;t}^j$.

Sufficiency We demonstrate sufficiency by construction. Let $\hat{a}_t = E[a_t | I_t]$ and consider the information set $I_t = S_t [f_t, g_t]$, where $S_t = \hat{a}_t + J_t = J_t$ is a signal that replicates the correlation structure of the expectation we wish to implement. Notice that I_t inherits recursiveness from S_t , ensuring consistency with Assumption 2.

From the law of iterated expectations, we have $E[a_t | S_t] = E[\hat{a}_t | S_t]$ as $S_t = I_t$. Projecting

\hat{a}_t onto s_t we obtain

$$\begin{aligned}
 E[a_t/s_t] &= \text{Cov}[\hat{a}_t; s_t] \text{Var}[s_t]^{-1} s_t \\
 &= \text{Cov}[s_t - \hat{a}_t; s_t] \text{Var}[s_t]^{-1} s_t \\
 &= \text{Var}[s_t] \text{Var}[s_t]^{-1} s_t \\
 &= s_t;
 \end{aligned} \tag{47}$$

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 $s_t = \sum_{t \in S_t} s_t$. Noting that by construction no other $t \in S_t$ can improve the forecast about a_t ,²⁹ we obtain

$$E[a_t/s_t] = E[a_tj/t] = E[a_tj/t] + \hat{a}_t$$

As the argument above applies to any $j_{i,t} \in T$, we have constructed exactly the information sets needed to satisfy (45) for all $(i; j; t)$:

²⁹To see this, note that the forecast error conditional on s_t is necessarily uncorrelated with any other $t \in S_t$: $\text{Cov}[a_t - E[a_tj/s_tg; t], s_t - \hat{a}_t] = \text{Cov}[a_t - s_t; t] = \text{Cov}[a_t - \hat{a}_t - \hat{a}_t; t] = \text{Cov}[\hat{a}_t; t] = 0$. Here the first equality follows from (47); the second one follows per the definition of \hat{a}_t ; the third one follows, because $a_t - \hat{a}_t$ defines the forecast error under full information I_t , so that any $t \in S_t - I_t$ must be orthogonal to it; and the last equality follows from the conditions of the theorem.

B Online Appendix: Additional Proofs and Results

B.1 Proof of Lemma 1

The characterization for \hat{y}_t is immediate. To solve for \hat{t}_t , let \hat{t}_t

Equipped with Lemma 2, our proof proceeds in two steps. First, we derive the mappings $(\cdot; \cdot)_{\mathcal{F}_s}$ and $(\cdot; \cdot)_{\mathcal{F}}_{\mathcal{F}_s}$ in closed form. Second, with this explicit characterization at hand, we complete the proof by constructing processes for \cdot and \cdot that for any given $(\cdot; \cdot)$ satisfy the conditions of Lemma 2.

Characterization of \cdot_s The mapping \cdot_s is immediate from (37),

$$\cdot_s(\cdot; \cdot) = \text{Cov}[\cdot; \cdot]_{\mathcal{F}_s}$$

Using (54) to eliminate $dy_{i;t+1}$ in (51), we have

$$(1 - \alpha) db_{i;t} + (L^{-1} - 1) (L^{-1})_+^{-1} i;t = \alpha^{-1} (L^{-1} - 1) B(L)_+^{-1} i;t \quad (56)$$

where $[\cdot]_+$ sends the negative powers of L to zero. Further using (56) to eliminate $db_{i;t}$ in (55) and applying the z -transform, we obtain the following functional equation

$$(1 - \alpha z)^{-1} y(z) = \alpha^{-1} [(1 - z)B(z) - B_0] + \alpha^{-1} (1 - \alpha)^{-1} (1 - \alpha z)^{-1} B(z)z \quad (57)$$

