Information-driven Business Cycles:
A Primal Approach[*]

Ryan Chahrour

Boston College

Robert Ulbricht

Toulouse School of Economics

March 8, 2017

Abstract

We develop a methodology to estimate DSGE models with incomplete information, free of parametric restrictions on information structures. First, we define a "primal" economy in which deviations from full information are captured by wedges in agents' equilibrium expectations. Second, we provide implementability conditions, which ensure the existence of an information structure that implements these wedges. We apply the approach to estimate a New Keynesian model in which firms, households and the monetary authority have dispersed information about business conditions and productivity is the only aggregate fundamental. The estimated model fits the data remarkably well, with informational shocks able to account for the majority of U.S. business cycles. Output is driven mainly by household sentiments, whereas firm errors largely determine inflation. Our estimation indicates that firms and the central bank learn the aggregate state of the economy quickly, while household confusion about aggregate conditions is sizable and persistent.

Keywords: Business cycles, dispersed information, DSGE models, primal approach, sentiments.

JEL Classification: E32, D84.

1   Introduction

Many prominent theories in macroeconomics are based on incomplete information. Among their applications, such theories offer a structural interpretation of cyclical fluctuations, formalizing the widespread idea that business cycles are driven by waves of optimism and pessimism among consumers and firms (Lorenzoni, 2009; Angeletos and La'O, 2013; Benhabib, Wang and Wen, 2015). However, very few of these models have been investigated quantitatively, mainly because of technical difficulties arising from the introduction of dispersed information in general equilibrium frameworks and the challenge of specifying ex ante plausible information structures to explore. This paper develops a new approach that avoids these difficulties, and uses it to explore the quantitative potential of dynamic stochastic general equilibrium (DSGE) models with incomplete information.

Our approach defines a "primal" economy in which deviations from full information are summarized by wedges in agents' equilibrium expectations. We then provide necessary and sufficient conditions that ensure the existence of an information structure that is consistent with the expectation errors implicit in these wedges. Subject to these implementability conditions, the set of dynamics spanned by expectational wedges in the primal economy is equivalent to the set of dynamics that is feasible in the incomplete-information economy. Exploiting this equivalence, we show how to estimate DSGE models with incomplete information using standard tools developed for full-information economies and without imposing parametric restrictions on information structures.

We apply our approach to estimate a dispersed-information version of an otherwise standard New-Keynesian model, in which shocks to productivity are the only fundamental source of aggregate volatility. The model allows households, firms, and the monetary authority to be imperfectly informed about both local and aggregate economic conditions. While the incomplete-information version of our economy is generally hard to solve, the corresponding primal economy permits a simple aggregate representation. The representation resembles a standard New Keynesian model with shocks to demand, markups and interest rates. Specifically, expectational errors by households map into a "demand wedge" in the New Keynesian IS curve, expectational errors by firms' map into a "supply wedge" in the Phillips curve, and errors made by the monetary authority replace the usual exogenously-specified shock to the interest target in the Taylor rule. The behavior of these wedges is constrained by the implementability conditions characterized by our approach.

We estimate the model using the generalized method of moments, minimizing the distance between the auto-covariance structure generated by the model and U.S. data on output, employment, inflation, and interest rates. We find that the estimated model does a remarkably

good job at matching the business cycle comovements found in the data, essentially replicating the vast majority of (auto) covariances within the confidence region of the data. This is in stark contrast to the full-information benchmark economy, driven only by technology shocks, which cannot replicate the joint comovement of output, employment and inflation seen in the data.

Essential for the empirical performance of the model is its ability to generate wedges that are correlated across equations. In particular, there is a strong negative correlation between the wedge entering the aggregate Philips curve and both the wedges in the aggregate Euler equation and the Taylor rule. While it is typically difficult for full-information models with structurally uncorrelated shocks to generate perturbations that are correlated across equations, expectational wedges are naturally correlated for two reasons. First, information can be correlated across different types of agents. E.g., a joint optimism regarding productivity could generate correlation across firms and household wedges. Second, because the objectives of agents are strategically interdependent, errors of one type of agent naturally translate also into errors of other types as long as agents are not fully aware of one another's errors. We show that the latter can account for virtually all of the estimated correlation patterns.

In our empirical specification, expectational wedges are driven by three independent innovations. First, they correlate with productivity shocks, reflecting imperfect information regarding the aggregate productivity process. Specifically, we find that households learn about productivity innovations only gradually, implying a slow adjustment of output in response to TFP shocks. By contrast, we cannot reject that firms and the monetary authority are fully aware of aggregate productivity, resulting in countercyclical movements in inflation and the Feds funds rate. The dynamics in response to productivity shocks account for roughly 10 percent of the business-cycle fluctuations in output and employment, and roughly 20 percent of inflation.[1]

The two remaining innovations reflect that agents may also anticipate changes in economic conditions that are never realized, leading to business cycle fluctuations that are driven purely by expectations. We find that the first of these two shocks is mainly driven by waves of optimism and pessimism on the part of households, inducing procyclical fluctuations in output, employment and inflation akin to a "demand shock". Again, we cannot reject that the monetary authority is fully aware of this shock, which explains the procyclical movement in interest rates. Firms, on the other hand, learn only gradually about the household error,

---

[1]The small contribution of productivity is consistent with both recent DSGE estimations and the structural VAR literature, which rarely finds that productivity shocks explain more than one quarter of output cyclicality (Shapiro and Watson, 1988; King et al., 1991; Cochrane, 1994; Gali, 1999; Christiano, Eichenbaum and Vigfusson, 2003; Smets and Wouters, 2007).

dampening the initial inflation response.

Finally, we find that the economy's response to the third innovation is fully driven by firms' beliefs about current and future business conditions, leading them to cut prices. The inflation-response to these expectational shocks is extremely short-lived, however, reflecting a quick error correction on the side of firms with an half-life of less than one quarter. In line with the short-lived nature of firm errors, the errors made by households and the monetary authority are such that we cannot reject that they are unaware of the shock to firms expectations. In particular the lack of adjustment in interest rates, in turn, dampens the output effect and reinforces the impact on prices, leading the shock be a main driver of inflation but not output.

Our baseline estimation imposes few restrictions regarding what information is potentially available to the agents in our economy. Instead we use our estimates to infer a plausible information structure that implements the estimated expectation processes. Our results indicate that expectational errors by firms and the central bank are largely confined to within a one-year window, consistent with a lack of reliable real time statistical information. Outside this window, we cannot reject that firms and the central bank have perfect information. By contrast, implementing household expectations requires that they be persistently confused about several aggregate variables.

The methodology developed in this paper is related to the literature on information-robust predictions by 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 primal approach developed in this paper is similar in that it also demonstrates the equivalence between a class of incomplete-information economies with another class of full-information models. It is more general, however, as it is not limited to static game environments, but equally applies to dynamic market economies. Moreover, the primal approach developed in this paper gives straightforward implementability conditions, and it extends to arbitrary "minimal information requirements" that can be imposed by the researcher from the outset.

On the applied side, our analysis relates to a recent literature exploring business cycle models with incomplete information. While the literature is mostly theoretical, there are now a few studies with a quantitative focus. In particular, Angeletos, Collard and Dellas (2015) explore a version of Angeletos and La'O (2013); Blanchard, L'Huillier and Lorenzoni (2013) estimate a simplified version of Lorenzoni (2009); Melosi (2014) estimates a variant of Woodford (2003); and Maćkowiak and Wiederholt (2015) calibrate a particular DSGE model with rational inattention. A notable difference with respect to these papers is the

flexibility of expectation dynamics considered in this paper. In particular, our approach does not require us to take 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. Instead it allows us to evaluate the empirical performance across *all* information structures and let the data decide on which provides the best fit.

At a methodological level, the closest to our approach are Jurado (2016)—who estimates a model with a version near-rational belief distortions—and Angeletos, Collard and Dellas (2015)—who bypass the computational difficulties of incomplete information by relaxing the common prior assumption. Our approach allows for a similar generality of application, while ensuring that expectation errors are consistent with rational expectations of all agents.

In its ability to reduce the computational burden of solving (and estimating) incomplete information models, our approach paper also relates to Rondina and Walker (2014), Acharya (2013) and Huo and Takayama (2015), who use frequency-domain techniques to obtain analytical solutions in certain models, and Nimark (2009) who explores the asymptotic accuracy of a finite-state approximation approach to a certain class of dispersed information models.

This paper is also related to the business cycle accounting literature in the tradition of Chari, Kehoe and McGrattan (2007). These papers consider simple economies augmented by a number of reduced-form wedges to equilibrium conditions. There are two important differences between papers in that tradition and the approach developed in this paper. First, in contrast with the business cycle accounting literature, we approach the wedges in our economy with a single structural interpretation in mind, supported by our equivalence result. Second, while the wedges in the business cycle accounting literature are exactly identified by the data, the information wedges in our economy are generally over-identified, both because we use more data series than shocks and because of the restrictions imposed by our structural interpretation of wedges as shocks to information. Conceptually the second source of over-identification is crucial, because the orthogonality conditions on our wedges precisely ensure that they can be implemented by a valid information structure.

The paper is structured as follows. Section 2 sets up the model economy. Section 3 describes the primal approach. Section 4 details our empirical strategy. Section 5 presents the baseline empirical results. Section 6 explores what the results imply about the informational primitives of our economy. Section 7 concludes.

2   The Model Economy

2.1   Setup

The model is a standard New Keynesian model where households, firms and the monetary authority have a generic set of (possibly incomplete) information regarding both local and aggregate economic conditions. 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 continuum of price-setting firms in a local labor market. Firms use the labor provided by the household to produce differentiated intermediate goods indexed by $j \in [0, 1]$. A competitive final goods and distribution sector, operating in the mainland, uses these goods as inputs to produce and distribute consumption to households.

Households   The preferences of the household on island i are given by

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

with

$$U(C_{i,t}, N_{i,t}) = \log C_{i,t} - \frac{1}{1+\zeta} N_{i,t}^{1+\zeta},$$

where $N_{i,t}$ is hours worked, $C_{i,t}$ is final good consumption, $\beta \in (0, 1)$ is the discount factor, $\zeta \geq 0$ is the inverse of the Frisch elasticity of labor supply, and $\mathcal{I}_{i,t}^{h}$ is the set of information available to household i at time t. The household's budget constraint is

$$P_{i,t}C_{i,t} + Q_{i,t}B_{i,t} \leq W_{i,t}N_{i,t} + B_{i,t-1} + \int_{0}^{1} D_{ij,t}\,\mathrm{d}j,$$

where $P_{i,t}$ is the price of the final good, $W_{i,t}$ is the nominal wage rate, $Q_{i,t}$ is the nominal price of a riskless one-period bond, $B_{i,t}$ are bond holdings, and $D_{ij,t}$ are the profits of firm j on island i. Prices for bonds and final consumption are island-specific due to the presence of idiosyncratic distribution costs (further detailed below). 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, leaving households exposed to idiosyncratic income risks.

Intermediate-goods producers Each good $(i, j) \in [0, 1]^2$ is produced by a monopolistically competitive firm which has access to the production technology

$$Y_{ij,t} = A_{i,t} N_{ij,t}^\alpha, \tag{1}$$

with $0 < \alpha \leq 1$. Firms compete in prices à la Calvo (1983). Each period, a firm resets its price with probability $1 - \lambda$ and keeps it unchanged with probability λ. The firms' objective is to maximize the expected market value of their profits, discounted at the market rate $Q_{i,t}$,

$$\mathbb{E} \left\{ \sum_{\tau=0}^\infty \lambda^\tau Q_{i,t+\tau} \left(P_{ij,t} Y_{ij,t+\tau|t} - W_{i,t+\tau} N_{ij,t+\tau|t} \right) | \mathcal{I}_{i,t}^f \right\},$$

where $Y_{ij,t+\tau|t}$ denotes the demand for good (i, j) in period $t + \tau$ when the price was last reset in period t, $N_{ij,t+\tau|t}$ is the corresponding labor input, and $\mathcal{I}_{ij,t}^f$ is the information available to firms on island i at time t.

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 productivity follows a random walk process given by

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

where 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 and distribution sector There is a competitive final-goods sector which aggregates intermediate input goods $(i, j) \in [0, 1]^2$, using the technology

$$Y_t = \left(\int_0^1 \int_0^1 Z_{i,t} Y_{ij,t}^{\frac{\theta-1}{\theta}} \, \mathrm{d}i \, \mathrm{d}j \right)^{\frac{\theta}{\theta-1}},$$

where $\theta > 1$ is the elasticity of substitution among goods, $Y_{ij,t}$ denotes the input of intermediate good (i, j) at time t, and $Z_{i,t}$ is an idiosyncratic demand-shifter with a time-invariant, stationary process that is i.i.d. across islands and satisfies $\int_0^1 \log(Z_{i,t}) \, \mathrm{d}i = 0$. Profit maxi-

mization implies that input demands are equal to

$$Y_{ij,t} = \left(\frac{P_{ij,t}}{P_t}\right)^{-\theta} Z_{i,t}^{\theta} Y_t, \tag{2}$$

where P_t is the aggregate price index given by

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

To distribute the consumption good to island i, final good firms must incur a random "iceberg" trade cost, such that $\exp(\nu_{i,t}^1)$ units of the final good must be shipped for 1 unit to arrive in island i. Similarly, nominal bonds are intermediated by a competitive distribution sector with random distribution costs $\exp(-\nu_{i,t}^2)$. Standard no-arbitrage conditions imply that the price of consumption in island i is given by

$$\log(P_{i,t}) = \log(P_t) + \nu_{i,t}^1,$$

and the price of nominal bonds is given by

$$\log(Q_{i,t}) = \log(Q_t) - \nu_{i,t}^2.$$

While the realizations of $\nu_{i,t}^1$ and $\nu_{i,t}^2$ are known by the distributors, they are not necessarily contained in $\mathcal{I}_{i,t}^h$. The role of these assumptions is to limit the ability of households to infer the aggregate shocks from their observations of the final goods and bonds prices.[2] We assume that $\nu_{i,t}^1$ and $\nu_{i,t}^2$ follow a time-invariant, stationary process that is i.i.d. across islands, ruling out any direct impact at the aggregate. The process is normalized so that $\int_0^1 \log(\nu_{i,t}^m) \, \mathrm{d}i = 0$ for $m \in \{1, 2\}$.

Monetary policy To close the model, we use a simple monetary policy rule to pin down the nominal bond price. Letting $i_t = -\log(Q_t)$, the central bank sets nominal bond prices

---

[2]As argued by Lorenzoni (2009), the ability of agents to infer the economy's aggregate state is in practice likely to be impaired by a larger number of shocks, by model misspecification, and by the possible presence of structural breaks. Introducing stochastic noises to the information of agents is a tractable way to mimic these complications in the context of a relatively simple model. Strategies with a similar effect include introducing random shocks to consumption baskets (Lorenzoni, 2009), limiting trade to only a finite set of suppliers (Angeletos and La'O, 2013), and using noise traders to perturb prices (e.g., Hellwig, 1980).

