QUANTIFYING CONFIDENCE

George-Marios Angeletos
Fabrice Collard
Harris Dellas

Quantifying Confidence
George-Marios Angeletos, Fabrice Collard, and Harris Dellas
NBER Working Paper No. 20807
December 2014
JEL No. E0

ABSTRACT

We develop a tractable method for augmenting macroeconomic models with autonomous variation in higher-order beliefs. We use this to accommodate a certain type of waves of optimism and pessimism that can be interpreted as the product of frictional coordination and, unlike the one featured in the news literature, regards the short-term economic outlook rather than the medium- to long-run prospects. We show that this enrichment provides a parsimonious explanation of salient features of the data; it accounts for a significant fraction of the business-cycle volatility in estimated models that allow for various competing structural shocks; and it captures a type of fluctuations that have a Keynesian flavor but do not rely on nominal rigidities.

George-Marios Angeletos
Department of Economics, E52-530
MIT
77 Massachusetts Avenue
Cambridge, MA 02139
and NBER
angelet@mit.edu

Fabrice Collard
University of Bern
Departement Volkswirtschaftslehre
Schanzeneckstrasse 1
CH-3001 Bern, Switzerland
fabrice.collard@vwi.unibe.ch

Harris Dellas
Department of Economics
University of Bern
VWI, Schanzeneckstrasse 1
CH 3012 Bern
Switzerland
harris.dellas@vwi.unibe.ch

An online appendix is available at http://www.nber.org/data-appendix/w20807

1 Introduction

As a recession sets in, confidence in the prospects of the economy sinks. Firms cut down on employment and investment as they turn pessimistic about the demand for their products; consumers reduce spending as they turn pessimistic about their job and income prospects; and the pessimism of one economic agent appears to justify, if not feed, that of others.

Workhorse macroeconomic models, especially those used for quantitative purposes, interpret such phenomena as the coordinated response of the agents to changes in payoff-relevant fundamentals such as the general level of know-how (technology shocks) or the efficacy of the financial sector (financial shocks). This leaves little room for expectations to play an autonomous role in driving the business cycle. This in turn is because such models assume away, not only multiple equilibria, but also frictional coordination in the form of higher-order uncertainty.[1] Formally, the economy is modeled as a game in which all players share a common prior and the same information at all times, face no uncertainty about one another's beliefs and behavior conditional on the fundamentals, and reach a perfect consensus about the current state and the future prospects of the economy.

These are strong assumptions, which are at odds with the heterogeneity of expectations evident in surveys. Once these assumptions are relaxed, the expectations of economic outcomes—for instance, firms' expectations of consumer spending and consumers' expectations of employment and income—can diverge from the expectations of fundamentals. This provides a novel explanation of the discrepancies between the predictions of the baseline RBC model and the data. It also accommodates phenomena akin to self-fulfilling fluctuations despite the uniqueness of equilibrium. In this paper, we provide a tractable formalization of these ideas and explore their quantitative potential.

Two contributions. We make two contributions, one methodological and one applied. We first develop a general method for enriching dynamic, general-equilibrium models with a tractable form of aggregate variation in higher-order beliefs (i.e., the beliefs of the beliefs of others). We then use this method to explore the macroeconomic implications of a certain type of waves of optimism and pessimism that can be interpreted as the product of frictional coordination and—unlike the one captured by the literature on news shocks—regards the short-term prospects of the economy.

We refer to these waves as variation in "confidence" and explore their implications within RBC and New Keynesian models of either the textbook or the medium-scale DSGE variety.[2] We show that they offer a parsimonious yet potent explanation of the business-cycle data. We also argue that they help capture a form of "demand-driven fluctuations" that does not rely on nominal rigidities and does not have to manifest as comovement between inflation and real economic activity.

Background and methodological contribution. We build heavily upon the macroeconomic literature on incomplete information and higher-order uncertainty. This literature goes back at least to Phelps (1971) and Townsend (1983) and has been revived recently by the influential contribu-

---

[1] Higher-order uncertainty refers to the uncertainty that the agents face about the beliefs of others.
[2] RBC is acronym for Real Business Cycles, DSGE for Dynamic Stochastic General Equilibrium.

tions of Morris and Shin (2001) and Woodford (2002). Within this literature, the closest precursor to our paper is Angeletos and La'O (2013), which has shown how higher-order uncertainty can help unique-equilibrium models accommodate forces akin to animal spirits and coordination failures.[3]

We borrow from this literature the insight that higher-order beliefs can deviate from first-order beliefs, but use heterogeneous priors instead of complex learning dynamics to engineer fluctuations in the gap between first- and higher-order beliefs. This approach entails a certain departure from Rational Expectations. But it also allows us to bypass the computational complications that have hindered progress in this literature on the quantitative front,[4] and to develop a general method for augmenting macroeconomic models with rich, yet tractable, higher-order beliefs.

In order to illustrate this point, consider the baseline RBC model. The equilibrium dynamics of this model can be summarized by a policy rule of the form $X_t = G(K_t, A_t)$, where A_t is the technology shock, K_t is the capital stock, and $X_t = (Y_t, N_t, C_t, K_{t+1})$ is a vector that collects the relevant macroeconomic outcomes, namely output, employment, consumption, and investment or, equivalently, the next-period capital stock. Adding incomplete information to this model allows higher-order beliefs to diverge from first-order beliefs but also increases the model's state space and considerably complicates its solution. By contrast, our heterogeneous-prior formulation captures a similar type of beliefs-driven fluctuations with only a minimal change in the state space: the equilibrium policy rule takes the form $X_t = G(K_t, A_t, \xi_t)$, where ξ_t is an exogenous random variable which, by construction, encapsulates the deviation of higher-order beliefs from first-order beliefs.

This gain in tractability is not limited to the baseline RBC model. For a large, essentially arbitrary, class of linear DSGE models, our approach guarantees a minimal increase in the state space and delivers the solution of the beliefs-augmented model as a relatively simple transformation of the solution of the original model. The beliefs-augmented model can thus be simulated, calibrated, and estimated with essentially the same facility as the original one.[5]

Applied contribution. By construction, the ξ_t shock represents variation in the gap between the first- and the higher-order beliefs of the exogenous fundamental (TFP). In equilibrium, this translates into waves of optimism and pessimism about aggregate output, employment, spending, and so on. We refer to these waves as variation in "confidence" and to ξ_t as the "confidence shock."

---

[3]See Angeletos and Lian (2016a) for a survey and evaluation of this literature. Let us also emphasize that we have in mind situations in higher-order uncertainty is both present and of consequence. This differentiates our paper, and the aforementioned literature more broadly, from both Lucas (1972), who considers a setting in which higher-order uncertainty is present but inconsequential, and Sims (2003), who abstracts from whether and how rational inattention can be conductive to higher-order uncertainty.

[4]These complications were first highlighted by Townsend (1983). They include the need for large state spaces in order to keep track of the dynamics of higher-order beliefs and the fixed point between the law of motion of the state and the agents' filtering problem. For detailed expositions of these complications and complementary attempts to make progress on the quantitative front, see Nimark (2017) and Huo and Takayama (2015a,b).

[5]The aforementioned gain may carry a cost: we abstract from the restrictions that the common-prior assumption, together with appropriate evidence, may impose on the magnitude and persistence of higher-order uncertainty. We elucidate this issue in Subsection 3.3 and argue that it may not matter for the applied contribution of our paper.

A distinct attribute of these waves is that they regard the short-term economic outlook. For instance, a negative innovation in ξ_t causes the firms to become pessimistic about profitability and returns over the next few quarters, and the consumers to become pessimistic about wages and income over the same horizon, without any change in expectations of either the exogenous fundamentals at any horizon or the endogenous outcomes in the medium to long run.