Evaluating (57) at $z = \alpha^{-1} (1 - \alpha)$, pins down y_0 and $y(z)$, from which we obtain the following equilibrium process for $dy_{i;t} = dy(L^{-1})_{i;t}$ and $dc_{i;t} = dc(L^{-1})_{i;t}$:

$$dy(z) = \alpha^{-1} (1 - z)B(z) + \alpha^{-1} (1 - \alpha)^{-1} B(\alpha^{-1} z) \quad (58)$$

and

$$dc_{i;t} = \alpha^{-1} (1 - z)B(z) + \alpha^{-1} (1 - \alpha)^{-1} B(\alpha^{-1} z) \quad (59)$$

Collecting equations, we obtain

$$s(\cdot; \mathcal{F}) = \text{Cov} \begin{pmatrix} 2 & 2 & 3 \\ 1 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 6 \\ 4 \\ 4 \end{pmatrix} \begin{pmatrix} 6 \\ 4 \\ 4 \end{pmatrix} \begin{pmatrix} 7 \\ 5 \\ 5 \end{pmatrix} (1 - L)B(L)_{i;t} \begin{pmatrix} 7 \\ 5 \\ 5 \end{pmatrix}$$

To begin, substitute (60) to (49), post-multiply both sides by

$$M = \begin{pmatrix} 1 & 1 & 0 \\ 4 & 1 & 5 \\ 0 & 1 & 1 \end{pmatrix};$$

and apply the z-transform, to obtain the equivalent functional equation

$$\tilde{\gamma}(z) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} B(z)(1 - z^{-1})B(z^{-1})^0 + \begin{pmatrix} 1 & 0 & 0 \\ 4 & 1 & 5 \\ 0 & 0 & 1 \end{pmatrix} + B(z)(1 - z^{-1})B(z^{-1})^0 \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 \\ 4 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \quad (61)$$

where $\tilde{\gamma}(z) = Zf_{sMg_s}$ is the (one-sided) z-transform of f_{sMg} , and where B parametrizes the joint process $(f_{i:t}; f_{i:t})$ as in the characterization of above. In particular, let

$$B(L) = \begin{pmatrix} B(L) \\ B_a(L) \\ B_z(L) \end{pmatrix}$$

where $B(z)$ is a lag-polynomial of size $2 - n$, $B_a(z)$ and $B_z(z)$ are each lag-polynomials of size $1 - n$, and n is an arbitrary number of innovations. Then (61) can be further rewritten as

$$\tilde{\gamma}_1(z) + \tilde{\gamma}(z) = (1 - z^{-1})B(z)B(z^{-1})^0 + \tilde{\gamma}(z) + B(z)B(z^{-1})^0 \quad (62)$$

and

$$\tilde{\gamma}_2(z) = (1 - z^{-1})B(z)B_a(z^{-1})^0; \quad (63)$$

where $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ correspond to the first two and third column of $\tilde{\gamma}$, respectively, and where

$$\tilde{\gamma}(z) = \begin{pmatrix} B(z) \\ B_a(z) \\ B_z(z) \end{pmatrix} (1 - z^{-1})^0$$

and

$$\begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 0 \\ 0 \end{pmatrix} ;$$

Fix $N = h$ as the largest non-zero power of z in \sim . Consider the following parametric structure for B , B_a , and B_z :

$$\begin{aligned} B(z) &= \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \\ B_a(z) &= \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \\ B_z(z) &= \begin{pmatrix} 0 & 0 \\ z_0 & z_1 z \end{pmatrix} \end{aligned}$$

with

$$(z) = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix} \begin{pmatrix} 0 \\ 0 \end{pmatrix} ;$$

and

$$z_0(z) = (1 - \beta) z_0 + \beta z_1 + \beta^2 z_2 + \dots$$

Notice that (i) the left-hand side, $z_0(z)$, is exogenously determined by the aggregate economy that we are trying to implement, and (ii) we have z_0 as a degree of freedom to induce an arbitrary unconditional covariance on the right-hand side. Writing out the right-hand side in the time-domain, we have

$$z_0 = z_0 + \beta z_1 + \frac{\beta^2}{1 - \beta} + \sum_{j=1}^{\infty} \beta^j z_j (1 + \beta) + \sum_{j=1}^{\infty} \beta^j z_j \quad (64)$$

jj in the orthogonality

where the first equality exploits that by Theorem 1 $\frac{p_i^h}{i;t} \geq \frac{h}{i;t} s$ and thus $\frac{p_i^h}{i;t} \geq (n_{i;t} s; C_{i;t} s)$ for all $s \geq 0$.