7

such that[3]

$$i_t = \mathbb{E}[\phi_y \hat{y}_t + \phi_\pi \pi_t | \mathcal{I}_t^b], \tag{3}$$

where $\pi_t \equiv \log(P_t/P_{t-1})$ is the inflation rate, $\hat{y}_t \equiv \log(Y_t/A_t)$ is the gap (up to an omitted constant) between actual output and potential output under flexible prices and full information, $\phi_y > 0$ and $\phi_\pi > 1$ parametrize the central bank's desired response to the output gap and inflation, and $\mathcal{I}_t^b$ is the central bank's information in period t. Note that the rule does not feature a smoothing parameter or an exogenous shock since both features may arise endogenously when the monetary authority has incomplete information.

Information Until now the only assumption on information that we have made is that within islands all firms share the same information set. Henceforth, we also assume symmetry across islands and time in the sense that the distribution of signals is identical for all $(\mathcal{I}_{i,t}^f, \mathcal{I}_{i,t}^h)$ and is stationary for all $(\mathcal{I}_{i,t}^f, \mathcal{I}_{i,t}^h)$ and $\mathcal{I}_t^b$.

While for the most part we do not further restrict these information sets from the outset, we do impose some minimal structure to guide our analysis. First, we assume that firms and households observe the local productivity $A_{i,t}$, local consumption $C_{i,t}$, local consumer and producer price indexes $P_{i,t}$ and $\bar{P}_{i,t}$, the local bond price $Q_{i,t}$, and the local wage $W_{i,t}$. In equilibrium, these statistics span most local variables such as $N_{i,t}$, $B_{i,t}$ and $D_{i,t}$, and ensure that households are aware of their feasible consumption sets. Second, we impose a bound on the horizon at which agents in the economy can be confused about past aggregate variables. In particular, letting $X_t \equiv (Y_t, P_t, Q_t, N_t)$, we impose that $X_{t-\bar{h}}$ is known at date t for some $\bar{h} > 0$. Third, it will be convenient to explicitly impose some basic principles of rationality at this point, namely that firms, households and the monetary authority are aware of their expectations (or, equivalently, that they are aware of their own actions[4]), and that all agents perfectly recall all previously acquired information.

Let $\Theta_{i,t}$ and $\Theta_{b,t}$ denote the resulting lower bounds on the date-t information available to firms, households and the monetary authority so that

$$\Theta_{i,t} \subseteq \mathcal{I}_{i,t}^f, \quad \Theta_{i,t} \subseteq \Theta_{i,t}^h \quad \text{and} \quad \Theta_{b,t} \subseteq \mathcal{I}_t^b. \tag{4}$$

Then given the above considerations, we have that

$$\Theta_t^b = \{Q_t, X_{t-\bar{h}}\} \cup \Theta_{t-1}^b \tag{5}$$

---

[3]The rule also contains a constant intercept, $-\log \beta - \log \alpha^{\alpha/(1+\varsigma)}$, ensuring consistency with the natural rate at the zero-inflation steady state. The term is omitted as it drops out after log-linearizing the model.

[4]For consumers and firms, knowledge of $C_{i,t}$ and $\bar{P}_{i,t}$ is sufficient. For banks we impose awareness of Q_t.

and

$$\Theta_{i,t} = \{A_{i,t}, C_{i,t}, P_{i,t}, \bar{P}_{i,t}, Q_{i,t}, W_{i,t}, X_{t-\bar{h}}\} \cup \Theta_{i,t-1}. \tag{6}$$

2.2 Equilibrium conditions

Following standard practice, we focus on a log-linear approximation to the model about its non-stochastic, zero-inflation steady state. In what follows, lower-case letters denote log-deviations of a variable from its steady-state value.

From the firms' maximization problem, it follows that reset-prices on island i satisfy the recursive relation

$$p_{i,t}^* = (1 - \lambda\beta)\mathbb{E}[s_{i,t}|\mathcal{I}_{i,t}^f] + \lambda\beta\mathbb{E}[p_{i,t+1}^*|\mathcal{I}_{i,t}^f], \tag{7}$$

where $s_{i,t}$ is the profit-maximizing, flexible-price target under full information. As firms' demand has constant elasticity, desired markups are constant and $s_{i,t}$ equates to marginal costs, $w_{i,t} + n_{ij,t} - y_{ij,t}$. Substituting for the local labor supply relation $\zeta n_{i,t} = w_{i,t} - p_{i,t} - c_{i,t}$ and the demand curve (2), we get

$$s_{i,t} = \bar{p}_{i,t} + \xi(\hat{y}_{i,t} + \mu_{i,t}),$$

where

$$\bar{p}_{i,t} = \lambda\bar{p}_{i,t-1} + (1 - \lambda)p_{i,t}^* \tag{8}$$

is the producer price-index on island i, $\hat{y}_{i,t} = y_{i,t} - a_{i,t}$ is the island-specific output gap, $\mu_{i,t} = (\alpha/(\zeta+1)) \cdot (c_{i,t} - y_{i,t} + p_{i,t} - \bar{p}_{i,t})$ is island i's nominal trade-balance (times a constant), and $\xi \equiv (\zeta + 1)/(\alpha + \theta(1 - \alpha))$ is the output elasticity of the flexible-price target. Equations (7) and (8) can be manipulated to get

$$\pi_{i,t} = \mathbb{E}[\kappa(\hat{y}_{i,t} + \mu_{i,t}) + \beta\pi_{i,t+1}|\mathcal{I}_{i,t}^f] \tag{9}$$

with $\pi_{i,t} \equiv \bar{p}_{i,t} - \bar{p}_{i,t-1}$ and $\kappa \equiv (1 - \lambda\beta)(1 - \lambda)\xi/\lambda$.

Equation (9) together with the households' Euler equation

$$c_{i,t} = \mathbb{E}[c_{i,t+1} - i_{i,t} + p_{i,t+1} - p_{i,t} \,|\, \mathcal{I}_{i,t}^h], \tag{10}$$

the market clearing condition $y_t = \int_0^1 c_{i,t}\,\mathrm{d}i$, the resource constraint $y_t = a_t + \alpha n_t$, the requirement that the aggregate price index is consistent with both consumer and producer prices $p_t = \int_0^1 p_{i,t}\,\mathrm{d}i = \int_0^1 \bar{p}_{i,t}\,\mathrm{d}i$, and the monetary policy rule (3) define the equilibrium to this economy.

3   A Primal Approach to Solving Incomplete Information Models

We now describe the methodology used to characterize the set of incomplete-information equilibria in the model economy. Our approach is quite different from the usual approach that first fixes an information structure for all agents and then searches for a fixed point between beliefs and equilibrium dynamics. In the presence of incomplete information, finding this fixed point can be challenging, often involves high-dimensional state spaces, and results are known to be sensitive to the precise informational assumptions entertained. We side-step these issues by providing an explicit characterization of all feasible belief dynamics that are implementable with some information structure. With this characterization at hand, we then can fix any feasible process for equilibrium beliefs and treat its deviation from the corresponding full-information beliefs as a *primitive* of the model. This procedure essentially transforms the incomplete-information economy into a full-information wedge-economy, allowing us to use standard tools to complete the characterization of equilibrium. In this section, we explain the approach in detail.

3.1   The primal economy

Definition   We begin by defining a primal version of our economy. Let $\mathbb{E}_t[\cdot] \equiv \mathbb{E}_t[\cdot|\mathcal{I}_t^*]$ denote the full-information expectations operator, where $\mathcal{I}_t^*$ contains the history of all variables realized at date t.[5] The primal economy is constructed by replacing for $m \in \{b, f, h\}$ all expectations operator $\mathbb{E}[\cdot|\mathcal{I}_{i,t}^m]$ by $\mathbb{E}_t[\cdot] + \tau_{i,t}^m$, where

$$\tau_{i,t}^m \equiv \mathbb{E}[\cdot|\mathcal{I}_{i,t}^m] - \mathbb{E}_t[\cdot]$$

defines an information wedge that is treated as an exogenous stochastic process in the context of the primal economy. That is, the primal economy treats the gap between equilibrium expectations and full-information as a primitive of the model, transforming the model economy into a wedge-economy in which all expectations are taken with respect to full information. Once we have characterized feasible processes for these wedges, we can specify any such feasible process and solve the equilibrium using standard full-information techniques.

---

[5]Notice that which variables are realized at date t is to some extend definitional, requiring the researcher to take a stand on what is potentially knowable at date t. In particular, $\mathcal{I}_t^*$ may well contain fundamental innovations dated in the future if these innovations are assumed to be realized at date t as in the news literature. Henceforth, we assume that for all aggregate random processes, innovation are not knowable in advance, so that there is no role for news in the full-information economy.

In the model economy, we have three types of expectational equations: the island-specific evolution of producer prices (9), the island-specific Euler equations (10), and the monetary policy rule (3). The corresponding counterparts in the primal economy are given by:[6]

$$\pi_{i,t} = \mathbb{E}_t[\kappa(\hat{y}_{i,t} + \mu_{i,t}) + \beta(\pi_{i,t+1} - \tau^f_{i,t+1})] + \tau^f_{i,t} \tag{11}$$

$$c_{i,t} = \mathbb{E}_t[(c_{i,t+1} - \tau^h_{i,t+1}) - i_{i,t} + p_{i,t+1} - p_{i,t}] + \tau^h_{i,t} \tag{12}$$

$$i_t = \mathbb{E}_t[\phi_y \hat{y}_t + \phi_\pi \pi_t] + \tau^b_t. \tag{13}$$

Here $\tau^h_{i,t}$ is the prediction error, relative to full information, made by household i about its effective wealth and the returns on saving. A positive prediction error lets households increase their consumption relative to the optimal level under full information. On the firms' side, $\tau^f_{i,t}$ is the prediction error, relative to full information, regarding present or future marginal cost of production, reflecting a misjudgment about either productivity or consumers' demand. Finally τ^b_t is the prediction error, relative to full information, made by the monetary authority regarding its interest rate target.

Equilibrium in the primal economy   The equilibrium in the primal economy is defined by the equations stated in Section 2.2, where (3), (9) and (10) are replaced by (11), (12) and (13). Unlike the incomplete-information economy, in which aggregation involves average expectation operators and hence depends on the cross-sectional distribution of beliefs, the primal economy permits a simple aggregate representation. Letting $\tau^f_t = \int_0^1 \tau^f_{i,t}\,\mathrm{d}i$ and $\tau^h_t = \int_0^1 \tau^h_{i,t}\,\mathrm{d}i$ and integrating over (11) and (12) we get

$$\pi_t = \kappa \hat{y}_t + \beta \mathbb{E}_t[\pi_{t+1} - \tau^f_{t+1}] + \tau^f_t \tag{14}$$

and

$$\hat{y}_t = \mathbb{E}_t[\hat{y}_{t+1} - \tau^h_{t+1} - i_t + \pi_{t+1}] + \tau^h_t. \tag{15}$$

Equations (14) and (15) closely resemble the standard New Keynesian IS and Phillips curve, augmented by informational demand and supply wedges τ^h_t and τ^f_t. The system is completed by the Taylor rule (13), yielding a three-equation system that can be used to solve for the equilibrium dynamics of π_t, $\hat{y}_t$ and i_t.

---

[6]In practice there are multiple isomorphic ways to define the wedges. Here we define the firm and household wedge after rewriting (9) and (10) in their non-recursive forms. With this normalization all our wedges capture the gap relative to the choices that each agent would take if he or she had full information at t and all future dates, taking as given the equilibrium choices of all other agents.

3.2 An Equivalence Theorem

We now discuss implementability of the information wedges. Let $\mathcal{T}$ denote a stochastic process for $\mathcal{T}_t \equiv \{\tau_{i,t}^f, \tau_{i,t}^h\}_{i=0}^1 \cup \tau_t^b$ and let $\mathcal{E}(\mathcal{T})$ denote an equilibrium in the primal economy induced by $\mathcal{T}$.[7] We assume $\mathcal{E}(\mathcal{T})$ to have a stationary Gaussian distribution (see below for a discussion on how our results extend to non-stationary and non-Gaussian cases). The following theorem states our main result.

Theorem 1. *Fix a $\mathcal{E}(\mathcal{T})$. Then there exists an information structure consistent with* (4) *that implements $\mathcal{T}$, and hence $\mathcal{E}(\mathcal{T})$, in the incomplete-information economy if and only if for all i and t it holds that (i) $\mathbb{E}[\mathcal{T}_t] = 0$ and (ii)*

$$\mathbb{E}[\tau_t^b \theta] = 0 \text{ for all } \theta \in \Theta_{b,t}, \tag{16a}$$

$$\mathbb{E}[\tau_{i,t}^f \theta] = 0 \text{ for all } \theta \in \Theta_{i,t}, \tag{16b}$$

$$\mathbb{E}[\tau_{i,t}^h \theta] = 0 \text{ for all } \theta \in \Theta_{i,t}. \tag{16c}$$

The theorem gives two conditions that are jointly necessary and sufficient for $\mathcal{T}$ to be implemented by some information structure. Condition (i) is a simple rationality requirement that agents cannot be systematically wrong in the long-run. Condition (ii) is an orthogonality requirement between all the information wedges and the minimal information sets $\Theta_{b,t}$ and $\Theta_{i,t}$. The necessity of this restriction is the familiar principle that expectation errors must be orthogonal to all available information. The novel part is the sufficiency of condition (ii). For any $\mathcal{E}(\mathcal{T})$ with $\mathbb{E}[\mathcal{T}_t] = 0$, we can always construct an information structure that implements the joint process $\mathcal{E}(\mathcal{T})$ as long as it satisfies (16). The following example illustrates this in a simple case. The general proof is given in Appendix A.1.

Example Consider an economy defined by a single equilibrium condition, $y_t = \mathbb{E}[a_t | \mathcal{I}_t]$, where $\mathbb{E}[a_t] = 0$, and let $\Theta_t = \{y_{t-s}\}_{s \geq 0}$. The primal economy is given by

$$y_t = a_t + \tau_t. \tag{17}$$

Let $\mathcal{E}_t = (y_t, a_t, \tau_t)$ be a stationary Gaussian process satisfying (17). Theorem 1 states that $\mathcal{E}_t$ is implementable by some $\{\mathcal{I}_t\}$, satisfying $y_t \in \mathcal{I}_t$ for all t, if and only if (i) $\mathbb{E}[\tau_t] = 0$ and (ii) $\mathbb{E}[\tau_t y_{t-s}] = 0$ for all $s \geq 0$. The necessity of conditions (i) and (ii) is immediate, since optimal

---

[7]There is no need for the equilibrium in the primal economy to be unique. If there are multiple $\mathcal{E}(\mathcal{T})$ for a given $\mathcal{T}$, our results hold with respect to each of them. Even if the equilibrium $\mathcal{E}(\mathcal{T})$ in the primal economy is unique given $\mathcal{T}$, the incomplete-information economy may still feature multiplicity as $\mathcal{T}$ itself may be driven by sunspot-realizations as in, e.g., Benhabib, Wang and Wen (2015).

inference requires that expectation errors are orthogonal to variables in the information set and are unpredictable.