This property underlies our preferred interpretation of the ξ_t shock as a vehicle for autonomous variation in expectations about the short-term economic outlook. It also distinguishes our contribution from the literature on news and noise shocks (Jaimovich and Rebelo, 2009; Lorenzoni, 2009; Barsky and Sims, 2011). That literature stresses beliefs of productivity and income in the medium to long run, a feature that, in the absence of appropriate bells and whistles, cannot generate realistic business cycles. By contrast, the emphasis on expectations about the short run allows our mechanism to produce realistic business cycles even within the textbook RBC model.

To understand why, augment the RBC model with our mechanism and consider a negative innovation in ξ_t. As firms expect the demand for their products to be weak in the short run, they find it optimal to lower their demand for labor and capital. In the eyes of households, this translates into a transitory fall in wages, capital returns, and overall income. Because this entails relatively weak wealth effects and relatively strong substitution effects, households react by working less and by reducing both consumption and saving. Variation in "confidence" thus generates strong positive comovement between employment, output, consumption, and investment at the business-cycle frequency, without commensurate movements in labor productivity and TFP at any frequency.

These predictions are in line with the comovements observed in the US data and cannot be easily replicated by alternative theories. We provide support for these claims by carrying out two empirical exercises. In the first, we consider the conditional moments in the data after removing the effects of an empirical proxy of the technology shock. One can think of the filtered data as representing the "residuals" between the data and the predictions of the baseline RBC model. Our theory does well on this front: not only does it capture the comovements in these residuals, but it also outperforms other parsimonious extensions of the RBC model. A similar picture emerges when considering the wedges along the lines of Chari, Kehoe, and McGrattan (2007).

In the second exercise, we estimate medium-scale DSGE models that include our confidence shock alongside several other shocks and also contain familiar bells and whistles from the DSGE literature, such as the specific types of consumption and investment adjustment costs popularized by Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2007). This exercise lacks parsimony—in particular, it allows business-cycle comovements to be accounted for by the combination of a plethora of shocks—but corresponds closer to standard practice. Despite the presence of multiple, competing shocks, the confidence shock emerges as the main driver of the business cycle, accounting for about one half of the volatility in the key macroeconomic quantities (GDP, hours, investment, consumption) and for the bulk of their comovements.

This finding is robust across two specifications. The first includes sticky prices, lets monetary policy follow a realistic Taylor rule, and is estimated using both real and nominal variables. The

3

second assumes flexible prices, abstracts from monetary policy and inflation, and is estimated using only real quantities. Irrespective of the specification, the posterior odds of the model that excludes the confidence shock are considerably smaller than those of the model that contains it. Last but not least, our mechanism allows for fluctuations that resemble those produced by aggregate demand shocks but do not require commensurate movements in inflation, a feature that seems consistent with the data and helps bypass the empirical failures of old and new Philips curves.

Because a direct, empirical counterpart to the confidence shock is hard, if possible at all, to obtain,[6] these findings only provide indirect support for our theory. They nevertheless indicate the quantitative potential of three elements that are missing from the DSGE literature: frictional coordination in the form of higher-order uncertainty; a prominent role for waves of optimism and pessimism about the short-term economic outlook; and demand-driven fluctuations outside the inflation-output nexus of the New Keynesian framework. Our contribution combines all three of these elements. Future work may narrow the focus to one or another of these elements.

Layout. The rest of the paper is organized as follows. Section 2 sets up the baseline model. Section 3 explains the recursive formulation of the equilibrium and our solution method. Section 4 derives, evaluates, and interprets the empirical properties of the baseline model. Section 5 extends the analysis to two richer, estimated, models. Section 6 concludes.

2   An RBC Prototype with Tractable Higher-Order Beliefs

In this section we set up our baseline model: an RBC prototype, augmented with a tractable form of higher-order belief dynamics. We first describe the physical environment, which is quite standard. We then specify the structure of beliefs, which constitutes the main novelty of our approach.

Geography, markets, and timing. There is a continuum of islands, indexed by i, and a mainland. Each island is inhabited by a firm and a household, which interact in local labor and capital markets. The firm uses the labor and capital provided by the household to produce a differentiated intermediate good. A centralized market for these goods operates in the mainland, alongside a market for a final good. The latter is produced with the use of the intermediate goods and is itself used for consumption and investment. All markets are competitive.

Time is discrete, indexed by $t \in \{0, 1, \ldots\}$, and each period contains two stages. The labor and capital markets of each island operate in stage 1. At this point, the firm decides how much labor and capital to demand—and, symmetrically, the household decides how much of these inputs to supply—on the basis of incomplete information regarding the concurrent level of economic activity

---

[6]A well-known empirical measure of expectations is the University of Michigan Index of Consumer Sentiment. This index comoves with, and in fact leads, the business cycle. Furthermore, this index is uncorrelated with utilization-adjusted TFP at all leads and lags. While these facts are in line with our theory, they do not rule out the possibility that the comovement of that index with the business cycle is driven by some other fundamental. The inherent difficulty is that the definition of what is a fundamental and what is not depends on the model under consideration.

on other islands. In stage 2, the centralized markets for the intermediate and the final goods operate, the actual level of economic activity is publicly revealed, and the households make their consumption and saving decisions on the basis of this information.

Households. Consider the household on island i. Her preferences are given by

$$\sum_{t=0}^{\infty} \beta^t U(c_{it}, n_{it})$$

where $\beta \in (0,1)$ is the discount factor, c_{it} is consumption, n_{it} is employment (hours worked), and U is the per-period utility function. The latter takes the form $U(c,n) = \frac{c^{1-\gamma}-1}{1-\gamma} - \frac{n^{1+\nu}}{1+\nu}$ where $\gamma \geq 0$ is the inverse of the elasticity of intertemporal substitution and $\nu \geq 0$ is the inverse of the Frisch elasticity of labor supply. Balanced growth requires $\gamma = 1$, a restriction that we impose in our quantitative exercises; letting $\gamma \neq 1$ helps accommodate a useful example in Section 3. The household's budget constraint is $P_t c_{it} + P_t i_{it} = w_{it} n_{it} + r_{it} k_{it} + \pi_{it}$, where P_t is the price of the final good, i_{it} is investment, w_{it} is the local wage, r_{it} is the local rent on capital, and π_{it} is the profit of the local firm. Finally, the law of motion for capital is $k_{i,t+1} = (1-\delta)k_{it} + i_{it}$, where $\delta \in (0,1)$ is the depreciation rate.

Intermediate-good producers. The output of the firm on island i is given by

$$y_{it} = A_t n_{it}^{1-\alpha} k_{it}^{\alpha}$$

where A_t is the aggregate TFP level and k_{it} is the local capital stock. The firm's profit is $\pi_{it} = p_{it} y_{it} - w_{it} n_{it} - r_{it} k_{it}$. For future reference, note that variation in expectations of p_{it} translates in variation in expectations of the returns to capital and labor.

Final-good sector. The final good is produced with a Cobb-Douglas technology, so that $\log Y_t = \int_0^1 \log y_{it} \, di$. By implication, the demand for the good of island i satisfies

$$\frac{p_{it}}{P_t} = \frac{Y_t}{y_{it}}. \tag{1}$$

Without any loss, we henceforth normalize the price level so that $P_t = 1$.[7]

The technology shock. TFP follows a random walk: $\log A_t = \log A_{t-1} + v_t$, where v_t is the period t innovation. The latter is drawn from a Normal distribution with mean 0 and variance σ_a^2.