Similarly, substituting for $w_{i;t}$ using the firm's labor demand and taking first differences, orthogonality of the firm wedge with respect to $dw_{i;t}$ requires

$$\text{Cov}[\frac{p_i^f}{i;t}; da_{i;t} + dp_{i;t} + d\frac{p_i^f}{i;t}] = \text{Cov}[\frac{p_i^h}{i;t}; d\frac{p_i^f}{i;t}] = 0; \tag{67}$$

Here the first equality follows as $\frac{p_i^f}{i;t} \geq \frac{f}{i;t} 0$ implies $\frac{p_i^f}{i;t} \geq n_{i;t} 0$ for all $s \geq 0$ and, hence, $\frac{p_i^f}{i;t} \geq (dy_{i;t} - dn_{i;t} + dp_{i;t})$ under the conditions of the proposition.

Subtracting (66) from (67), we have

$$\text{Cov}[\frac{p}{i;t}; d\frac{p}{i;t} - \frac{p}{i;t}] = 0$$

or

$$(1 - \text{Corr}[\hat{y}_t; \hat{y}_{t-1}])^{-1} \text{Var}[\hat{y}_t] - \text{Cov}[\hat{y}_t; \hat{y}_{t-1}] = (1 - \text{Corr}[\frac{p}{i;t}; \frac{p}{i;t-1}]) \text{Var}[\frac{p}{i;t}] \geq 0;$$

which implies the bound given in the statement of the proposition. □

C Online Appendix: Estimation of Unrestricted Wedge Process

Here we describe the methodology for estimating the unrestricted wedges $\hat{\omega}_t$ used in Section 5.

C.1 Description of Methodology

We model the unrestricted wedges as a MA(14) process, which loads on two intrinsic innovations, represented by the 2×1 vector u_t , in addition to the productivity shock ϵ_t ,

$$\hat{\omega}_t = \phi(L)\epsilon_t + \psi(L)u_t;$$

where $\phi(L)$ and $\psi(L)$ are square-summable lag polynomials in non-negative powers of L , and ϵ_t and u_t are orthogonal white noise. W.l.o.g., we normalize $\text{Var}[u_t] = I_2$, leaving us to estimate $\psi_{ma}(\phi; \psi)$. For this purpose, we use the generalized method of moments (GMM) to minimize the distance between the model's covariance structure and U.S. data on

real per-capita output, inflation, nominal interest rates, and per-capita hours.³¹ Let

$$\tilde{\gamma}_T = \text{vech} \sqrt{\text{Var}[(q_t^{\text{data}}; \dots; q_t^{\text{data}})]} g;$$

denote the empirical auto-covariance matrix of frequency-filtered quarterly US data for q ($y_t; i_t; i_t; n_t$). We target auto-covariances between zero and $k = 8$ quarters. For the filtering, we use the [Baxter and King \(1999\)](#) approximate high-pass filter with a truncation horizon of 32 quarters, which we denote by $q_t = BK_{32}(q_t)$.³²

To conserve on the 91 parameters that characterize γ_{ma} , we make two observations, documented in [Figure 4](#) below. First, $\tilde{\gamma}_T$ is well-described by a VAR(1) process for y_t . Second, a MA(14) truncation of the VAR(1) process that best replicates $\tilde{\gamma}_T$ is almost indistinguishable (in terms of second moments) from the VAR(1) process itself. Accordingly, we construct γ_{ma} by first estimating y_t as a VAR(1) that is driven by u_t and i_t , and then constructing $\hat{\gamma}_{ma}$ as the MA(14) truncation of the estimated process.³³

Let γ_{ar} denote the 10 parameters characterizing the VAR(1) and γ . Then the estimator is given by

$$\hat{\gamma}_{ar} = \underset{ar}{\text{argmin}} (\tilde{\gamma}_T - \tilde{\gamma}(ar))^{\theta} W^{-1} (\tilde{\gamma}_T - \tilde{\gamma}(ar)); \quad (68)$$

where $\tilde{\gamma}(ar)$ is the model analogue to $\tilde{\gamma}_T$ and W is a diagonal matrix with the bootstrapped variances of $\tilde{\gamma}_T$ along the main diagonal. To avoid the issues detailed in [Gorodnichenko and Ng \(2010\)](#), our model analogue $\tilde{\gamma}(ar)$ is computed after applying the same filtering procedure to the model that we have applied to the data.