To see why the conditions are also sufficient, 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 $(\omega_t, \mathcal{E}_t)$ 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

$$\mathbb{E}[a_t | \mathcal{I}_t] = \mathrm{Cov}(a_t, y^t)[\mathrm{Var}(y^t)]^{-1} y^t. \tag{18}$$

Notice that

$$\mathrm{Cov}(y_t, y^t) = \begin{bmatrix} 1 & 0 & 0 & \cdots \end{bmatrix} \mathrm{Var}(y^t). \tag{19}$$

Further notice that (17) in combination with condition (ii) gives $\mathrm{Cov}(a_t, y^t) = \mathrm{Cov}(y_t - \tau_t, y^t) = \mathrm{Cov}(y_t, y^t)$. We can thus use (19) to substitute out $\mathrm{Cov}(a_t, y^t)$ in (18) to get

$$\mathbb{E}[a_t | \mathcal{I}_t] = y_t.$$

We conclude that as long as conditions (i) and (ii) hold, there exists a simple information-structure $\{\mathcal{I}_t\}$ that implements $\mathcal{T}_t$. 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 $\mathcal{T}_t$.[8]

The full proof in the appendix generalizes to dynamic economies involving many equations, variables, and information sets. As $\mathcal{E}(\mathcal{T})$ is an equilibrium in the primal economy, we can use the logic above to implement the beliefs implied by the primal economy for each information set, assured that the remaining equilibrium conditions of the incomplete-information economy hold by construction. Moreover, the full proof also allows for arbitrary minimal information requirements.

Remarks   Although our notation in presenting Theorem 1 is motivated by our model economy, the proof of the theorem is generic and can be applied to virtually any rational expectations DSGE model. Nevertheless we make a few assumptions, some of them implicit, that are worth discussing.

First, we require stationarity of $\mathcal{E}$. On the one hand, this rules out non-stationary processes of $\mathcal{T}_t$. On the other hand, this requires the primal economy to be stationary. In many

---

[8]The argument is related to the one given in Bergemann and Morris (2016) who show the equivalence between Bayes correlated equilibria and static Bayesian games with incomplete information. Our approach of formulating a primal economy and characterizing implementability in terms of a simple orthogonality condition is more general, however, as it straightforwardly applies to dynamic economies and allows for arbitrary minimal information requirements.

cases, an appropriate transformation can be used to induce stationarity in the primal economy, even when the economy is fundamentally non-stationary. E.g., in our case, it suffices to define the primal economy in terms of the output gap $\hat{y}_t$ as in (13)–(15), ensuring stationarity of $\mathcal{E}$ as long as $\mathcal{T}$ is stationary.

Second, while we assume $\mathcal{E}$ to be Gaussian, the assumption is not needed when one is only interested in implementing the auto-covariance structure of $\mathcal{E}_t$. In our empirical application, we make sure that our estimator indeed only uses information regarding the covariance structure, so that we do not need to make any distributional assumptions regarding $\mathcal{T}$ to invoke our theoretical results.

Third, while in our case each wedge corresponds to a unique information set, Theorem 1 applies to settings in which multiple wedges are associated with a single information set. Formally, the theorem extends to such cases by treating all expectations and wedges as vectors. An immediate corollary is that whenever two distinct information sets $\mathcal{I}_1$ and $\mathcal{I}_2$ share the same minimal information requirement Θ, the set of information-wedges supported by $(\mathcal{I}_1, \mathcal{I}_2)$ is identical to the one supported under the additional requirement that $\mathcal{I}_1 = \mathcal{I}_2$. In our case this implies that imposing common knowledge within islands places no additional restrictions on $\mathcal{T}_t$.[9]

Forth, while we take the minimal information-sets $\Theta_{b,t}$ and $\Theta_{i,t}$ to be as in (5) and (6), the theorem applies to any minimal information requirement desired by the researcher. The only restriction is that for each wedge τ the corresponding set Θ must at least contain the full history of all equilibrium expectations corresponding to the wedge (unless the researcher wants to relax the assumption of perfect recall).

3.3   Wedges in the aggregate economy

If the researcher is interested in the ability of the incomplete information model to match aggregate data, as we are in this paper, then the crucial question is to what degree the orthogonality conditions in Theorem 1 restrict the set of feasible dynamics for the aggregate wedges $\bar{\mathcal{T}}_t \equiv (\tau_t^b, \tau_t^f, \tau_t^h)$. From (5), we have

$$\mathbb{E}[\tau_t^b x_{t-\bar{h}-j}] = 0, \ \forall j \geq 0. \tag{20}$$

$$\mathbb{E}[\tau_t^b i_{t-j}] = 0, \ \forall j \geq 0. \tag{21}$$

---

[9]Despite it being w.l.o.g. in terms of $\mathcal{T}_t$, we do not impose common knowledge here, because it restricts the means to implement a given $\mathcal{T}_t$, ruling out the particular narrative that we find most plausible in light of our estimation below.

Similarly, orthogonality between $X_{t-\bar{h}}$ and $(\tau_{i,t}^f, \tau_{i,t}^h)$ implies

$$\mathbb{E}[\tau_t^f x_{t-\bar{h}-j}] = 0, \ \forall j \geq 0 \tag{22}$$

$$\mathbb{E}[\tau_t^h x_{t-\bar{h}-j}] = 0, \ \forall j \geq 0. \tag{23}$$

The presence of local price and productivity shocks guarantees that subject to these restrictions, any aggregate wedge process $\bar{\mathcal{T}}_t$ can be supported by sufficiently large local shocks.

To build an intuition why this is the case, consider a variant of the example given above, where $y_i = a_i + \tau_i$ and $\Theta_i = \{y_i\}$ for $i \in [0,1]$. The orthogonality condition in Theorem 1 imposes two restrictions on the distribution of (y_i, a_i, τ_i). First, it implies that $\mathrm{Var}[a_i] = \mathrm{Var}[y_i] + \mathrm{Var}[\tau_i]$, constraining the variance of τ_i to satisfy $\mathrm{Var}[\tau_i] \leq \mathrm{Var}[a_i]$. Second, it pins down the covariation of τ_i with a_i, $\mathrm{Cov}[a_i, \tau_i] = -\mathrm{Var}[\tau_i]$. Now consider $(\bar{y}, \bar{a}, \bar{\tau}) = \int_0^1 (y_i, a_i, \tau_i) \mathrm{d}i$. It can be shown that by varying the correlation of (τ_i, a_i) in the cross-section, one can implement any distribution over $(\bar{y}, \bar{a}, \bar{\tau})$ that satisfies

$$\mathrm{Var}[\tau_i] \geq \frac{(\mathrm{Cov}[\bar{\tau}, \bar{a}] + \mathrm{Var}[\tau_i])^2}{\mathrm{Var}[a_i] - \mathrm{Var}[\bar{a}]} + \mathrm{Var}[\bar{\tau}]$$

and $\bar{y} = \bar{a} + \bar{\tau}$. Clearly for any $\mathrm{Var}[\tau_i] > \mathrm{Var}[\bar{\tau}]$, the condition is non-binding for some sufficiently volatile a_i.[10] Intuitively, when $(a_i - \bar{a})$ is sufficiently volatile, we can support aggregate information-wedges using correlated errors about the idiosyncratic variations in the agents' objectives, regardless of what is the aggregate uncertainty about $\bar{a}$. The literature has proposed various channels that may give rise to such correlation patterns, including correlated noise shocks (Lorenzoni, 2009), correlated shocks to higher-order expectations (Angeletos and La'O, 2013), and the presence of informational sunspots (Benhabib, Wang and Wen, 2015).

Our empirical strategy in this paper is to proceed with the minimal restrictions on the aggregate wedges given in (20)–(23). With the estimated model at hand, we then construct processes for $\nu_{i,t}^1$, $\nu_{i,t}^2$, $z_{i,t}$, $\Delta a_{i,t}$ and for the island-specific wedges $\tau_{i,t}^f$ and $\tau_{i,t}^h$ that support the estimated processes for τ_t^f and τ_t^h subject to the restrictions in Theorem 1. The details are provided in Appendix A.2.

---

[10]E.g., fix $\mathrm{Var}[\tau_i] = 2\mathrm{Var}[\bar{\tau}] + \mathrm{Cov}[\bar{\tau}, \bar{a}]$. Then the condition holds for any $\mathrm{Var}[a_i] \geq \mathrm{Var}[\bar{a}] + 4(\mathrm{Var}[\bar{\tau}] + \mathrm{Cov}[\bar{\tau}, \bar{a}])$, which also suffices to satisfy the idiosyncratic variance bound $\mathrm{Var}[\tau_i] \leq \mathrm{Var}[a_i]$ (the covariance condition between τ_i and a_i holds by construction).

4 Econometric Methodology

In this section, we describe our strategy for estimating the incomplete-information economy. Building on the equivalence result in Theorem 1, our empirical approach formulates the information structure directly in terms of a stochastic process for the information wedge $\bar{\mathcal{T}}_t$. While practical concerns lead us to adopt a specific parametric specification for $\bar{\mathcal{T}}_t$, the approach is essentially non-parametric regarding the underlying structure of information. Since we do not have strong priors with regard to the precise channels through which agents collect their information (and even less with regard to the distributional properties of the noise terms associated with these channels), we believe that this is a natural starting point for an empirical investigation.

The relevant structural parameters in our model are β, α, ϕ_y, ϕ_π and the composite parameter κ. We fix these parameters in our estimation procedure as they are only weakly identified given the specification of $\bar{\mathcal{T}}$ adopted below. The discount factor β is set equal to 0.99, consistent with an average real return on assets of 4 percent per year. The production parameter α is estimated from (1) via OLS, using the utilization-adjusted productivity series from Fernald (2014) to measure a_t, yielding an estimate of 0.9. The Taylor rule coefficients ϕ_y and ϕ_π are set to 0.005 and 2, broadly consistent with the estimated parameter values in Blanchard, L'Huillier and Lorenzoni (2013) and Smets and Wouters (2007). The Philips curve coefficient κ is set equal to 0.025.[11]

As is common in the DSGE literature we restrict the stochastic process for $\bar{\mathcal{T}}_t$ to be first-order auto-regressive, so

$$\bar{\mathcal{T}}_t = \Lambda \bar{\mathcal{T}}_{t-1} + \omega_t,$$

where ω_t is i.i.d. across time with zero mean and covariance matrix Ψ. In our baseline setting, we consider matrices Ψ that are rank two as a third independent innovation to ω_t is found to only marginally improve the fit and is statistically insignificant. To allow agents to be potentially unaware of productivity shocks, we also allow ω_t to correlated with the productivity innovations ϵ_t. The joint covariance matrix is denoted $\tilde{\Psi} = \mathrm{Var}[(\omega_t, \epsilon_t)]$. In total, this gives us 18 parameters that are to be estimated, $\gamma \equiv \{vec(\Lambda), vech(\tilde{\Psi})\}$.

Let Γ denote the set of parameters consistent with the implementability conditions given in (20)–(23) for $\bar{h} = 32$. We estimate the model parameters γ using the generalized method of moments (GMM) to minimize the distance between the model's covariance structure and

---

[11]The value is based on an output elasticity of the flex-price target ξ set equal to 0.15 as suggested by Woodford (2011, Ch. 3) and a value for λ equal to 2/3 implying an average price duration of 3 quarters.

the data, subject to $\gamma \in \Gamma$.[12] Let

$$\tilde{\Omega}_T = vech\{\mathrm{Var}[(\tilde{x}_t^d, \dots, \tilde{x}_{t-k}^d)]\},$$

denote the empirical auto-covariance matrix of quarterly, frequency-filtered data $\tilde{x}_t^d$ on US real output, inflation, the FED funds rate and employment, where auto-covariances are computed for up to $k = 8$ quarters.[13] For the filtering, we use the Baxter and King (1999) approximate high-pass filter with a truncation horizon of 32 quarters; i.e., $\tilde{x}_t^d = HP_{32}(x_t^d)$ where x_t^d is the data-equivalent to $x_t = (\hat{y}_t, \pi_t, i_t, n_t)$.[14] Our estimator is then given by

$$\hat{\gamma} = \operatorname*{argmin}_{\gamma \in \Gamma}(\tilde{\Omega}_T - \tilde{\Omega}(\gamma))'W(\tilde{\Omega}_T - \tilde{\Omega}(\gamma)), \qquad (24)$$

where $\tilde{\Omega}(\gamma)$ is the model analogue to $\tilde{\Omega}_T$ and W is a weighting matrix set to an estimate of $[\mathrm{Var}\{T^{1/2}\tilde{\Omega}_T\}]^{-1}$ (see Appendix B.2 for details). To avoid biasing our results by a mismatch between model and data frequencies as discussed in Gorodnichenko and Ng (2009), we compute $\tilde{\Omega}(\gamma)$ by applying the same filtering procedure to the model as we do to the data. In Appendix B.1, we provide a closed-form transformation from $\Omega \equiv vech\{\mathrm{Var}[(\mathrm{d}x_t, \dots, \mathrm{d}x_{t-K})]\}$ to $\tilde{\Omega} = \Xi\Omega$ for a constant matrix Ξ and $K = k + 2\bar{\tau}$. Using the transformation, we can equivalently express (24) as

$$\hat{\gamma} = \operatorname*{argmin}_{\gamma \in \Gamma}(\Omega_T - \Omega(\gamma))'\tilde{W}(\Omega_T - \Omega(\gamma)), \qquad (25)$$

where now the unfiltered model is estimated (in first differences) on unfiltered data and the filtering is achieved by replacing W with $\tilde{W} \equiv \Xi'W\Xi$. Using (25) in place of (24), estimation becomes straightforward as the mapping from γ to $\Omega(\gamma)$ is available in closed form.

All confidence intervals and hypothesis tests are based on a bootstrapped distribution, $\{\hat{\gamma}_b\}_{b=1}^{B}$, with $B = 500$ replications (and analog distributions for the benchmark models discussed below). As bootstrap data generating process we use a VAR(10) estimated on $\mathrm{d}x_t$. In each sample b, we first construct W_b according to the steps described in Appendix B.2, and then use (25) to estimate $\hat{\gamma}_b$ where the target moments $(\Omega_b - \Omega(\gamma))$ are recentered about their population mean to adjust for overidentification.

---

[12] As for our parametric specification for $\bar{\mathcal{T}}_t$, (22)–(20) can not hold exactly unless $\bar{\mathcal{T}}_t = 0$ for all t, we allow for a small numerical deviation. The estimated model satisfies all orthogonality conditions within a tolerance of $2 \cdot 10^{-6}$.

[13] All data are defined quarterly and ranges from 1960 to 2012. Inflation is defined as the log-difference in the GDP deflator and employment is given by hours worked in the non-farm sector.

[14] The Baxter and King (1999) filter requires specification of a maximal lag-length $\bar{\tau}$ for the approximation. We set $\bar{\tau}$ to their recommended value of 12.

5 Empirical Results

In this section we explore the properties of the estimated model and assess its ability to account for business cycle comovements in the data. A full listing of the estimated parameters values is given in Table 5 in the appendix.