A tractable form of higher-order uncertainty. We open the door to a gap between first- and higher-order beliefs by removing common knowledge of A_t in stage 1 of period t: each island i observes only a private signal of the form $z_{it} = \log A_t + \varepsilon_{it}$, where ε_{it} is an island-specific error. We then engineer the desired variation in higher-order beliefs by departing from the common-prior assumption and letting each island believe that the signals of others are biased: for every i, the prior of island i is that $\varepsilon_{it} \sim \mathcal{N}(0, \sigma^2)$ and that $\varepsilon_{jt} \sim \mathcal{N}(\xi_t, \sigma^2)$ for all $j \neq i$, where ξ_t is a random variable

---

[7]This only applies to the present model, which abstract from nominal rigidity and monetary policy. In the New Keynesian variant studied in Section 5, P_t is determined jointly with the real allocations.

that becomes commonly known in stage 1 of period t and that represents the perceived bias in one another's signals. These priors are commonly known: the agents "agree to disagree".

We have in mind a sequence of models in which first- and higher-order beliefs converge to Dirac measures as $\sigma \to 0$. But instead of studying the case with $\sigma \approx 0$, we only study the case with $\sigma = 0$. This guarantees that the agents act as if they were perfectly informed about the underlying state of Nature and that the pair (A_t, ξ_t) is a sufficient statistic for the entire hierarchy of beliefs about both current and future fundamentals. Together with the assumption that the aggregate capital stock (the endogenous state variable) becomes common knowledge at the end of each period, this guarantees that the model admits a tractable recursive solution, as shown in Section 3.

The confidence shock. We finally let ξ_t follow an $AR(1)$ process: $\xi_t = \rho \xi_{t-1} + \zeta_t$, where $\rho \in [0, 1)$ and ζ_t is drawn from a Normal distribution with mean 0 and variance σ_ξ^2. This helps mimic an elementary property of common-prior settings: in such settings, any innovation in the gap between first- and higher-order beliefs can last for a while but must eventually vanish as old information gets replaced by new. See Subsection 3.3 for an example that illustrates this point.

Remarks and Interpretation. Our heterogeneous-prior specification puts strains on the rationality of the agents. First, it lets the impact of ξ_t on n-th order beliefs increase with n. Second, it ties the persistence of higher-order beliefs to the persistence of the ξ_t shock. Finally, it implies a systematic bias in equilibrium expectations: although the firms and the consumers predict correctly the sign of the equilibrium impact of ξ_t on the relevant economic outcomes, they systematically overestimate its magnitude, and they also fail to learn from their past mistakes.

One does not have to take these properties literally. Common-prior settings such as those studied in Angeletos and La'O (2013), Benhabib, Wang, and Wen (2015), Huo and Takayama (2015a), Nimark (2017), and Rondina and Walker (2014) can accommodate similar fluctuations in higher-order beliefs. In effect, what is "bias" in our setting becomes "rational confusion" in those settings. Furthermore, higher-order beliefs can be persistent in both cases, although the persistence is endogenous to the learning that takes place over time in the latter case. We illustrate these points in Subsection 3.3 by establishing an observational equivalence, from the point of view of aggregate data, between a special case of our model and a common-prior variant.

Most importantly, the subsequent analysis will reveal that the empirical performance of our theory hinges, not on the precise micro-foundations of the belief waves considered, but rather on the property that the beliefs regard firm profitability and household income in the short run as opposed to the medium and long run. We thus encourage the reader to adopt a flexible interpretation of ξ_t as a modeling device that helps capture more generally this kind of belief waves.[8]

---

[8]We can imagine at least two variants of our framework that can also capture such waves. The one replaces our heterogeneous-prior specification with Knightean uncertainty (ambiguity) about the information of others and allows for, possibly endogenous, variation in the level of this uncertainty; such a model could build a bridge between our work and a literature on ambiguity and robust control (Hansen and Sargent, 2007, 2012; Woodford, 2010). The other allows directly for irrational shifts in expectations of profitability and income in the short run; such a model would fit well with the narratives in Akerlof and Shiller (2009) and Burnside, Eichenbaum, and Rebelo (2016).

3 Equilibrium characterization and solution method

In this section, we characterize the equilibrium of the model and present our solution method. We also use an example to illustrate the idea that our heterogeneous-prior formulation can be seen as convenient proxy for belief fluctuations in common-prior settings with rich information structures.

3.1 Recursive equilibrium

As behavior is forward looking, the optimal choices any agent (or island) makes at any point of time depend on her beliefs, not only of the concurrent behavior of others, but also of their behavior in the future. This suggests a high-dimensional fixed-point relation between actual behavior and the expectations that agents form at any time about future economic outcomes, including expectations of the future terms of trade (the prices of the island-specific goods), wages, and interest rates. In general, the introduction of higher-order uncertainty can perturb this kind of expectations in a sufficiently significant manner as to render a low-dimensional recursive representation infeasible. With our formulation, however, such a representation is feasible and, indeed, relatively straightforward.

To start with, note that the equilibrium allocations on any given island can be obtained by solving the problem of a fictitious local planner. The latter chooses local employment, output, consumption and savings so as to maximize local welfare subject to the following resource constraint:

$$c_{it} + k_{i,t+1} \quad = \quad (1-\delta)k_{it} + p_{it}y_{it} \tag{2}$$

Note that this constraint depends on p_{it} and, thereby, on aggregate output, objects that are endogenous in general equilibrium but are taken as given by the fictitious local planner (or, equivalently, in the partial equilibrium of the given island). This dependence captures the type of aggregate-demand externalities and other general-equilibrium effects that are at the core of DSGE models.

To make her optimal decisions at any given point of time, the aforementioned planner must form beliefs about the value of p_{it} (or, equivalently, of Y_t) at all future points of time. These beliefs encapsulate the beliefs that the local firm forms about the evolution of the demand for its product and of the costs of its inputs, as well as the beliefs that the local consumer forms about the dynamics of local income, wages, and capital returns. The fact that the various beliefs are tied together underscores the cross-equations restrictions that discipline the exercises conducted in this paper: if expectations were "completely" irrational, the beliefs of different endogenous objects would not be tied together. The observable implications of these restrictions will be revealed in what follows. For now, we emphasize that ξ_t matters for equilibrium outcomes because, and only because, it triggers comovement in the expectations of the various actors in our model.

In a recursive equilibrium, these expectations can be tracked with the help of a small number of functions, which themselves encapsulate the fixed-point relation between behavior and beliefs. For the model under consideration, this means that we can define a recursive equilibrium as a collection of four functions, denoted by **P**, **G**, V_1, and V_2, such that the following is true:

- $\mathbf{P}(z, \xi, K)$ captures the price (the terms of trade, or equivalently, the demand) expected by an island in stage 1 of any given period when the local signal is z, the confidence shock is ξ, and the capital stock is K; and $\mathbf{G}(A, \xi, K)$ gives the aggregate capital stock next period when the current realized value of the aggregate state is (A, ξ, K).

- V_1 and V_2 solve the following Bellman equations:

$$
\left.
\begin{aligned}
V_1(k; z, \xi, K) = \max_{n} \quad & V_2(\hat{m}; z, \xi, K) - \tfrac{1}{1+\nu} n^{1+\nu} \\
s.t. \quad & \hat{m} = \hat{p}\hat{y} + (1 - \delta)k \\
& \hat{y} = z k^{\alpha} n^{1-\alpha} \\
& \hat{p} = \mathbf{P}(z, \xi, K)
\end{aligned}
\right\} \tag{3}
$$

$$
\left.
\begin{aligned}
V_2(m; A, \xi, K) = \max_{c, k'} \quad & \tfrac{c^{1-\gamma}-1}{1-\gamma} + \beta \int V_1(k'; A', \xi', K') \mathrm{d}f(A', \xi'|A, \xi) \\
s.t. \quad & c + k' = m \\
& K' = \mathbf{G}(A, \xi, K)
\end{aligned}
\right\} \tag{4}
$$

- $\mathbf{P}$ and $\mathbf{G}$ are consistent with the policy rules that solve the local planning problem in (3)-(4).