A final challenge for estimating the model is that filtering the model can be computationally expensive. We address this issue by proving the following equivalence results (see [Appendix C.3](#) for proof).

Lemma 3. *Estimator (68) is equivalent to*

$$\hat{\gamma}_{ar} = \underset{ar}{\text{argmin}} (\gamma_T - \gamma(ar))^{\theta} W^{-1} (\gamma_T - \gamma(ar)); \quad (69)$$

where $\gamma_T = \text{vech} \sqrt{\text{Var}[(ds_t; \dots; ds_t)]} g$ and $W = (\theta W^{-1})^{-1}$ for $K = k + 2$. The trans-

³¹Data range from 1960Q1 to 2012Q4. Real output is given by nominal output divided by the GDP deflator. Inflation is defined as the log-difference in the GDP deflator. Interest rates are given by the Federal Funds Effective rate. Hours are given by hours worked in the non-farm sector. Variables are put in per-capita terms using the non-institutional population over age 16.

³²The [Baxter and King \(1999\)](#) filter requires specification of a lag-length for the approximation. We set to their recommended value of 12.

³³Our estimator penalizes excessively persistent dynamics beyond the usual business cycle horizon by imposing a numerical penalty on impulse responses beyond 32 quarters.

dence intervals (depicted by the shaded areas). The solid blue and red lines show the corresponding moments for the estimated model for the VAR(1) and MA(14) truncation of the wedges, respectively. Each row i and column j in the table of plots shows the covariances between φ_t^i and φ_{t-k}^j with lags $k \in \{0, 1, \dots, 8\}$ depicted on the horizontal axis. Despite the parametric restriction on φ_t and a_t and the fact that we have less shocks than data series, the unrestricted-wedge model does a very good job at capturing the auto-covariance structure of the four time series. In addition, there is no notable difference between the VAR(1) and MA(14) truncation of φ_t .

C.3 Proof of Lemma 3

where

$$= P_0 \begin{matrix} 2 \\ 6 \\ 4 \end{matrix} \begin{matrix} BL_0 \\ \\ \\ \end{matrix} : \begin{matrix} 3 \\ 7 \\ 5 \end{matrix} \begin{matrix} BL_0 \\ \\ \\ \end{matrix} P_1 \quad (74)$$

with B and L_j as in (72) and (73). Substitution in (70) yields (71).

D Online Appendix: Comparative Statics With Countercyclical Inflation

In analogue to Figure 2, we explore comparative statics with respect to the parametrization of the micro-shocks, but for the case where inflation is countercyclical with $\gamma_y = .3$. The results, shown in Figure 5, display the same qualitative pattern as for the procyclical case explored in the main text. While the maximal volatility is higher, we again see a clear positive relationship between σ_y^{\max} and the volatilities of the micro shocks. As before, the impact of idiosyncratic demands shocks is most relevant, paralleling their key role in the procyclical case.

Here we do not include the cases without demand uncertainty ($p_{i,t} \geq \bar{p}_{i,t}$), because in line with our discussion in the main text, in these cases inflation is necessarily procyclical (see Appendix B.3 for a formal proof). Intuitively, this reflects again the discrepancy in propagation underlying the pro- and countercyclical inflation cases: While procyclical inflation is tied to nominal misperception and expectation errors about aggregate prices, countercyclical inflation is tied to expectation errors regarding local demand, and thus is impossible to implement when $p_{i,t}$ is observed by firms. (See also the explanations given in the context of Figure 1.)