5.1 Predicted Moments

We begin by assessing the empirical performance of the estimated model. Figure 1 compares the predicted model moments with the targeted data moments. The dashed black lines show the empirical covariance structure $\tilde{\Omega}_T$ along with 90 percent confidence intervals (depicted by the shaded areas). The solid blue lines show the corresponding moments predicted by the model. Each row i and column j in the table of plots shows the covariances between $\tilde{x}_t^i$ and $\tilde{x}_{t-k}^j$ with lags $k \in \{0, 1, \ldots, 8\}$ depicted on the horizontal axis.

It is evident that the estimated model does a remarkable job at capturing the auto-covariance structure of the four time series. In particular, the model captures the positive contemporaneous comovement of output, hours, and inflation visible in the data, which is typically difficult for productivity-driven models to accommodate. The estimated model also does an excellent job at capturing the autocorrelation structure found in the data, in particular the rising profile of inflation's comovement with lagged GDP (second row, first column), as well as the falling autocorrelations of GDP, inflation, and hours (along the diagonal).

A test of the model's over-identifying restrictions confirms the visual impressions. Based on the estimated distance to the data ($J = 26.0$), we do not reject the validity of the model at a p-value of 24.4 percent (obtained via the bootstrap described above). In comparison, the full-information model where we restrict γ to satisfy $\Lambda = \Psi = 0$ is clearly unable to account for the empirical covariance structure.[15] Concordantly, the parameter restriction $\Lambda = \Psi = 0$ is rejected at a p-value of less than 0.01 percent.

Essential for the empirical performance of the model is that the estimated information-wedges can be correlated across time and among each other (see Table 1). In particular, there is a strong negative correlation between τ_t^f and both τ_t^h and τ_t^b. While most business cycle models generate shocks that are correlated across time, it is typically difficult for full-

---

[15]With the exception of output, all targeted moments are zero under the full-information model since potential output is constant when productivity follows a random walk process. We also explored whether a more flexible process for a_t that allows for fluctuations in the natural rate can improve the performance of the full-information model, but rejected this possibility as the predicted comovement patterns are at odds with the empirical covariance structure.
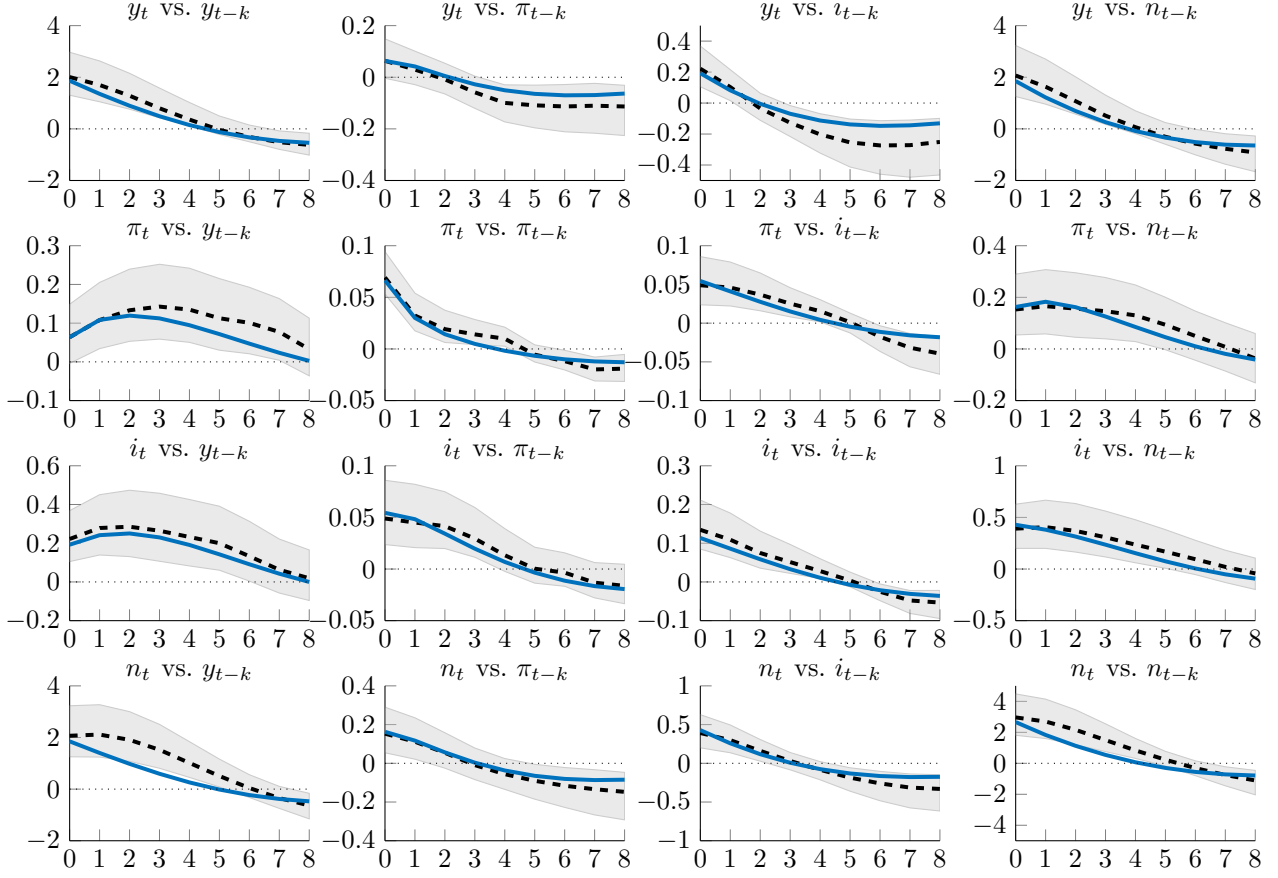
Figure 1: Business cycle comovements in the data and predicted by the estimated model. Note.—Dashed black lines show the empirical covariance structure $\tilde{\Omega}_T$ together with 90 percent confidence intervals depicted by the shaded areas. Solid blue lines show the corresponding model moments $\tilde{\Omega}(\hat{\gamma})$. Each row i and column j in the table shows the covariances between $\tilde{x}_t^i$ and $\tilde{x}_{t-k}^j$ with lags $k \in \{0, 1, \ldots, 8\}$ depicted on the x-axis.

information models with structurally uncorrelated shocks to generate perturbations that are correlated across equations. By contrast, the incomplete-information wedges are naturally correlated, both because information can be correlated across different types of agents and because the objectives (3), (9) and (10) are linked through strategic interdependences. To see how strategic links may give rise to a correlation in the wedges, suppose, e.g., that $\hat{y}_t$ increases due to optimistic households (an increase in τ_t^h). If they were fully aware of household's optimism, firms and the monetary authority would optimally increase, respectively, prices and the interest rate. When incompletely aware of the change in household expectations, firms and the central bank underreact, leading to negative movements in τ_t^b and τ_t^f. In Section 6.2, we demonstrate that the correlation patterns in Table 1 can indeed be accounted for by an imperfect awareness of one another's errors.

Table 1: Correlation structure of estimated information wedges

	Standard deviation	First-order autocorr.	Contemporaneous correlation			
			with τ_t^b	with τ_t^f	with τ_t^h	with ϵ_t
τ_t^b	0.43	0.28	1.00	·	·	·
τ_t^f	0.32	0.44	-0.70	1.00	·	·
τ_t^h	0.69	0.81	0.06	-0.31	1.00	·
ϵ_t	0.69	–	-0.14	-0.02	-0.24	1.00

5.2 Properties of the estimated model

We now decompose the process of the estimated information-wedges to derive impulse responses to productivity innovations and different types of expectational shocks. The estimated process for $\bar{\mathcal{T}}_t$ reflects two distinct channels through which incomplete information affects the dynamics of the economy. First, agents may have incomplete information regarding aggregate productivity, potentially modifying the economy's response to productivity shocks. Second, with incomplete information, agents can also make correlated errors that introduce an independent source of business cycle fluctuations.

Under the assumption that productivity is exogenous to $\bar{\mathcal{T}}_t$, we can separate out the two roles of incomplete information by projecting $\bar{\mathcal{T}}_t$ on current and past productivity shocks ϵ_t. We get

$$\bar{\mathcal{T}}_t = A(L)\epsilon_t + u_t, \tag{26}$$

where $A(L)\epsilon_t \equiv \mathbb{P}[\bar{\mathcal{T}}_t | \epsilon_t, \epsilon_{t-1}, \dots]$ denotes the projection. The lag-polynomial A identifies the average expectation errors in (11)–(13) due to agents being unaware about productivity innovations. The remaining residuals, u_t, identify purely expectational business cycle shocks.

Recall that under our baseline assumption, u_t is driven by only two independent innovations. Accordingly, let

$$u_t = B_1(L)\eta_{1,t} + B_2(L)\eta_{2,t}, \tag{27}$$

where B_1 and B_2 are lag-polynomials in nonnegative powers of L, and $\eta_{1,t}$ and $\eta_{2,t}$ are orthogonal white-noise processes. As in the structural VAR literature, B_1 and B_2 cannot be uniquely identified without additional identifying assumptions. To provide an economically interesting interpretation of our shocks, we identify the first shock $\eta_{1,t}$ as the shock that contributes most to unconditional output variation after controlling for the effect of the technology shock. The second shock $\eta_{2,t}$ then captures the remaining movements in the economy. The identification strategy is closely related to Uhlig's (2003; 2004) approach of

Table 2: Unconditional variance decomposition

	GDP	Inflation	Hours	Interest Rate
$\eta_{1,t}$	0.90	0.15	0.78	0.36
$\eta_{2,t}$	0.01	0.68	0.01	0.06
ϵ_t	0.09	0.18	0.22	0.58

Note.—Contributions are to unconditional variances in the frequency filtered model.

identifying shocks based on their contribution to some finite forecast error variance.

Table 2 shows the contributions of the three shocks to the unconditional variances of the high-pass filtered model variables. While productivity contributes substantially to fluctuations in interest rates, the majority of the fluctuations in output, inflation and employment are accounted for by the expectational shocks $\eta_{1,t}$ and $\eta_{2,t}$. Specifically, the first-ordered expectational shock is the dominant driver of output and employment fluctuations, whereas the second-ordered shock is the primary driver of inflation.

Impulse responses to productivity shocks   The solid lines in Figure 2 represent the responses of the estimated model to a one-standard-deviation innovation in productivity. The shaded confidence bands are based on the 5 and 95 percentile response of the bootstrapped distribution. The random walk assumption we have placed on productivity implies that under full information, output would immediately jump to the new potential output level (depicted by the dotted black line). As unanticipated permanent changes in potential output have no effect on the output gap, there would be no responses in inflation, hours, and the interest rate under full information.

With incomplete information, the responses are quite different. Output hardly moves on impact and then only slowly adjust to its new potential level over a course of roughly four years. The main cause of this slow response is a negative response in the household's information-wedge τ_t^h, reflecting slow learning of households about the change in productivity. In contrast to τ_t^h, the responses in τ_t^f and τ_t^b are not significantly different from zero, so that we cannot reject the hypothesis that firms and the monetary authority are perfectly aware of the change in potential output when it hits the economy. As a result, firms and the monetary authority are also approximately aware of the negative output gap induced by τ_t^h, so that inflation and interest rates fall in response to the change in household demand. Overall, the picture shows a delayed response to the productivity shock driven by slow adjustment of households, consistent with the VAR-based evidence by Basu, Fernald and Kimball (2006).
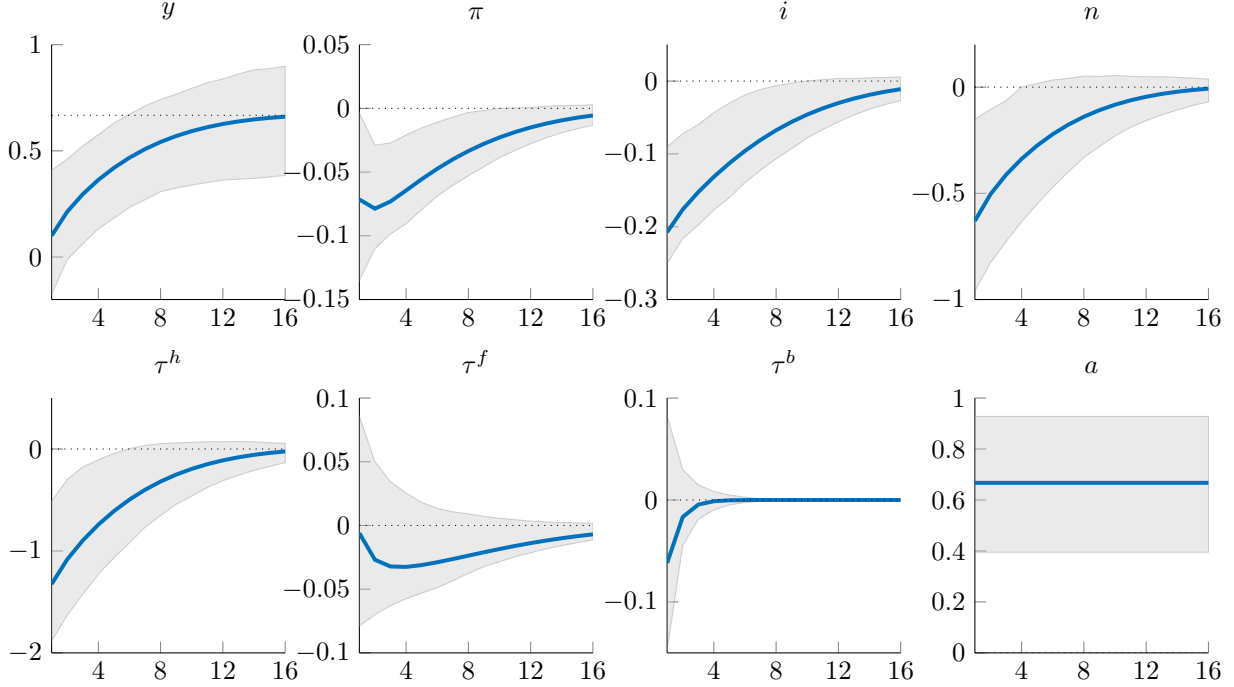
Figure 2: Impulse responses to a productivity shock. Note.—Responses are for a one-standard-deviation shock to ϵ_t and are depicted in percentage deviations from the steady state. Shaded regions are bootstrapped confidence intervals, bounded by the 5 and 95 percentiles of the bootstrapped distribution. Dotted black lines give the full-information responses.