To interpret (3) and (4), note that V_1 and V_2 denote the local planner's value functions in, respectively, stages 1 and 2; m denotes the quantity of the final good acquired in stage 2; and the *hat* symbol over a variable indicates the stage-1 belief of that variable. Next, note that the last constraint in (3) embeds the belief that the price of the local good is governed by the function $\mathbf{P}$, while the other two constraints are the local production function and the local resource constraint. The problem in (3) therefore describes the optimal employment and output choices in stage 1, when the local capital stock is k, the local signal of the aggregate state is (z, ξ, K), and the local beliefs of "aggregate demand" are captured by the function $\mathbf{P}$. The problem in (4), in turn, describes the optimal consumption and saving decisions in stage 2, when the available quantity of the final good is m, the realized aggregate state is (A, ξ, K), and the island expects aggregate capital to follow the policy rule $\mathbf{G}$.

The decision problem of the local planner treats the functions $\mathbf{P}$ and $\mathbf{G}$ as exogenous. In equilibrium, however, these functions must be consistent with the policy rules that solve this problem. Let $\mathbf{n}(k, z; \xi, K)$ be the optimal choice for employment that obtains from (3) and $\mathbf{g}(m; A, \xi, K)$ be the optimal policy rule for capital that obtains from (4). Next, let $\mathbf{y}(z; A, \xi, K) \equiv A\mathbf{n}(z, \xi, K)^{1-\alpha} K^{\alpha}$ be the output level that results from the aforementioned employment strategy where the realized TFP is A and the local capital stock coincides with the aggregate one. The relevant equilibrium-consistency conditions can then be expressed as follows:

$$
\mathbf{P}(z, \xi, K) = \frac{\mathbf{y}(z + \xi, z, \xi, K)}{\mathbf{y}(z, z, \xi, K)} \tag{5}
$$

$$
\mathbf{G}(A, \xi, K) = \mathbf{g}\Big(\mathbf{y}(A, A, \xi, K) + (1 - \delta)K \; ; \; A, \xi, K \Big). \tag{6}
$$

To interpret condition (5), recall that in stage 1 each island believes that, with probability one, TFP satisfies $A = z$ and the signals of all other islands satisfy $z' = A + \xi = z + \xi$. Together with the fact that

8

all islands make the same choices in equilibrium and that the function **y** captures their equilibrium production choices, this implies that the local beliefs of local and aggregate output are given by, respectively, $\hat{y} = \mathbf{y}(z, z, \xi, K)$ and $\hat{Y} = \mathbf{y}(z+\xi, z, \xi, K)$. By the demand function in (1), it then follows that the local belief of the price must satisfy $\hat{p} = \hat{Y}/\hat{y}$, which gives condition (5). To interpret condition (6), recall that all islands end up making identical choices in equilibrium, implying that the available resources of each island in stage 2 coincide with $Y + (1 - \delta)K$, where Y is the aggregate quantity of the final good (aggregate GDP). Note next that the realized production level of *all* islands is given by $\mathbf{y}(A, A, \xi, K)$ and, therefore, Y is also given by $\mathbf{y}(A, A, \xi, K)$. Together with the fact that **g** is the optimal savings rule, this gives condition (6).

Summing up, an equilibrium is a fixed point of the Bellman equations (3)-(4) and the consistency conditions (5)-(6). In principle, one can obtain the global, non-linear, solution of this fixed-point problem with numerical methods. As in the DSGE literature, however, we find it useful to concentrate on the log-linear approximation of the solution around the deterministic steady state.

3.2 Log-linear Solution

To obtain the log-linear solution, we first log-linearize the equilibrium equations around the deterministic steady state. With a slight abuse of notation, we henceforth re-interpret all the variables in terms of the log-deviations of these variables from their steady-state values.

The terms of trade faced by island i are $p_{it} = Y_t - y_{it}$. The associated marginal revenue products of labor and capital are, respectively, $MRPL_{it} \equiv p_{it} + y_{it} - n_{it}$ and $MRPK_{it} \equiv p_{it} + y_{it} - k_{it}$. The optimal behavior of island i is thus characterized by the following system:

$$\nu n_{it} = \mathbb{E}_{it}[MRPL_{it}] - \gamma \mathbb{E}_{it} c_{it} \tag{7}$$

$$\gamma \left(\mathbb{E}'_{it} c_{i,t+1} - c_{it} \right) = (1 - \beta(1-\delta)) \mathbb{E}'_{it}[MRPK_{i,t+1}] \tag{8}$$

$$p_{it} + y_{it} = (1-s)c_{it} + s\iota_{it} \tag{9}$$

$$y_{it} = A_t + \alpha k_{it} + (1-\alpha)n_{it} \tag{10}$$

$$k_{i,t+1} = \delta i_{it} + (1-\delta)k_{it} \tag{11}$$

where $s \equiv \frac{\alpha\beta\delta}{1-\beta(1-\delta)}$ denotes the steady-state investment-to-GDP ratio. The interpretation of these conditions is straightforward: (7) is the labor-supply condition; (8) is the Euler condition; (9) is the resource constraint; and (10) is the production function; and (11) is the law of motion for capital.

To convey the basic idea behind our solution method, consider momentarily a special case that can be solved by hand: let utility be linear in consumption and assume away capital ($\gamma = \alpha = 0$). In this case, the equilibrium can be reduced to the following fixed-point relation:

$$n_{it} = \mathbb{E}_{it}[\chi A_t + \omega N_t], \tag{12}$$

where $\chi \equiv \frac{1}{1+\nu}$ and $\omega \equiv \frac{1}{1+\nu} \in (0, 1)$.[9] Equilibrium employment can therefore be understood as the solution to a static beauty-contest game, of the type found in Morris and Shin (2002) and Angeletos

---

[9]To obtain condition (12), note first that, when $\alpha = 0$, investment is zero, output is given by $y_{it} = A_t + n_{it}$, and the

and Pavan (2007).[10] In this game, a player is an island, her action is local employment, the fundamental is the underlying TFP, χ measures the direct effect of the fundamental on individual outcomes holding constant the aggregate outcomes, and ω measures the degree of strategic complementarity.[11] Importantly, an island responds to ξ_t because, and only because, this shock influences its beliefs about aggregate employment (and thereby its beliefs about its terms of trade).

To solve (12), start by guessing the following policy rule at the individual level:

$$n_{it} = \Lambda_z z_{it} + \Lambda_\xi \xi_t \tag{13}$$

Aggregation gives $N_t = \Lambda_z \bar{z}_t + \Lambda_\xi \xi_t$. Next, note that, due to our specification of priors,

$$\mathbb{E}_{it}[A_t] = z_{it} \quad \text{and} \quad \mathbb{E}_{it}[\bar{z}_t] = \mathbb{E}_{it}[A_t + \xi_t] = z_{it} + \xi_t.$$

It follows that $\mathbb{E}_{it}[N_t] = \Lambda_z z_{it} + (\Lambda_\xi + \Lambda_z^n)\xi_t$. Using this fact in (12), we infer that whenever i expects the others to play according to the rule given by (13), his best response is to set

$$n_{it} = (\chi + \omega\Lambda_z)\, z_{it} + \omega(\Lambda_\xi + \Lambda_z)\xi_t.$$

Matching the coefficients obtained above with those in the proposed policy rule implies that the latter is part of an equilibrium if and only if the following is true:

$$\Lambda_z = (\chi + \omega\Lambda_z) \quad \text{and} \quad \Lambda_\xi = \omega(\Lambda_\xi + \Lambda_z).$$

Solving these two equations for the coefficients Λ_z and Λ_ξ gives

$$\Lambda_z = \frac{\chi}{1-\omega} = \frac{1}{\nu} \quad \text{and} \quad \Lambda_\xi = \frac{\omega}{1-\omega}\Lambda_z = \frac{1}{\nu^2}. \tag{14}$$

We infer that there exists a unique equilibrium and that the policy rule for local employment in this equilibrium is given by (13) along with (14). Finally, using the fact that $z_{it} = A_t$ with probability one, we conclude that the realized aggregate level of output is given by

$$Y_t = A_t + N_t = \Lambda_A A_t + \Lambda_\xi \xi_t,$$

with $\Lambda_A = 1 + \Lambda_z$ and with (Λ_z, Λ_ξ) given as in (14).

Two properties of this solution are worth noting. First, the coefficient Λ_A, which governs the response of Y_t to A_t, is the same as the one in the version of the model that imposes common knowledge of A_t and shuts down the ξ_t shock. Second, the coefficient Λ_ξ, which governs the effect of the ξ_t shock, is proportional to Λ_A by a factor that is itself increasing in ω. That is, the impact of

resource constraint reduces to $c_{it} = p_{it} + y_{it}$, Using these facts into the labor-supply condition (9) gives $(1 + \nu)n_{it} = \mathbb{E}_{it}[p_{it} + y_{it}]$. Finally, using $p_{it} = Y_t - y_{it}$ and $Y_t = A_t + N_t$ results into condition (12).

[10] The term "beauty contests" is often used to refer to a class of coordination games with linear best responses.

[11] In the example under consideration, ω happens to coincide with χ, but this is not generally true. It is therefore best to think of ω and χ as two distinct objects.

the confidence shock relative to that of the technology shock increases with the degree of strategic complementarity. This is because ξ_t matters only by influencing the beliefs of the actions of others.

Go back now to the general case $(\alpha, \gamma > 0)$. The presence of an endogenous state variable (capital) and of forward-looking behavior implies that the economy can be thought of as a *dynamic game* in which the best response of a player (or island) today depends both on past outcomes and on expectations of future outcomes. This complicates the fixed-point problem that needs to be solved. The essence, however, is similar to that in the above example.

We thus start by guessing the following policy rules at the island level:

$$n_{it} = \Lambda_k^n(k_{it} - K_t) + \Lambda_K^n K_t + \Lambda_z^n z_{it} + \Lambda_\xi^n \xi_t \tag{15}$$

$$c_{it} = \Gamma_k^c(k_{it} - K_t) + \Gamma_K^c K_t + \Gamma_z^c z_{it} + \Gamma_{\bar{z}}^c \bar{z}_t + \Gamma_a^c A_t + \Gamma_\xi^c \xi_t \tag{16}$$

$$k_{it+1} = \Omega_k^k(k_{it} - K_t) + \Omega_K^k K_t + \Omega_z^k z_{it} + \Omega_{\bar{z}}^k \bar{z}_t + \Omega_a^k A_t + \Omega_\xi^k \xi_t \tag{17}$$

where Λ^n, Γ^c, and Ω^k are coefficients that remain to be determined. We then proceed to solve for the equilibrium values of these coefficients by solving the fixed-point problem between the individual policy rules and the associated aggregate outcomes imposed by conditions (7)-(11).

To generate data from the model, we set $z_{it} = \bar{z}_t = A_t$ and compute the aggregate outcomes implied by (15)-(17). This gives N_t, C_t and K_{t+1} as functions of the vector (K_t, A_t, ξ_t), verifying that the latter is the state variable for the aggregate outcomes. Note that setting $z_{it} = \bar{z}_t = A_t$ corresponds to invoking the objective truth. However, to solve the fixed-point problem between the individual policy rules (or strategies) and the aggregate outcomes, we have to treat z_{it}, $\bar{z}_t$ and A_t as distinct objects. This is necessary in order to keep track of the difference between the first- and the higher-order beliefs of the underlying fundamental and, thereby, between objective and subjective beliefs.

The details are spelled out in Online Appendix O.5. The bottom line is that we can obtain the solution of our model as a tractable transformation of that of the standard RBC model. Furthermore, this solution has two key properties. First, the coefficients $(\Lambda_K^n, \Gamma_K^c, \Omega_K^k)$ and $(\Lambda_A^n, \Gamma_A^c, \Omega_A^k)$, which determine the impact of the capital stock and the technology shock on aggregate outcomes coincide with those in the standard RBC model. Second, the coefficients $(\Lambda_\xi^n, \Gamma_\xi^c, \Omega_\xi^k)$, which determine the impact of the confidence shock, can be solved as functions of the aforementioned coefficients and a few other coefficients, which themselves capture the degree of strategic complementarity in the economy. These properties mirror those noted in the example above.

The solution strategy described above and the aforementioned properties extend to a large class of linear DSGE models; see Online Appendix O.5. The beliefs-augmented model can thus be simulated and estimated with the same ease as the original model. This facilitates the quantitative explorations conducted in Sections 4 and 5.

3.3 Heterogeneous vs Common Priors

As already mentioned, the main advantage of our approach relative to common-prior, incomplete-information models is its flexibility and its straightforward applicability to macroeconomic models.

A potential cost is that it bypasses the restrictions that the common-prior assumption imposes on the size and dynamics of higher-order uncertainty. We now use an example to illustrate this trade off and to corroborate the claim that our heterogeneous-prior specification can be thought as a proxy for higher-order uncertainty in common-prior settings.

In particular, we show that the tractable example considered in the previous section is observationally equivalent to a common-prior variant, in a sense that will be made precise below. We then derive the restrictions that the common-prior variant imposes on the volatility and the persistence of the kind of belief-driven fluctuations we are interested in.

We start by showing that the special case of our model that was solved by hand in the previous section (namely, the one with $\alpha = \gamma = 0$) is observationally equivalent, in a sense that will be made precise below, to a common-prior variant. This variant is obtained by introducing heterogeneity in TFP and letting trade be done according to random, pairwise, matching across the islands. As in Angeletos and La'O (2013), these modifications allow fluctuations to obtain from correlated noise in the rational beliefs that islands form about their pairwise terms of trade.

Let us fill in the details. TFP in island i is given by $A_{it} = A_t + a_i$, where A_t is the aggregate TFP shock and a_i is an island-specific fixed effect. The former follows a random walk with the same variance as in the heterogeneous-prior economy; the latter is distributed in the cross-section of islands according to a Normal distribution with mean zero and variance $\tilde{\sigma}_a^2$. The aggregate TFP shock is assumed to be common knowledge. Nevertheless, higher-order uncertainty is still present because each island is uncertain about the productivity and the information of its trading partner when choosing employment and production. In particular, the information that island i has in the morning of period t about its current-period match is summarized by the following two signals:

$$z_{it} = a_{m(i,t)} + \tilde{\xi}_t \qquad \text{and} \qquad w_{it} = \tilde{\xi}_t + u_{i,t},$$

where $m(i,t)$ denotes the trading partner of island i in period t, $u_{i,t}$ is orthogonal to $a_{m(i,t)}$, i.i.d. across islands and unpredictable on the basis of past information, and $\tilde{\xi}_t$ is an aggregate shock that is orthogonal to the aggregate TFP shock and that follows an AR(1) process. More specifically,

$$\tilde{\xi}_t = \tilde{\rho}\tilde{\xi}_{t-1} + \tilde{\sigma}_\xi \tilde{\zeta}_t$$

where $\tilde{\zeta}_t \rightsquigarrow \mathcal{N}(0,1)$, $\tilde{\sigma}_\xi > 0$, and $\tilde{\rho} \in [0,1]$. Literally taken, z_{it} is i's private signal about the idiosyncratic TFP of its trading partner; this signal is contaminated by common noise, given by $\tilde{\xi}_t$; and w_{it} is a private signal that is informative about this noise.[12] Clearly, the $\tilde{\xi}_t$ shock plays the same role in this common-prior setting as the ξ_t shock in our heterogeneous-prior setting.