Impulse responses to expectational shocks Figure 3 shows impulse response functions to a one-standard-deviation shock in $\eta_{1,t}$ (the shock contributing most to the observed output fluctuations). The propagation of the shock is mainly through a sharp and fairly persistent surge in household optimism. The resulting increase in household demand leads to a joint increase in output, inflation, interest rates, and employment. In terms of output, the shock implies a peak response of approximately 1 percent and has a half-live of two years. The responses in inflation and the interested rate are initially dampened due to a brief and comparably minor decline in τ_t^f. In Section 6.2 we demonstrate how this is in line with firms being unaware at first about the optimism of households. As firms learn about the presence of the household error, their responses converge to the full-information response, inducing an overall hump-shaped pattern in prices. Throughout, the bank-wedge τ_t^b is economically and statistically insignificant so that interest rates are set as if the monetary authority is perfectly informed about the household error τ_t^h. Finally, because productivity did not change, the response in employment mirrors the response in output, corresponding to a significant decrease in the labor wedge.

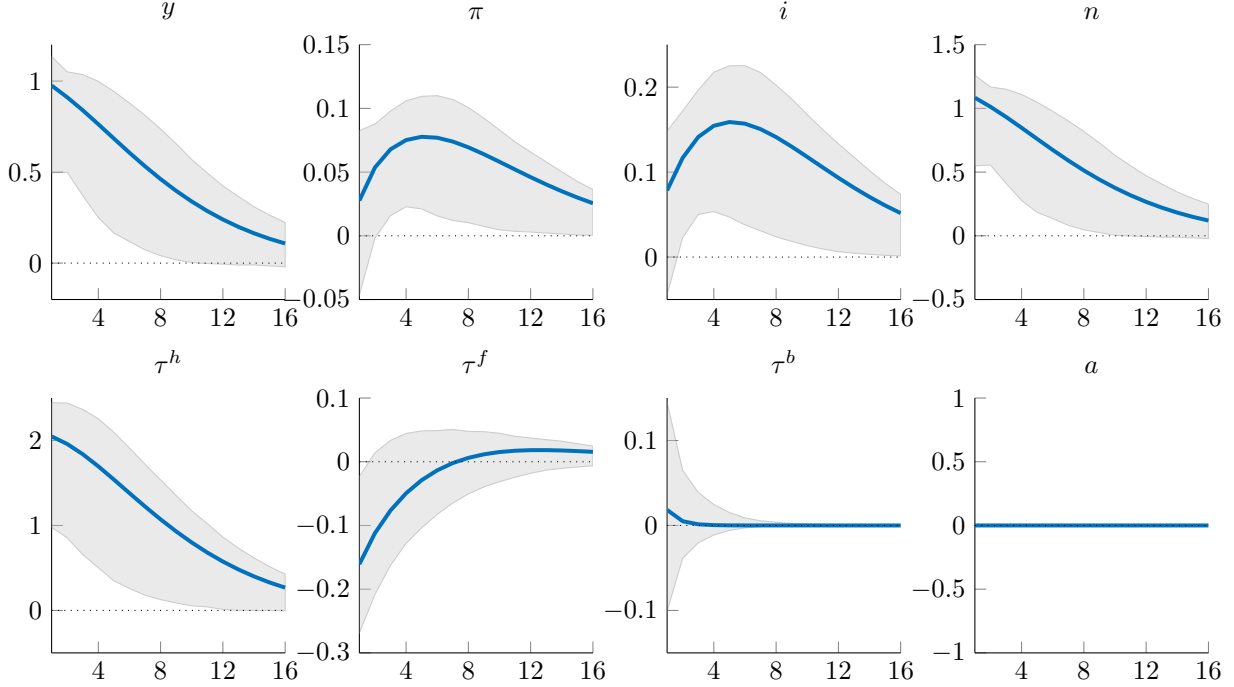We now turn to the second expectational shock $\eta_{2,t}$. Figure 4 gives the impulse responses.

Figure 3: Impulse responses to an expectational shock in $\eta_{1,t}$. Note.—Responses are for a one-standard-deviation shock to $\eta_{1,t}$ and are depicted in percentage deviations from the steady state. Shaded regions are bootstrapped confidence intervals, bounded by the 5 and 95 percentiles of the bootstrapped distribution.

The dominant driver of this shock is an expectational shock to firms yielding a sharp drop in inflation. The responses in output and interest rates are, by contrast, not significant. This is because the impact of τ_t^f on the optimal consumption and interest targets are almost exactly offset by corresponding responses in τ_t^h and τ_t^b, suggesting that households and the monetary authority are unaware of the shock to firms' expectations. A possible factor explaining this is the short-lived nature of the expectation-error made by firms, which has a half-life of less than one quarter and cedes to be statistically significant within two quarters, making it difficult for households and the monetary authority to learn about the firms' error in time.

6 Informational Primitives

In this section, we explore what the primal economy can teach us about the informational primitives in our model. We proceed in two steps. First, we use our estimates to test for the presence of various aggregate statistics in the information sets of agents. Second, we argue for a particular narrative of information transmission in the economy and demonstrate how, using the narrative, the estimated information-wedges can be implemented with a fully parametric information structure.
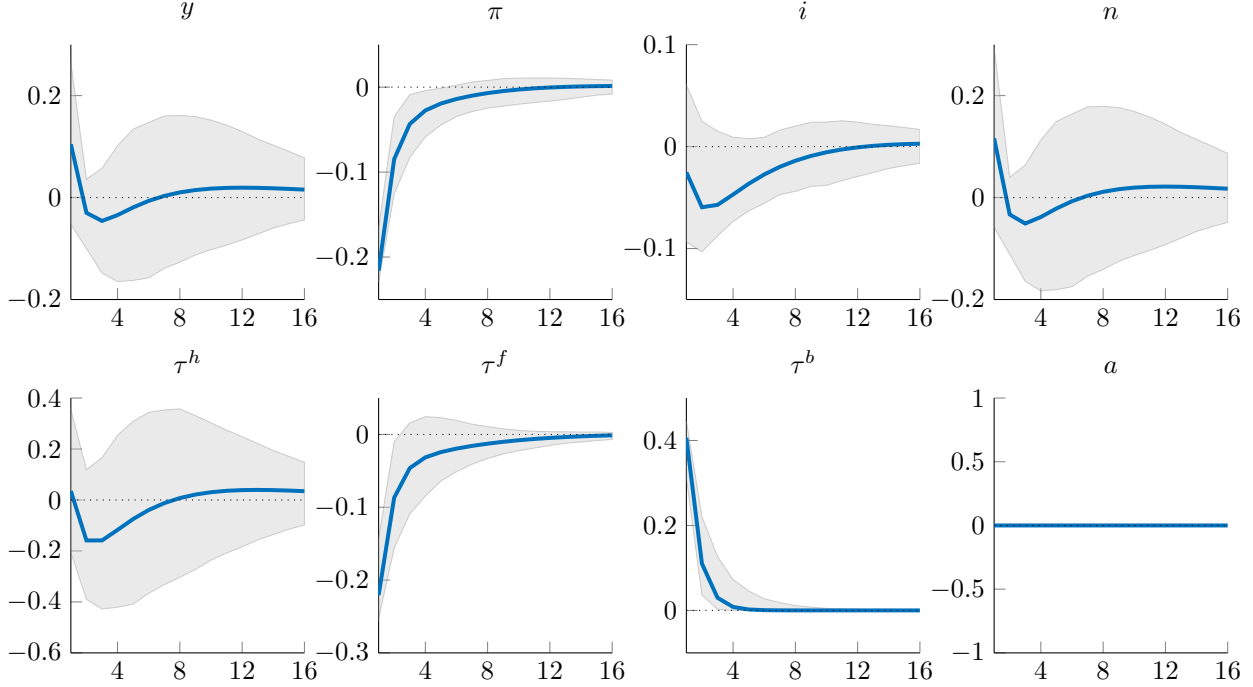
Figure 4: Impulse responses to an expectational shock in $\eta_{2,t}$. Note.—Responses are for a one-standard-deviation shock to $\eta_{2,t}$ and are depicted in percentage deviations from the steady state. Shaded regions are bootstrapped confidence intervals, bounded by the 5 and 95 percentiles of the bootstrapped distribution.

6.1   Knowledge of aggregate conditions

Figure 5 reports the (auto) correlation coefficients of the information wedges vis-à-vis output growth, inflation, the Federal funds rate and productivity growth. The shaded regions correspond to 90 percent confidence intervals. Since informational errors must be orthogonal to variables in the corresponding information set, these plots indicate what variables may and may not be observed by agents populating our estimated economy. By design the monetary authority is aware of the prevailing interest rate so that τ_t^b is uncorrelated with i_t at all lags. Since the correlation between τ_t^b and $(\mathrm{d}y_t, \mathrm{d}a_t)$ is statistically insignificant at all lags, we find that the monetary authority is also likely to be aware of contemporaneous output and productivity growth. By contrast, our test rejects that the monetary authority is aware of inflation within a 1-year horizon. Consistent with a lack of reliable real-time inflation statistics, the correlation is largest for current-quarter inflation and quickly drops to economically insignificant levels as data becomes available over the course of one year.

In the private sector, our test suggests that firms are unaware of aggregate demand within a 1-year horizon. With this exception, we cannot reject orthogonality of the firms' information wedges $\tau_{i,t}^f$ with respect to any of the aggregate statistics, indicating again that including these variables in the information of firms is consistent with an informational account of the
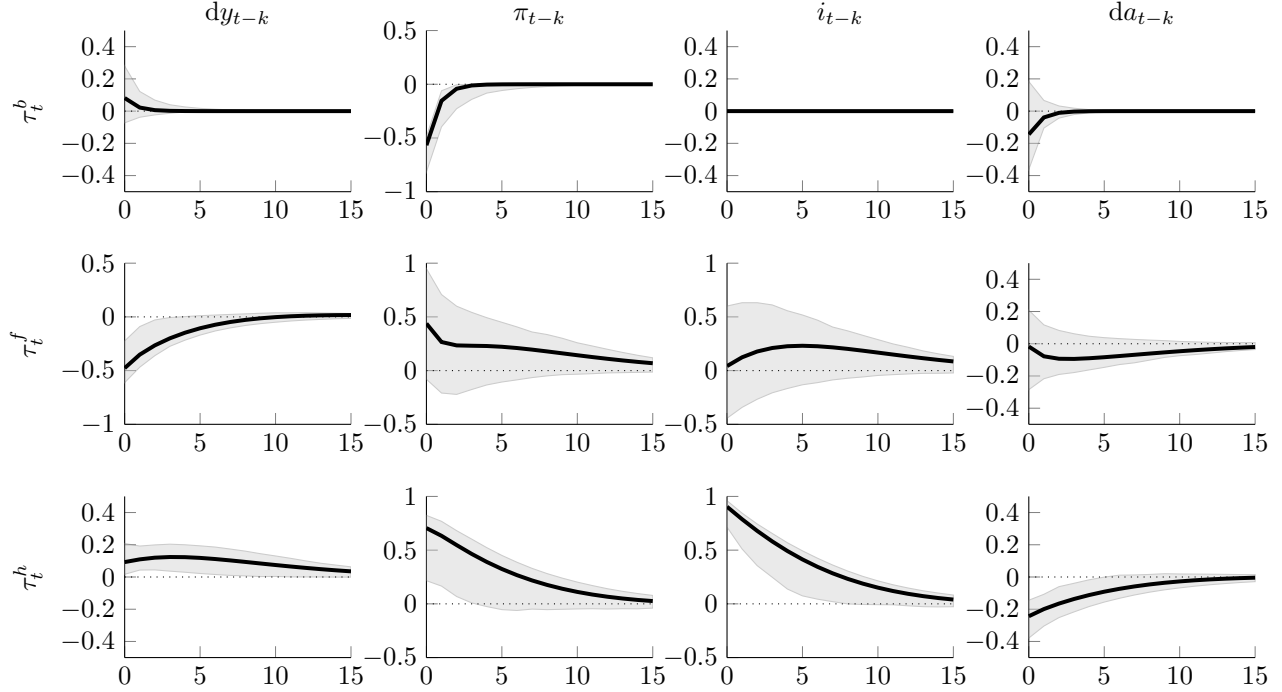
Figure 5: Correlation between information-wedges and aggregate statistics. Note.—The plot shows the (auto) correlation coefficients of the estimated information-wedges vis-à-vis output growth, inflation, the fed funds rate and productivity growth. The order of the autocorrelation is on the x-axis. Shaded areas depict 90 percent confidence intervals.

business cycle. In stark contrast, the households' information wedges significantly violate orthogonality with respect to all aggregate statistics at horizons of up to 2 years. While firms and the monetary authority have a relative good understanding of the aggregate state of the economy, our results thus suggest that households are generally unaware of contemporaneous economic conditions. The results are in line with forecast-based evidence which finds that households appear to be less informed than other agents in the economy (e.g., Carroll, 2003).

6.2 Interpretation and implementation

There are in general many information structures implementing a given process for $\bar{\mathcal{T}}_t$. Our description of the estimated economy above already hints at a particular narrative, in which the two expectational shocks $\eta_{1,t}$ and $\eta_{2,t}$ reflect intrinsic shocks to households' and firms' expectations, respectively. We now formalize this narrative and demonstrate how it can be implemented by an appropriate choice of information structure. We proceed in two steps. First, we develop the narrative by taking a stand on the origin of the estimated expectation errors inherent in $\bar{\mathcal{T}}_t$. This allows us to interpret the estimated equilibrium expectations, which are about endogenous objectives, in terms of expectations regarding certain exogenous

fundamentals. Second, we demonstrate how the expectation dynamics that we develop in the first step can be implemented by a particular information structure and characterize the island-specific noises needed to support them.

To streamline the narrative, our implementation focuses on the case where ϵ_t is perfectly known by firms, $(\epsilon_t, \eta_{1,t})$ is perfectly known by the monetary authority, and $\eta_{2,t}$ is fully unknown to households and the monetary authority. In addition, we truncate the response of τ_{t+j}^f to $\eta_{1,t}$ for $j \geq 8$ at zero to prevent overshooting of firms expectations in response to the household sentiment shock. All these modifications are within the confidence set of our original estimate. Figure 8 in the appendix shows the implemented impulse response functions in comparison to the original ones.

Interpreting the primal economy From (26) and (27), the aggregate wedges are given by

$$\tau_t^m = A^m(L)\epsilon_t + B_1^m(L)\eta_{1,t} + B_2^m(L)\eta_{2,t},$$

where for $m \in \{b, f, h\}$, A^m, B_1^m and B_2^m are lag-polynomials defined by the corresponding rows of A, B_1, and B_2. In the narrative that we develop, we interpret $B_1^h(L)\eta_{1,t}$ and $B_2^f(L)\eta_{2,t}$ as intrinsic fluctuations in household and firms expectations, driven by the "sentiment" shocks $\eta_{1,t}$ and $\eta_{2,t}$. By contrast, we interpret all remainder fluctuations in $\bar{\mathcal{T}}_t$, defined by $\{A^m\}_{m \in \{b,f,h\}}$, $\{B_1^m\}_{m \in \{b,f\}}$ and $\{B_2^m\}_{m \in \{b,h\}}$, as gradual learning on the part of (other) agents about these sentiments and aggregate productivity. According to this narrative, the economy is thus driven by two sentiment shocks and a shock to aggregate productivity, whereas the estimated comovement between τ_t^h, τ_t^f, τ_t^b and ϵ_t is accounted for by agent's being imperfectly aware of one another's errors.