In the absence of the aforementioned shocks, the two economies reduce to the same underlying RBC benchmark and thus give rise, in equilibrium, to the same observables at the aggregate level. Let Y_t^* denote the level of output in that benchmark. From the results of the previous subsection, we have that the equilibrium level of output in the heterogeneous-prior economy is given by $Y_t = Y_t^* + \Lambda_\xi \xi_t$,

---

[12]This signal can be recast as a signal extracted from past trades; see Angeletos and La'O (2013) for details.

with Λ_ξ as in (14). And since ξ_t is an AR(1) process, we conclude that the "output gap" relative to the RBC benchmark is also an AR(1) process:

$$Y_t - Y_t^* = \varphi(Y_{t-1} - Y_{t-1}^*) + \psi\varepsilon_t, \tag{18}$$

where $\varepsilon_t \rightsquigarrow \mathcal{N}(0,1)$ is i.i.d. over time and independent of the technology shock,[13] and where

$$\varphi = \rho \qquad \text{and} \qquad \psi = \frac{\omega\sigma_\xi}{(1-\omega)^2}. \tag{19}$$

Consider next the common-prior economy and let $\tilde{\theta} \equiv (\tilde{\rho}, \tilde{\sigma}_\xi, \tilde{\sigma}_u, \tilde{\sigma}_a)$ collects its informational parameters. Its solution is far from trivial, but can be obtained by adapting Theorem 1 in Huo and Takayama (2015b).[14] We thus have that the output gap in this economy also follows an AR(1) process as in (18), except that now φ and ψ are given by the following:

$$\left.\begin{aligned}
\varphi &= \Phi\left(\tilde{\theta}, \omega\right) \equiv \frac{1}{2}\left[\left(\frac{1}{\tilde{\rho}} + \tilde{\rho} + \frac{1-\omega}{\tilde{\rho}}\frac{\tilde{\sigma}_a^2+\tilde{\sigma}_u^2}{\tilde{\sigma}_a^2\tilde{\sigma}_u^2}\tilde{\sigma}_\xi^2\right) - \sqrt{\left(\frac{1}{\tilde{\rho}} + \tilde{\rho} + \frac{1-\omega}{\tilde{\rho}}\frac{\tilde{\sigma}_a^2+\tilde{\sigma}_u^2}{\tilde{\sigma}_a^2\tilde{\sigma}_u^2}\tilde{\sigma}_\xi^2\right)^2 - 4}\right] \\
\psi &= \Psi\left(\tilde{\theta}, \omega\right) \equiv \frac{\omega\Phi(\tilde{\theta},\omega)}{\tilde{\rho}\left(1-\omega^2\frac{\tilde{\rho}\tilde{\sigma}_a^2+\Phi(\tilde{\theta},\omega)\tilde{\sigma}_u^2}{\tilde{\rho}\tilde{\sigma}_a^2+\tilde{\rho}\tilde{\sigma}_u^2}\right)}\tilde{\sigma}_a
\end{aligned}\right\} \tag{20}$$

By comparing (19) and (20), we can readily prove that the two economies are observationally equivalent in the following sense.

Proposition 1 *Let $\theta \equiv (\rho, \sigma_\xi)$, $\tilde{\Theta} \equiv [0,1] \times \mathbf{R}_+^3$, and $\Theta \equiv [0,1] \times \mathbf{R}_+$; and let $\mathcal{C}(\tilde{\theta})$ and $\mathcal{H}(\theta)$ denote, respectively, the common-prior economy parameterized by $\tilde{\theta}$ and the heterogeneous-prior economy parameterized by θ. Then:*

(i) For any $\tilde{\theta} \in \tilde{\Theta}$, there exists a $\theta \in \Theta$ such that $\mathcal{H}(\theta)$ implies the same stochastic process for the output gap and all the macroeconomic quantities as $\mathcal{C}(\tilde{\theta})$.

(ii) The converse is also true: for any $\theta \in \Theta$, there exists a $\tilde{\theta} \in \tilde{\Theta}$ such that $\mathcal{C}(\tilde{\theta})$ implies the same stochastic process for the output gap and all the macroeconomic quantities as $\mathcal{H}(\theta)$.

The intuition behind this result is that the two economies feature exactly the same variation in the expectations of the relevant economic outcomes: in either economy, a positive (resp., negative) output gap obtains if and only if the firms and the households of each island are optimistic (resp., pessimistic) about the terms of trade, or the demand, that their island is likely to face in the short run. What differs between the two economies is the way these waves of optimism and pessimism are captured: in one economy, they are engineered with the help of a specific departure from rational expectations; in the other, they are instead sustained by rational confusion. Accordingly, whereas the higher-order belief shock is allowed to be common knowledge in the heterogeneous-prior economy,

---

[13]Note that $\varepsilon_t \equiv \frac{1}{\sigma_\xi}\zeta_t$, with ζ_t being the innovation in the confidence shock.

[14]The result contained in Huo and Takayama (2015b) abstracts from the aggregate TFP shock. By adding such a shock but assuming that it is always common knowledge, we guarantee that the same solution applies to the gap $Y_t - Y_t^*$.

it has to be imperfectly observed in the common-prior one. Nevertheless, by choosing the parameters that govern the dynamics of that shock and of the quality of learning in the latter economy, we can always match the stochastic process for the aforementioned expectations in the former economy, and can therefore also generate the same observables at the aggregate level.

This result is subject to the following qualification: the ability to replicate a heterogeneous-prior economy with a common-prior one relies on the freedom to choose a sufficient high $\tilde{\sigma}_a$ in the latter. This is because the level of fundamental, or first-order uncertainty in the common-prior economy—parameterized here by $\tilde{\sigma}_a$—imposes certain bounds on the persistence and the volatility of higher-order beliefs and, equivalently, on φ and ψ. For the heterogeneous-prior economy to respect the same bounds, ρ and σ_ξ must satisfy certain restrictions. Proposition 2 below describes the bounds on φ and ψ; Corollary 1 gives the corresponding restrictions on ρ and σ_ξ.

Proposition 2 *For any $\varphi \in [0, 1)$ and any $\omega \in (0, 1)$, let*

$$B(\varphi, \omega) \equiv \max_{\hat{\rho} \in [0,1], \hat{\sigma}_u \geq 0, \hat{\sigma}_\xi \geq 0} \left\{ \Psi\left(\hat{\rho}, \hat{\sigma}_u, \hat{\sigma}_\xi, 1, \omega\right) \ s.t. \ \Phi\left(\hat{\rho}, \hat{\sigma}_u, \hat{\sigma}_\xi, \omega\right) = \varphi \right\}.$$

A process for the output gap as in condition (18) can be obtained in the equilibrium of a common-prior economy $\mathcal{C}(\tilde{\theta})$ if and only if (i) $0 \leq \varphi < 1$ and (ii) $0 \leq \psi \leq B(\varphi, \omega)\tilde{\sigma}_a$.