The dynamics of the sentiments are evident from the impulse response of τ_t^h to a shock in $\eta_{1,t}$ as seen in Figure 3 and the response of τ_t^f to a shock in $\eta_{2,t}$ as seen in Figure 4. To see how we can interpret the remaining fluctuations in $\bar{\mathcal{T}}_t$ as gradual learning, consider the response in the average firms' error τ_t^f to the household sentiment shock $\eta_{1,t}$, given by $B_1^f(L)\eta_{1,t}$. Our approach attributes $B_1^f(L)$ to imperfect information among firms regarding $\eta_{1,t}$. Let $\{\pi_j\}$ denote the projection coefficients of inflation π_t onto $\{\eta_{1,t-j}\}$ as estimated in the primal economy, and analogously let $\{\pi_j^*\}$ denote the projection coefficients of firms' optimal response $(\pi_t - \tau_t^f)$ onto $\{\eta_{1,t-j}\}$. Using the definition of the aggregate wedge, the projection coefficients must satisfy

$$\pi_j \eta_{1,t-j} = \pi_j^* \bar{\mathbb{E}}[\eta_{1,t-j} | \mathcal{I}_{i,t}^f],$$
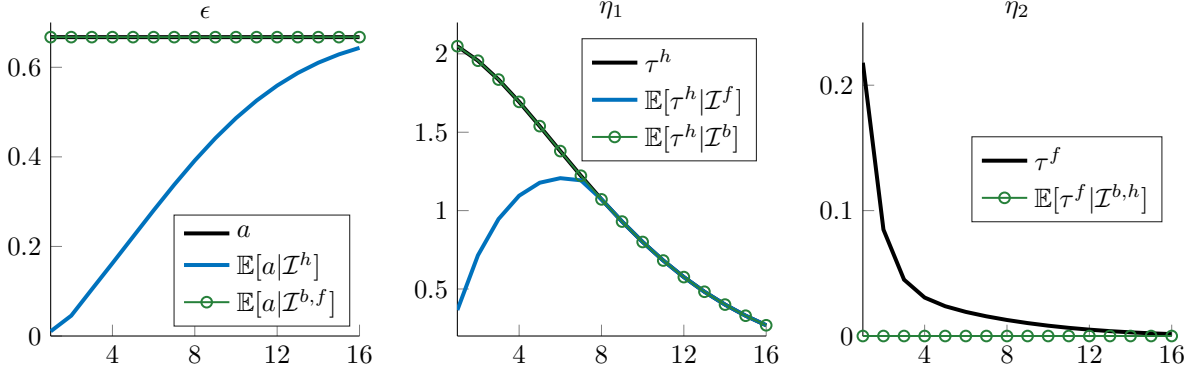
Figure 6: Expectation dynamics to the three model shocks.

where $\bar{\mathbb{E}}[\cdot|\mathcal{I}_{i,t}^f]$ denotes average expectations across firms.[16] Note that there is no uncertainty about the projection coefficients $\{\pi_j^*\}$. Intuitively, in equilibrium, firms understand how their objective is affected by a shock to household sentiments, taking into account the equilibrium response of households, the monetary authority and other firms. Their only uncertainty is about the realization $\eta_{1,t-j}$. Using the estimated inflation response $\{\pi_j\}$ to back out this uncertainty, we get

$$\bar{\mathbb{E}}[\eta_{1,t-j}|\mathcal{I}_{i,t}^f] = \frac{\pi_j}{\pi_j^*}\eta_{1,t-j}. \tag{28}$$

Analogous steps deliver the (average) expectations of households in response to productivity shocks, whereas the remaining expectations are trivial given the modifications described above.

Figure 6 displays the dynamics responses of agents' (average) expectations to one-standard-deviation innovations in ϵ_t, $\eta_{1,t}$ and $\eta_{2,t}$. The left panel reports the responses to a productivity shock. The expectations of firms and the monetary authority trace the response of productivity, reflecting an estimated response in inflation and interest rates that is consistent with full information. By contrast, households learn about productivity shocks only gradually, being on average aware of 7 percent of the realized innovation upon impact.

Next consider the responses to $\eta_{1,t}$. Again, the estimated response in interest rates is consistent with the monetary authority being perfectly aware of the intrinsic fluctuations in household sentiments. The response of firms' average expectations corresponds to an impact awareness of 13 percent and is subsequently converging to the realized response in household

---

[16]Here we exploit that according to our interpretation, $\mathbb{P}\{\bar{\mathbb{E}}[\Delta\pi_{i,t}^*|\mathcal{I}_{i,t}^f]|\eta_{1,t-j}\} = 0$; i.e., $\eta_{1,t}$ induces no correlated errors across firms about the island-specific components in their target prices. In attributing $B_1^f(L)\eta_{1,t}$ exclusively to incomplete information regarding $\eta_{1,t}$, we also implicitly assume that learning is independent across innovations $\eta_{1,t}$, $\eta_{1,t-1}$ and so on. In our implementation below, we provide a specific signal structure for which this is the case.

sentiments over the course of 8 quarters. Finally, the right panel depicts the response of firm sentiments to $\eta_{2,t}$. As discussed before, the estimated responses in y_t and i_t are consistent with households and the monetary authority being fully ignorant of the change in firm sentiments.

Implementation   We now demonstrate a specific information structure that implements the expectations shown in Figure 6 in the incomplete-information economy. To implement the expectations of the monetary authority, it suffices to set $\mathcal{I}_t^b = \{\epsilon_{t-s}, \eta_{1,t-s}\}_{s\geq 0} \cup \Theta_{b,t}$. The presence of ϵ_t and $\eta_{1,t}$ in the information set provides full information about productivity and household sentiments, whereas none of the available signals provides any useful information about $\eta_{2,t}$.[17]

To implement households' and firms' expectations, we need to design an information structure that on the one hand gives rise to intrinsic sentiment fluctuations and on the other hand induces gradual learning in line with the characterization above. In order to do this, we again use projections to isolate fluctuations due to each agent type's own sentiment shock, and then provide agents with noisy signals about remaining sources of aggregate volatility. As the implementation strategy is the same for firms and households, we only demonstrate it for firms.

Define

$$\pi_{i,t}^* \equiv \mathbb{E}_t[\kappa(\hat{y}_{i,t} + \mu_{i,t}) + \beta\pi_{i,t+1}^*]$$

as the objective of firms in island i at date t, and let

$$\mathring{\pi}_{i,t}^* \equiv \pi_{i,t}^* - \mathbb{P}[\pi_{i,t}^*|\eta_{1,t}, \eta_{1,t-1}, \dots] - \mathbb{P}[\pi_{i,t}^*|\epsilon_t, \epsilon_{t-1}, \dots]$$

be the objective after projecting out all variations driven by $\eta_{1,t}$ and ϵ_t. By construction, $\pi_{i,t} = \pi_{i,t}^* + \tau_{i,t}$, which given our decomposition can be rewritten as

$$\pi_{i,t} = \mathbb{E}[\mathring{\pi}_{i,t}^*|\mathcal{I}_{i,t}^f] + \mathbb{E}\left\{\mathbb{P}[\pi_{i,t}^*|\eta_{1,t}, \eta_{1,t-1}, \dots]|\mathcal{I}_{i,t}^f\right\} + \mathbb{E}\left\{\mathbb{P}[\pi_{i,t}^*|\epsilon_t, \epsilon_{t-1}, \dots]|\mathcal{I}_{i,t}^f\right\}. \qquad (29)$$

The first term in equation (29) above, $\mathring{\pi}_{i,t}^*$, is the portion of inflation not accounted for by firms' mistaken forecasts of other agents' beliefs, and will serve as the basis for our implementation of firm sentiments. The second and third term in (29) accordingly capture the learning dynamics of firms regarding household sentiments and aggregate productivity.

---

[17] While $\Theta_{b,t}$ contains information regarding the aggregate state of the economy with a lag of $\bar{h}$ quarters, this information is useless for the purpose of learning about (aggregate) firm sentiments, because τ_t^f is only driven by innovations within a horizon of $\bar{h}$ quarters.

To implement the firm sentiment shock, define

$$\mathring{\tau}_{i,t} = \mathbb{E}[\mathring{\pi}^*_{i,t}|\mathcal{I}_{i,t}] - \mathring{\pi}^*_{i,t}$$

as the expectation error regarding $\mathring{\pi}^*_{i,t}$. By construction, any distribution of $\mathring{\tau}_{i,t}$ across islands that satisfies

$$\int_0^1 \mathring{\tau}_{i,t}\,\mathrm{d}i = B_2^f(L)\eta_{2,t} \tag{30}$$

will implement the firm sentiments $B_2^f(L)\eta_{2,t}$. In order to implement such a $\{\mathring{\tau}_{i,t}\}$, we use the precise logic of our proof to Theorem 1 and equip firms with a signal

$$\omega_{i,t} = \mathring{\pi}^*_{i,t} + \mathring{\tau}_{i,t}.$$

As is demonstrated in the example in Section 3.2, this suffices to implement any desired $\mathring{\tau}_{i,t}$ as long as $\omega_{i,t} \perp \mathring{\tau}_{i,t}$ and $\mathcal{I}^f_{i,t} \perp \mathring{\tau}_{i,t}$. Subject to these restrictions, we can construct a particular process for $\mathring{\tau}_{i,t}$ that satisfies (30) as long as there are sufficiently large local shocks (see Appendix A.2 for details). The properties of the local shocks needed to support our implementation strategy are further discussed below.

We are left to implement the learning dynamics regarding household sentiments and productivity. Suppose $\mathcal{I}^f_{i,t} = \{a_{t-s}, \omega_{i,t-s}, \chi_{i,t}\}_{s\geq 0}\cup\Theta_{i,t}$. The presence of a_t in the information set induces perfect information regarding productivity, whereas the signals $\{\chi_{i,t-s}\}_{s\geq 0}$ govern learning about $\eta_{1,t}$ (by construction no other element in $\mathcal{I}^f_{i,t}$ contains useful information about $\eta_{1,t}$). Let

$$\chi_{i,t} \sim \mathcal{N}\left\{\begin{pmatrix} \eta_{1,t} \\ \vdots \\ \eta_{1,t-\bar{h}} \end{pmatrix}, \begin{pmatrix} \sigma^2_{f,0} & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \sigma^2_{f,\bar{h}} \end{pmatrix}\right\},$$

where we assume that the signal noises are i.i.d. across islands. Then, assuming $\eta_{1,t}$ to be Gaussian, we have that

$$\bar{\mathbb{E}}[\eta_{t-j}|\mathcal{I}^f_{i,t}] = \frac{\sum_{k=0}^j \sigma^{-2}_{f,k}}{1 + \sum_0^j \sigma^{-2}_{f,k}}\eta_{t-j}, \tag{31}$$

where $\mathrm{Var}\{\eta_{1,t}\}$ has been normalized to unity. Equating equation (31) with (28), we find the signal variances $\{\sigma^2_{f,j}\}_{j=0}^{\bar{h}}$ to implement the process for firm sentiments,

$$\sigma^{-2}_{f,j} = \frac{\pi_j}{\pi^*_j - \pi_j}\sigma^{-2}_{\eta,1} - \sum_{k=0}^{j-1}\sigma^{-2}_{f,k}.$$
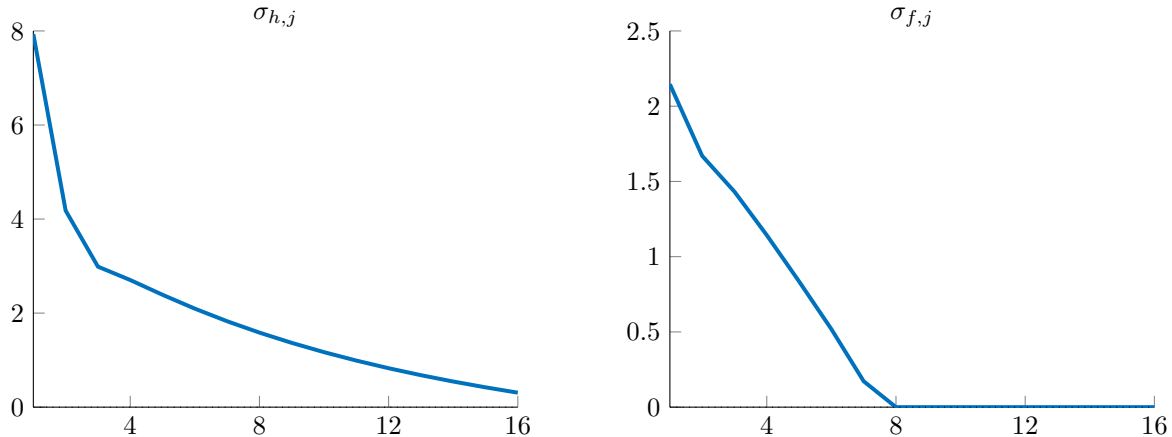
Figure 7: Signal structure implementing the estimated learning dynamics. The left panel depicts the standard deviation of the new signals available to households at time t regarding a productivity innovation at time $t - j$. The right panel depicts the standard deviation of the new signals available to firms at time t regarding an innovation in household sentiments at time $t - j$.

Figure 7 reports the standard deviations of the signals that implement the learning dynamics of firms regarding to $\eta_{1,t}$ and of households regarding to a_t. The latter are computed following precisely the same steps as outlined for firms above. It can be seen that the newly arriving signals are more precise over time, reflecting learning dynamics that converge to the truth at a faster rate than the one implied where an independently but identical distributed signals regarding past innovations becomes available at each date.

Volatility of islands-specific noise terms   We now return to the question of what volatility for the idiosyncratic noise terms, $\{\Delta a_{i,t}, z_{i,t}, \nu_{i,t}^1, \nu_{i,t}^2\}$, is required to implement the processes for aggregate expectation described above. In practice, many different processes for these shocks are able to support the estimated belief processes. For our implementation, we treat these shocks as independent MA (32) processes that can be arbitrarily correlated with the island-specific component of the information wedges (see the appendix for details). We then search numerically for a feasible implementation of estimated belief processes that minimizes the sum of unconditional variances of these shocks. Table 3 reports the corresponding standard deviations for these shocks.[18] While local conditions must exhibit substantial volatility, the order of magnitude for these volatilities are on par with those of aggregate conditions in the economy. The values also fall well within the plausible range for the volatility of local conditions, which are often calibrated or estimated to be much larger.

---

[18]We also repeat the exercise for the originally estimated economy, yielding very similar numbers (reported in Table 4 in the appendix).

Table 3: Standard deviations of island-specific noise terms

Noise process	$\Delta a_{i,t}$	$z_{i,t}$	$\nu_{i,t}^1$	$\nu_{i,t}^2$
Standard deviation	0.89	4.34	1.49	1.28

7   Summary

We have established the equivalence between a primal economy characterized by a set of reduced-form wedges and the class of economies driven by incomplete information. Applying our result, we have shown how to estimate a macroeconomic model with incomplete information without parametric assumptions on information structures. Our approach is, at once, straightforward to use and can be easily adapted to myriad contexts. We use the approach to identify ex post the empirically relevant information-frictions. In the context of a new-Keynesian model, our analysis points to incomplete information on the part of households as the most important ingredient for explaining the data.