Corollary 1 *A heterogeneous-prior economy $\mathcal{H}(\theta)$ can be replicated by a common-prior economy $\mathcal{C}(\tilde{\theta})$ if and only if (i) $0 \leq \rho < 1$ and (ii) $\sigma_\xi \leq \frac{\omega}{(1-\omega)^2} B(\rho, \omega)\tilde{\sigma}_a$.*

Part (i) of Proposition 2 states that the beliefs-driven fluctuations in the common-prior economy are necessarily transitory. This would be true even if we allowed $\tilde{\rho} > 1$, meaning an explosive process for the $\tilde{\xi}_t$ shock. The reason is that these fluctuations are sustained only by rational confusion, which itself fades away as additional information arrives over time. Part (ii), on the other hand, provides a tight upper bound on the volatility of these fluctuations. This bound is proportional to $\tilde{\sigma}_a$, because, as already explained, this parameter pins down the level of first-order uncertainty, which in turn binds the level of higher-order uncertainty.

Corollary 1 converts the above properties into restrictions on the parameters of the heterogeneous-prior specification. Part (i) justifies our earlier assertion that letting $\rho < 1$ helps capture within our framework the property that the fluctuations sustained by higher-order uncertainty have to be transient. Part (ii), on the other, provides an upper bound on σ_ξ.

To recap, we have established two lessons in the context of the example considered. First, the heterogeneous-prior setting is observationally equivalent to a common-prior variant in terms of beliefs-driven fluctuations. Second, the common-prior setting imposes a joint restriction between the magnitude and persistence of these fluctuations and the underlying fundamental uncertainty. Translating this restriction to the heterogeneous-prior setting yields a bound on σ_ξ.

How tight is this bound? In Online Appendix O.1, we use a back-of-the-envelope exercise to argue the following: if we were to approach the US data with the simple model considered in this

subsection, the bound would be large enough to allow for the entire business cycle to be driven by the confidence shock. And although a similar result is not readily available for the estimated models of Section 5, we suspect that our quantitative findings are consistent with realistic common-prior models. The recent work of Huo and Takayama (2015b) seems to corroborate this conjecture.

4   Empirical Properties of the Confidence Shock

We now use a parametrized version of our model to illustrate the comovements induced by the confidence shock on the key macroeconomic quantities. We also explain why these patterns are consistent with salient features of the data and why they are not shared by other parsimonious extensions of the RBC model.[15] We finally elaborate on the sense in which the confidence shock can be thought of as an aggregate demand shock whose ability to generate realistic business cycles does not require either the presence of nominal rigidities or the comovement of the real quantities with inflation.

4.1   Parameterization and IRFs

The parameters are set as follows: the discount factor is 0.99; the elasticity of intertemporal substitution is 1; the Frisch elasticity of labor supply is 2; the capital share in production is 0.3; the depreciation rate is 0.015; and the persistence of the confidence shock is $\rho = 0.75$. The last choice is somewhat arbitrary, but can motivated as follows. First, the implied forecast errors have a half life of less than a year, which is broadly in line with survey evidence in Coibion and Gorodnichenko (2012). Second, the value of ρ assumed here is close to the one estimated in the next section in the context of two medium-scale, DSGE models. Finaly, to the extent that the fluctuations induced by ξ_t in our model resemble either the "demand shock" identified in Blanchard and Quah (1989) or the "main business cycle shock" identified in Angeletos, Collard, and Dellas (2017), our parametrization is consistent with the evidence in those papers as well.[16]

Figure 1 reports the Impulse Response Functions (IRFs) of the model's key quantities to a positive innovation in ξ_t. It is evident that the shock causes output, hours, consumption and investment move in the same direction. But why?

We address this question in two steps. In the rest of this subsection, we explain how the variation in higher-order beliefs of the exogenous fundamental (TFP) translates into variation in the expectations of the aggregate economic activity and the terms of trade. In the next subsection, we clarify how the empirical performance of the theory hinges on the horizon of the latter kind of expectations.

Start by inspecting conditions (7)-(11), which determine the equilibrium behavior. The following property is evident: the optimal behavior of an island depends on its higher-order beliefs of aggregate

---

[15]The performance of our mechanism within richer, medium-scale, DSGE models is addressed in Section 5.

[16]Note that we have not specified σ_a and σ_ξ, the standard deviations of the two shocks. This is not necessary for the purposes of this section, because we focus on comovement patterns and do not attempt to match the overall volatility in the data. See, however, the remarks in footnote 21.
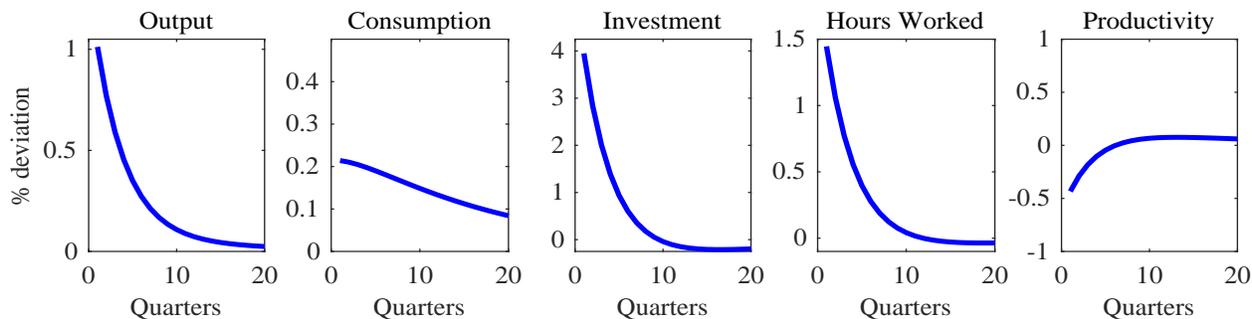
Figure 1: Impulse responses to a positive confidence shock

TFP *only* through its first-order beliefs of its terms of trade, which in turn coincide with its first-order beliefs of aggregate output. This reveals the ultimate modeling role of the ξ_t shock, which is to capture extrinsic variation in the expectations of the relevant economic outcomes.

This perspective applies more generally. In the class of models we are interested in, the equilibrium expectations of the endogenous outcomes can be expressed as a function of the hierarchy of beliefs about the underlying fundamentals regardless of the information structure. However, different assumptions about the information structure lead to different predictions about the stochastic properties of the expectations of economic outcomes. In the standard practice, these expectations are spanned by the expectations of fundamentals because higher-order beliefs collapse to first-order beliefs. By contrast, our approach leaves room for autonomous variation in the expectations of economic outcomes by letting the higher-order beliefs to deviate from the first-order beliefs.

4.2 The Key Mechanism: Beliefs about the Short-Term Economic Outlook

So far, we have argued that it is best to think of ξ_t shock as a modeling device for introducing autonomous variation in the expectations of the relevant economic outcomes. This is important, but it is not the whole story. Because behavior is forward looking, the horizon of these expectations is a crucial determinant of how actual outcomes respond to shifts in them. We now build on this basic observation to explain why the comovement patterns seen in Figure 1 hinge on the property that the ξ_t shock captures expectations of the short-term economic outlook, as opposed to expectations of the medium- or long-run prospects.