A   Mathematical appendix

A.1   Proof of Theorem 1

Consider any information-wedge $\tau \in \mathcal{T}$. Let $(\hat{a}, \mathcal{I}^*)$ denote the corresponding expectation target and full-information set contained in $\mathcal{E}$, so that

$$y \equiv \mathbb{E}[\hat{a}|\mathcal{I}^*] + \tau \tag{32}$$

defines the equilibrium "belief" implied by the primal economy. Also let Θ be the corresponding minimal information requirement on $\mathcal{I}$ in the incomplete information economy. We want to show that conditions (i) and (ii) are jointly necessary and sufficient for the construction of some $\mathcal{I} \in \Theta$ such that

$$\mathbb{E}[\hat{a}|\mathcal{I}] = y. \tag{33}$$

Necessity   The necessity is immediate, since optimal inference requires that expectation errors are orthogonal to variables in the information set and are unpredictable. E.g., combining (32) and (33), implementation requires

$$\tau = \mathbb{E}[\hat{a}|\mathcal{I}] - \mathbb{E}[\hat{a}|\mathcal{I}^*]. \tag{34}$$

Taking expectations over (34) yields $\mathbb{E}[\tau] = 0$. Similarly, postmultiplying (34) by $\Theta' \subseteq \mathcal{I} \subseteq \mathcal{I}^*$ gives $\mathbb{E}[\tau\Theta'] = \mathbb{E}[\hat{a}\Theta'|\mathcal{I}] - \mathbb{E}[\hat{a}\Theta'|\mathcal{I}^*]$ and, hence, $\mathbb{E}[\tau\Theta'] = 0$.

Sufficiency   We demonstrate sufficiency by construction. Let $\mathcal{I} = \Theta$. Notice that in dynamic settings, the constructed information-set $\mathcal{I}$ is recursive whenever Θ is recursive (as we are assuming in this paper). Recursivity of Θ thus ensures that $\mathbb{E}[\hat{a}|\mathcal{I}]$ is fully consistent with the dynamic nature of agents' belief process.

Let $\tilde{\Theta} \equiv \Theta \setminus \{y\}$, where, by assumption, $y \in \Theta$ (see the discussion in the main text), and let $a \equiv \mathbb{E}[\hat{a}|\mathcal{I}^*]$. From the law of iterated expectations, we have $\mathbb{E}[\hat{a}|\mathcal{I}] = \mathbb{E}[a|\mathcal{I}]$ as $\mathcal{I} \subseteq \mathcal{I}^*$. Projecting a onto $(y, \tilde{\Theta})$ we thus have[19]

$$\mathbb{E}[\hat{a}|\mathcal{I}] = \begin{bmatrix} \Sigma_{ay} & \Sigma_{a\tilde{\Theta}} \end{bmatrix} \begin{bmatrix} \Sigma_{yy} & \Sigma_{y\tilde{\Theta}} \\ \Sigma'_{y\tilde{\Theta}} & \Sigma_{\tilde{\Theta}\tilde{\Theta}} \end{bmatrix}^{-1} \begin{bmatrix} y \\ \tilde{\Theta} \end{bmatrix},$$

---

[19]When the vector Θ contains co-linear variables, the proof follows after replacing $\Sigma_{\Theta\Theta}^{-1}$ with the generalized inverse $\Sigma_{\Theta\Theta}^\dagger$ and using the standard properties of the projection matrix, $\Sigma_{\Theta\Theta}\Sigma_{\Theta\Theta}^\dagger$.

where we use Σ_{ab} as shorthand for $\text{Cov}[a, b']$. Combining (32) with condition (ii) yields

$$\text{Cov}(a, y') = \text{Cov}(y - \tau, y') = \text{Cov}(y, y')$$
$$\text{Cov}(a, \tilde{\Theta}') = \text{Cov}(y - \tau, \tilde{\Theta}') = \text{Cov}(y, \tilde{\Theta}').$$

Noting that

$$\begin{bmatrix} \Sigma_{yy} & \Sigma_{y\tilde{\Theta}} \end{bmatrix} = \begin{bmatrix} I & 0 \end{bmatrix} \begin{bmatrix} \Sigma_{yy} & \Sigma_{y\tilde{\Theta}} \\ \Sigma'_{y\tilde{\Theta}} & \Sigma_{\tilde{\Theta}\tilde{\Theta}} \end{bmatrix},$$

we therefore get

$$\mathbb{E}[\hat{a}|\mathcal{I}] = \begin{bmatrix} I & 0 \end{bmatrix} \begin{bmatrix} y \\ \tilde{\Theta} \end{bmatrix} = y.$$

As the proof applies to any $\tau \in \mathcal{T}$, we conclude that as long as conditions (i) and (ii) holds, we can replicate $\mathcal{T}$ by including an exogenous signal $a + \tau$ into each information set that has the same distributional properties as the primal "belief" y. Moreover, because $\mathcal{E}(\mathcal{T})$ is an equilibrium in the primal economy, all equilibrium conditions in the incomplete-information economy hold by construction, concluding the proof of the theorem.

A.2 Implementation of aggregate wedges

For the aggregate firm and household wedges τ_t^f and τ_t^h to be implementable, they must be supported by a set of island-specific wedge processes $\{\tau_{i,t}^f, \tau_{i,t}^h\}$ that is consistent with the additional orthogonality conditions imposed by the island-specific portion of $\Theta_{i,t}$ (orthogonality to the aggregate portion is ensured by (20)–(23)). Exploiting the structure of the local economy, the island-specific portion $\{a_{i,t}, c_{i,t}, p_{i,t}, \bar{p}_{i,t}, i_{i,t}, w_{i,t}\}$ is informationally equivalent to

$$\mathcal{S}_{i,t} \equiv \begin{pmatrix} \mathrm{d}a_t \\ \hat{y}_t \\ \pi_t \\ \pi_t \\ i_t \\ \theta^{-1}\mathrm{d}y_t + \pi_t \end{pmatrix} + \begin{pmatrix} \mathrm{d}\Delta a_{i,t} \\ \Delta c_{i,t} - \Delta a_{i,t} \\ \mathrm{d}\nu_{i,t}^1 \\ \Delta \pi_{i,t} \\ \nu_{i,t}^2 \\ \mathrm{d}z_{i,t} \end{pmatrix},$$

where we use $\mathcal{S}_t$ and $\Delta\mathcal{S}_{i,t}$ to refer to the first (aggregate) and second (island-specific) term, respectively. Similar, let $\Delta\tau_{i,t}^m \equiv \tau_{i,t}^m - \tau_t^m$ denote the island-specific component of firm and

household wedges, $m \in \{f, h\}$. The orthogonality condition then requires

$$\text{Cov}\{(\Delta\tau_{i,t}^f, \Delta\tau_{i,t}^h), \Delta\mathcal{S}_{i,t}\} = -\text{Cov}\{(\tau_t^f, \tau_{i,t}^h), \mathcal{S}_{t-j}\} \quad \text{for all} \quad j \geq 0, \tag{35}$$

where the right-hand side is pinned down by our estimation results. Notice that the only endogenous terms in $\Delta\mathcal{S}_{i,t}$ are $\Delta c_{i,t}$ and $\Delta\pi_{i,t}$. The dynamics of $\Delta c_{i,t}$ and $\Delta\pi_{i,t}$ are governed by the stochastic properties of $(\Delta\tau_{i,t}^f, \Delta\tau_{i,t}^h, \Delta a_{i,t}, z_{i,t}, \nu_{i,t}^1, \nu_{i,t}^2)$. Implementation of (τ_t^f, τ_t^h) therefore essentially amounts to constructing a joint-process for $(\Delta\tau_{i,t}^f, \Delta\tau_{i,t}^h, \Delta a_{i,t}, z_{i,t}, \nu_{i,t}^1, \nu_{i,t}^2)$, which is consistent with (35).

In particular, the dynamics of the "Δ-economy" are determined by the following set of equilibrium conditions[20]:

$$\Delta c_{i,t} = \mathbb{E}_t[b\Delta c_{i,t+1} - \Delta\tau_{i,t+1}^h + \mathrm{d}\nu_{i,t+1}^1 - \nu_{i,t}^2] + \tau_{i,t}^h \tag{36}$$

$$\Delta\pi_{i,t} = \mathbb{E}_t[\kappa(\Delta\hat{y}_{i,t} + \mu_{i,t}) + \beta(\Delta\pi_{i,t+1} - \Delta\tau_{i,t+1}^f)] + \tau_{i,t}^f \tag{37}$$

where

$$\mu_{i,t} = \frac{\alpha}{\zeta+1}\left[\Delta c_{i,t} - \Delta\hat{y}_{i,t} - \Delta a_{i,t} + \nu_{i,t}^1 - \Delta\bar{p}_{i,t}\right] \tag{38}$$

$$\Delta\hat{y}_{i,t} = \theta(z_{i,t} - \Delta\bar{p}_{i,t}) - \Delta a_{i,t} \tag{39}$$

$$\Delta\bar{p}_{i,t} = \Delta\bar{p}_{i,t-1} + \Delta\pi_{i,t}. \tag{40}$$

Given a stochastic process for $(\Delta\tau_{i,t}^f, \Delta\tau_{i,t}^h, \Delta a_{i,t}, z_{i,t}, \nu_{i,t}^1, \nu_{i,t}^2)$, the rational expectations equilibrium to the system (36)–(40) can be solved using standard methods. For our implementation, we let $(\Delta\tau_{i,t}^f, \Delta\tau_{i,t}^h, \Delta a_{i,t}, z_{i,t}, \nu_{i,t}^1, \nu_{i,t}^2)$ be a MA (32) process[21], where we restrict innovations so that the fundamentals $(\Delta a_{i,t}, z_{i,t}, \nu_{i,t}^1, \nu_{i,t}^2)$ are independent. We then numerically search for the process $(\Delta\tau_{i,t}^f, \Delta\tau_{i,t}^h, \Delta a_{i,t}, z_{i,t}, \nu_{i,t}^1, \nu_{i,t}^2)$ that minimizes the sum of unconditional variances of the exogenous noise terms $\Delta a_{i,t}$, $z_{i,t}$, $\nu_{i,t}^1$ and $\nu_{i,t}^2$ subject to the implementability constraint (35). The resulting minimal volatilities are reported in Table 4.

In Table 3 in the main text, we also report minimal volatilities for the approximate economy described in Section 6.2. The strategy to derive those volatilities is similar to the

---

[20]To keep the Δ-economy stationary, we assume that the financial distribution shock $\nu_{i,t}^2$ is elastic with respect to expected consumption. Let $\tilde{\nu}_{i,t}^2$ denote the adjusted distribution shock. Then formally we assume that $\tilde{\nu}_{i,t}^2 = \nu_{i,t}^2 + (1-b)\mathbb{E}_t\Delta c_{i,t+1}$, the adjustments in consumption resulting from the second term are made with a delay of 32 quarters in order to avoid interference with the learning problem. Throughout, we set $b = 0.9995$.

[21]Notice that by construction the right-hand side of (35) evaluates to zero for all $j \geq 32$, so that a MA (32) process driving the Δ-economy is sufficient to implement the aggregate wedges.

Table 4: Standard deviations of island-specific noise terms (original economy)

Noise process	$\Delta a_{i,t}$	$z_{i,t}$	$\nu^1_{i,t}$	$\nu^2_{i,t}$
Standard deviation	1.03	4.22	1.52	1.91

one described above, with a slightly modified Δ-economy that for $m \in \{f, h\}$ differentiates between $\Delta\mathring{\tau}^m_{i,t}$ and $\Delta\tau^m_{i,t}$ and adds additional implementability conditions corresponding to $\omega^m_{i,t} \perp \mathring{\tau}_{i,t}$ as explained in the main text (for the purpose of the implementation, the gap between $\Delta\mathring{\tau}^m_{i,t}$ and $\Delta\tau^m_{i,t}$ is exogenously given by the noises in the signals $\chi^m_{i,t}$ as specified by (31)).

B Details of the econometric methodology

B.1 Applying the frequency-filter

Let

$$J = \left(\tilde{\Omega}_T - \tilde{\Omega}(\gamma)\right)' W \left(\tilde{\Omega}_T - \tilde{\Omega}(\gamma)\right) \tag{41}$$

denote the penalty function in terms of BK-filtered moments, where as suggested by Gorodnichenko and Ng (2009) the filter is applied to both the data and the model. In this appendix, we demonstrate how the penalty can be expressed in terms of the variance over unfiltered moments, $\Omega \equiv vech\left\{\text{Var}\left(\mathrm{d}x^t_{t-K}\right)\right\}$, where d is the first-difference operator, and $K \equiv k + 2\bar{\tau}$ with $\bar{\tau}$ denoting the approximation horizon of the BK-filter.[22] Specifically, for any positive-semidefinite W we show that J in (41) is equivalent to

$$J = \left(\Omega_T - \Omega(\gamma)\right)' \tilde{W} \left(\Omega_T - \Omega(\gamma)\right), \tag{42}$$

with $\tilde{W} = \Xi'W\Xi$ replacing W.

The Baxter and King (1999) filtered version of x_t is given by

$$\tilde{x}_t = \sum_{j=-\bar{\tau}}^{\bar{\tau}} a_j x_{t-j}$$

where $\tilde{x}_t$ is stationary by construction. For the high-pass filter used in this paper, the weights $\{a_j\}$ are given by

$$a_j = \tilde{a}_j - \theta, \quad \theta = \frac{1}{2\bar{\tau}+1} \sum_{j=-\bar{\tau}}^{\bar{\tau}} \tilde{a}_j$$

---

[22]The first-difference filter is applied to ensure stationarity for variables that have a unit root.

with

$$\tilde{a}_0 = 1 - \bar{\omega}/\pi, \quad \tilde{\alpha}_{j \neq 0} = -\sin(j\bar{\omega})/(j\pi), \quad \bar{\omega} = 2\pi/32.$$

To construct the filter-matrix Ξ, rewrite $\tilde{x}_t$ in terms of growth rates to get

$$\tilde{x}_t = \sum_{j=-\bar{\tau}}^{\bar{\tau}} \sum_{l=0}^{\infty} a_j \mathrm{d}x_{t-j-l}.$$

Noting that $\sum_{j=-\bar{\tau}}^{\bar{\tau}} a_j = 0$, we can simplify to get

$$\tilde{x}_t = B \mathrm{d}x_{t-\bar{\tau}-j}^{t+\bar{\tau}}$$

where

$$B = [b_{-\bar{\tau}}, \ldots, b_{\bar{\tau}}] \otimes I_n, \tag{43}$$

$n = 4$ is the number of variables in $\tilde{x}_t$, and $b_s = \sum_{j=-\bar{\tau}}^{s} \alpha_j$.

Letting L_j define the backshift matrix

$$L_j = \left[0_{n(2\bar{\tau}+1),nj}, \quad I_{n(2\bar{\tau}+1)}, \quad 0_{n(2\bar{\tau}+1),n(k-j)} \right], \tag{44}$$

we then have that

$$\tilde{\Sigma}_j \equiv \mathrm{Cov}(\tilde{x}, \tilde{x}_{t-j}) = B L_0 \Sigma^K L_j' B',$$

or, equivalently,

$$vec(\tilde{\Sigma}_j) = (BL_j \otimes BL_0) \, vec(\Sigma^K).$$

To complete the construction of Ξ, define selector-matrices P_0 and P_1 such that

$$vech(\tilde{\Sigma}^k) = P_0 \begin{bmatrix} vec(\tilde{\Sigma}_0) \\ \vdots \\ vec(\tilde{\Sigma}_k) \end{bmatrix}$$

and

$$vec(\Sigma^K) = P_1 vech(\Sigma^K).$$

Stacking up $vec(\tilde{\Sigma}_j)$, we then get

$$\tilde{\Omega} = \Xi \Omega$$

where

$$\Xi = P_0 \begin{bmatrix} BL_0 \otimes BL_0 \\ \vdots \\ BL_k \otimes BL_0 \end{bmatrix} P_1 \tag{45}$$

with B and L_j as in (43) and (44). Substitution in (41) yields (42).

B.2 Estimation of the optimal weighting matrix

Our estimation of

$$S \equiv \mathrm{Var}\left\{ T^{1/2}\left(\tilde{\Omega}_T - \tilde{\Omega}(\gamma_0) \right) \right\} = \mathrm{Var}\left\{ T^{1/2}\tilde{\Omega}_T \right\}$$

is based on a bootstrap identical to the one described in the main text (with 5000 replications). Let $\tilde{S} = \mathrm{Var}\{T^{1/2}\tilde{\Omega}_b\}$ where the variance is across bootstrap samples with $\tilde{\Omega}_b = \Xi\Omega_b$ being the target moments in a given sample $b \in \{1, \dots, 5000\}$. It is well-known that estimations of covariances of covariance structures are prone to small-sample bias due to the estimation of fourth moments, which tend to correlate with the targeted covariance structure (e.g., Abowd and Card 1989 and Altonji and Segal, 1996). In addition, we find that $\tilde{S}$ is near singular. We follow Christiano, Trabandt and Walentin (2010) and dampen the off-diagonal elements of $\tilde{S}$ relative to the diagonal to improve the small-sample efficiency of $\tilde{S}$. Specifically, let $\tilde{S}_{i,j}$ denote the (i,j)-th block of $\tilde{S}$ corresponding to the cross-sample covariance between $\mathrm{Cov}\{dx_t, dx_{t-i}\}$ and $\mathrm{Cov}\{dx_t, dx_{t-j}\}$. We replace $\tilde{S}$ by $\tilde{S}^{(\nu_1,\nu_2)}$ where each block $\tilde{S}_{i,j}^{(\nu_1,\nu_2)}$ in $\tilde{S}^{(\nu_1,\nu_2)}$ is given by $\varsigma_{i,j}(\nu_1)(M(\nu_2) \circ \tilde{S}_{i,j})$ with

$$\varsigma_{i,j} = \left(1 - \frac{|i-j|}{k+1}\right)^{\nu_1}, \quad \nu_1 \geq 0$$

and

$$M(\nu_2) = 1 - \nu_2 + \nu_2 I_{n^2}, \quad 0 \leq \nu_2 \leq 1,$$

where $\circ$ is the element-wise (Hadamard) product and $n = 4$ is the number of elements in x_t. Intuitively, ν_2 is a dampening factor applied to the off-diagonal elements within each $\tilde{S}_{i,j}$ block and ν_1 is a dampening applied to the covariance between different auto-covariance-blocks that is increasing in $|i-j|$. For $\nu_1 = \nu_2 = 0$, the resulting matrix $\tilde{S}^{(\nu_1,\nu_2)}$ equals $\tilde{S}$. For $\nu_1 \to \infty$, $\tilde{S}^{(\nu_1,\nu_2)}$ becomes a block-diagonal version of $\tilde{S}$, so that the GMM criterion does not depend on the cross-block co-variation $\mathrm{Cov}\left\{\mathrm{Cov}\left[\tilde{x}_t, \tilde{x}_{t-i}\right], \mathrm{Cov}\left[\tilde{x}_t, \tilde{x}_{t-j}\right]\right\}$ for $i \neq j$. For $\nu_2 = 1$, each block $\tilde{S}_{i,j}^{(\nu_1,\nu_2)}$ becomes diagonal, so that the GMM criterion does not depend on the cross-variable co-variation $\mathrm{Var}\left\{\mathrm{Cov}\left[\tilde{x}_t^m, \tilde{x}_{t-i}^n\right]\right\}$ for any $m \neq n$. In either case, the criterion

continues to make full use of all targeted moments $\tilde{\Omega}_T$. To have a consistent estimator of S, we need that $\nu_1 \to 0$ and $\nu_2 \to 0$ as $T \to \infty$, but do not restrict the small sample behavior of ν_1 and ν_2. Our choice of ν_1 and ν_2 is aimed at maximizing the small sample efficiency of $\tilde{S}^{(\nu_1,\nu_2)}$. Specifically, we set ν_1 and ν_2 to minimize the RMSE in a simulation experiment where we generate time series of the length of our original data sample, treating the bootstrap DGP described in the main text as the truth. The efficient estimator is given by $\nu_1 = 5$ and $\nu_2 = 0.5$, which also suffices to make $\tilde{S}^{(\nu_1,\nu_2)}$ well-conditioned. Collecting, we have $W = [\tilde{S}^{(\nu_1,\nu_2)}]^{-1}$ and $\tilde{W} = \Xi'[\tilde{S}^{(\nu_1,\nu_2)}]^{-1}\Xi$.

C Additional tables and figures

Table 5: Parameters of the estimated VAR(1) process for the information wedges

Coefficient matrix Λ on lagged states			Coefficient matrix R, where $\tilde{\Psi} = RR'$		
0.874	−1.171	−1.103	2.049	0	1.325
(0.730, 0.896)	(−2.406, 1.538)	(−2.406, 1.538)	(0.997, 2.442)		(0.506, 1.882)
0.006	0.803	0.222	−0.165	−0.218	0.006
(−0.005, 0.017)	(0.436, 0.867)	(−0.004, 0.338)	(−0.285, −0.020)	(−0.252, −0.123)	(−0.086, 0.079)
0.000	0.000	0.271	0.025	0.406	0.062
(−0.001, 0.001)	(−0.172, 0.023)	(0.039, 0.616)	(−0.069, 0.160)	(0.306, 0.447)	(−0.082, 0.146)
			0	0	−0.667
					(−0.928, −0.395)

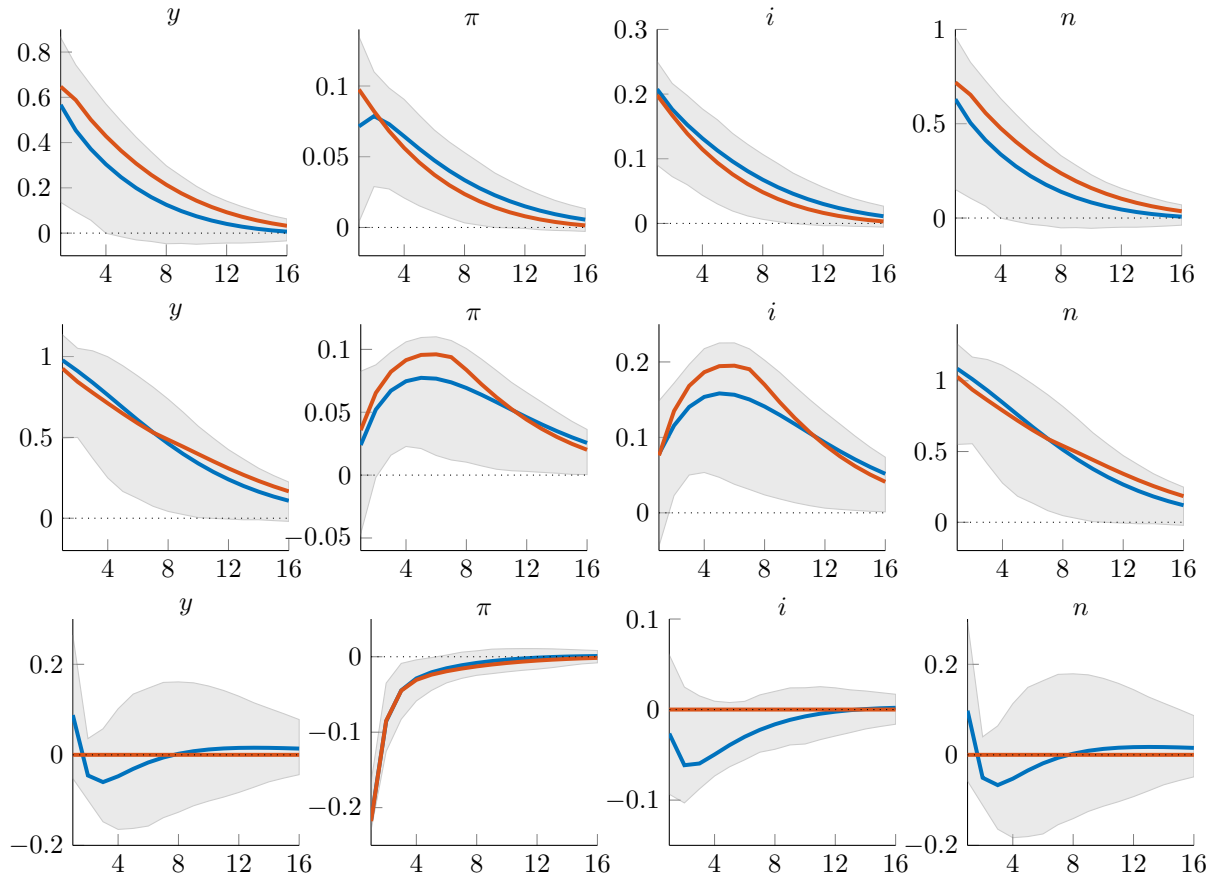Note.—Numbers in parenthesis are 90% confidence intervals.

Figure 8: Impulse responses in the estimated economy (blue) and the implemented learning economy (red). The first row depicts impulse responses to ϵ_t, the second row, responses to $\eta_{1,t}$, and the third row, responses to $\eta_{2,t}$.

References

Abowd, John M, and David Card. 1989. "On the Covariance Structure of Earnings and Hours Changes." *Econometrica*, 57(2): 411–445.

Acharya, Sushant. 2013. "Dispersed Beliefs and Aggregate Demand Management." *working paper.*

Altonji, Joseph G, and Lewis M Segal. 1996. "Small-sample bias in GMM estimation of covariance structures." *Journal of Business & Economic Statistics*, 14(3): 353–366.

Angeletos, George-Marios, and Jennifer La'O. 2013. "Sentiments." *Econometrica*, 81(2): 739–779.

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas. 2015. "Quantifying Confidence." *working paper.*

Basu, Susanto, John G Fernald, and Miles S Kimball. 2006. "Are Technology Improvements Contractionary?" *American Economic Review*, 96(5): 1418–1448.

Baxter, Marianne, and Robert G King. 1999. "Measuring business cycles: Approximate band-pass filters for economic time series." *Review of Economics and Statistics*, 81(4): 575–593.

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.

Bergemann, Dirk, and Stephen Morris. 2016. "Bayes correlated equilibrium and the comparison of information structures in games." *Theoretical Economics*, 11(2): 487–522.

Bergemann, Dirk, Tibor Heumann, and Stephen Morris. 2014. "Information and Volatility." *mimeo.*

Blanchard, Olivier J, Jean-Paul L'Huillier, and Guido Lorenzoni. 2013. "News, Noise, and Fluctuations: An Empirical Exploration." *American Economic Review*, 103(7): 3045–3070.

Calvo, Guillermo. 1983. "Staggered Prices in a Utility Maximizing Framework." *Journal of Monetary Economics*, 12(1978): 383–398.

Carroll, Christopher D. 2003. "Macroeconomic Expectations of Households and Professional Forecasters." *Quarterly Journal of Economics*, 118(1): 269–298.

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

Christiano, Lawrence J, Martin Eichenbaum, and Robert Vigfusson. 2003. "What happens after a technology shock?" National Bureau of Economic Research.

Christiano, Lawrence J., Mathias Trabandt, and Karl Walentin. 2010. "DSGE Models for Monetary Policy Analysis." *Handbook of Monetary Economics, Edited by Benjamin Friedman and Michael Woodford*, 3: 285–367.

Cochrane, John H. 1994. "Shocks." Vol. 41, 295–364, Elsevier.

Fernald, John. 2014. "A quarterly, utilization-adjusted series on total factor productivity." *Manuscript, Federal Reserve Bank of San Francisco*, , (April).

Gali, Jordi. 1999. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review*, 89(1): 249–271.

Gorodnichenko, Yuriy, and Serena Ng. 2009. "Estimation of DSGE Models when the Data are Persistent." *working paper*.

Hellwig, Martin F. 1980. "On the aggregation of information in competitive markets." *Journal of Economic Theory*, 22(3): 477–498.

Huo, Zhen, and Naoki Takayama. 2015. "Higher Order Beliefs, Confidence, and Business Cycles." *working paper*.

Jurado, Kyle. 2016. "Advance Information and Distorted Beliefs in Macroeconomic and Financial Fluctuations." Working Paper.

King, Robert G, Charles I Plosser, James H Stock, and Mark W Watson. 1991. "Stochastic Trends and Economic Fluctuations." *The American Economic Review*, 819–840.

Lorenzoni, Guido. 2009. "A Theory of Demand Shocks." *The American Economic Review*, 99(5): 2050–2084.

Maćkowiak, Bartosz, and Mirko Wiederholt. 2015. "Business cycle dynamics under rational inattention." *Review of Economic Studies*, forthcoming.

Melosi, Leonardo. 2014. "Estimating models with dispersed information." *American Economic Journal: Macroeconomics*, 6(1 C): 1–31.

Nimark, Kristoffer P. 2009. "Dynamic higher order expectations." *mimeo.*

Rondina, Giacomo, and Todd B Walker. 2014. "Dispersed Information and Confounding Dynamics." *mimeo.*

Shapiro, Matthew, and Mark Watson. 1988. "Sources of business cycles fluctuations." In *NBER Macroeconomics Annual 1988, Volume 3.* 111–156. MIT Press.

Smets, Frank, and Rafael Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review*, 97(3): 586–606.

Uhlig, Harald. 2003. "What moves real GNP?" *Unpublished manuscript.*

Uhlig, Harald. 2004. "Do Technology Shocks Lead to a Fall in Total Hours Worked?" *Journal of the European Economic Association*, 2(2-3): 361–371.

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.* , ed. Philippe Aghion, Roman Frydman, Joseph Stiglitz and Michael Woodford, 25–58. Princeton:Princeton University Press.

Woodford, Michael. 2011. *Interest and prices: foundations of a theory of monetary policy.* Princeton University Press.