To reveal the short-term nature of the belief waves triggered by the ξ_t shock, we present the effects of the shock on the "term structure of expectations". Consider, in particular, the forecasts that island i forms in period t about about its terms of trade k periods ahead, namely, $\mathbb{E}_{i\tau}[p_{i,\tau+k}]$, for all $k \geq 1$. As already noted, these forecasts are tied to the forecasts that the firms make about their sales, that the households form about their income, and that everybody forms about aggregate output. Figure 2 draws the average forecast at different horizons (namely, for $k \in \{1, \ldots, 12\}$), both at the moment the shock hits the economy (solid line) and four quarters later (dashed line).
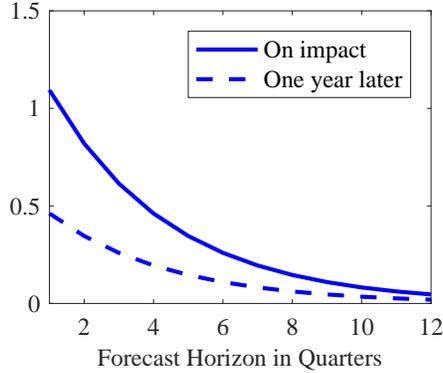
Figure 2: Forecasts of terms of trade, following a confidence shock

As is evident in the figure, a positive innovation in ξ_t raises the expected expected terms of trade in the next few quarters *without* moving the corresponding expectations at longer horizons. The same point applies to the forecasts of the aggregate levels of output, hours, consumption and investment. As time passes, the optimism fades away and the curve in Figure 2 shifts down. Nevertheless, the curve remains downward-sloping, underscoring that the waves of optimism (and pessimism) accommodated in our paper regard *exclusively* the short-term economic outlook.

This property is key to understanding the comovement patterns seen in Figure 1. In the eyes of the firms, a positive innovation in ξ_t means a short-lived increase in the expected demand for their product. To take advantage of this, the firms raise their demand for both labor and capital, pushing the wage and the rental rate of capital up. As a result, the households experience a transitory increase in their income and in the returns to labor and capital. Because this entails only a small increase in permanent income, the wealth effect on labor supply is easily dominated by the competing substitution effect. This guarantees that hours, and hence also output and income, increase in equilibrium. Finally, because the boom is expected to be transitory, the households find it optimal to consume only a fraction of the realized increase in their income and to save the rest. All in all, the shock therefore causes a joint increase in hours, output, consumption and investment, and without a commensurate shift in TFP and labor productivity, just as seen in Figure 1.

As noted in the introduction, this mechanism is different from the one in the literature on news and noise shocks (Beaudry and Portier, 2006; Jaimovich and Rebelo, 2009; Lorenzoni, 2009; Barsky and Sims, 2011). To illustrate the difference, consider Barsky and Sims (2012), an example of that literature that accommodates both news and noise shocks. Aggregate TFP is given by $A_t = A_{t-1} + \gamma_{t-1} + \varepsilon_{a,t}$, where $\gamma_t = \rho_\gamma \gamma_{t-1} + \varepsilon_{\gamma,t}$, $\rho_\gamma \in (0, 1)$, and $\varepsilon_{a,t} \sim \mathcal{N}(0, \sigma_a^2)$ and $\varepsilon_{\gamma,t} \sim \mathcal{N}(0, \sigma_\gamma^2)$ are independent of one another and serially uncorrelated. Furthermore, the representative agent observes A_t perfectly, but only receives a noisy signal of γ_t. Finally, this signal is given by $z_t = \gamma_t + \eta_t$, where $\eta_t \sim \mathcal{N}(0, \sigma_\eta^2)$ is uncorrelated over time and independent of the current and past values of the innovations $\varepsilon_{a,t}$ and $\varepsilon_{\gamma,t}$. In this formulation, $\varepsilon_{\gamma,t}$ moves both the expectations and the actual realizations of future TFP,

17

whereas $\varepsilon_{\eta,t}$ moves the expectations without moving the actual realizations. The former represents a news shock, the latter a noise shock.
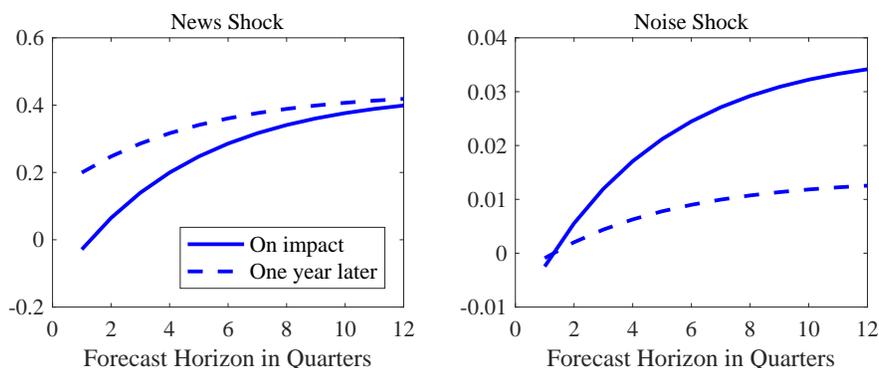
Figure 3: Forecasts of output at different horizons, following a news and a noise shock.

Figure 3 reports the impact of these shocks on the expectations of aggregate output at different horizons, both right after the realization of the shock (solid line) and four quarters later (dashed line). The left panel corresponds to the news shock, the right to the noise shock. By comparing the two panels, we see that the two shocks have qualitatively similar effects on impact. As time passes and more information arrives, the agents can tell whether the initial shift in their beliefs was due to a true increase in the long-run level of TFP or due to noise. This explains why the effects of the news shock get reinforced with the passage of time, while those of the noise fade away. The nature of optimism, however, is the same across these two cases—and it is very different from the one seen in Figure 2. While the confidence shock shifts the expectations of the short-term economic outlook, the news and noise shocks shift expectations of the medium- and long-run prospects.

It is precisely this difference that accounts for the superior quantitative performance of our mechanism. As already explained, the confidence shock triggers small shifts in expected permanent income and large shifts in the expected short-run returns to capital and labor. The opposite is true with the kind of news and noise shocks studied in the extant literature. When a positive news or noise shock hits the economy, the firms do not change their demand for labor and capital because they perceive no immediate change in their short-term returns, but the households reduce both their supply of labor and their saving because they expect higher wages and higher income in the future: a positive news shock is a good time both to consume more and to take a vacation. As a result, the equilibrium levels of employment and investment move in the opposite direction than that of consumption, which in turn explains why these shocks fail to generate realistic business cycles within baseline versions of either the RBC or the New Keynesian model.[17]

---

[17] To overcome this challenge, Jaimovich and Rebelo (2009) augment the baseline RBC model with adjustment costs that makes investment today increase in anticipation of higher investment in the future; and with a particular form of internal habit that generates a negative income effect on leisure in the short run. Lorenzoni (2009), on the other hand, abstracts from investment, adds nominal rigidity, and lets monetary policy induce pro-cyclical output gaps. Yet, this mechanism

4.3 Conditional Moments

We have shown that the confidence shock produces transitory comovements in the key macroeconomic quantities, without commensurate movement in TFP and labor productivity. We have also offered the economic intuition for this result. But is this prediction consistent with the data?

One could imagine answering this question by obtaining an empirical counterpart of ξ_t from surveys of higher-order beliefs of TFP. However, such surveys are not available. But even if they were available, they would only help under a literal, narrow interpretation of ξ_t, which is not our preferred way to think about the *applied* contribution of our paper. Instead, we believe that this contribution is maximized by interpreting ξ_t as a proxy for autonomous variation in the *first-order* beliefs of the endogenous economic outcomes over the business cycle—think of the expectations that the firms form about the demand for their products, or those that the consumers in turn form about their employment and income.

Because these expectations are part of the equilibrium and are jointly determined with the actual outcomes, it is unclear how one could identify ξ_t through, say, a SVAR approach analogous to those used in the identification of technology and monetary shocks. Lacking a better alternative, we thus proceed to evaluate the empirical performance of our theory in a more indirect way, by comparing two sets of conditional moments: those generated in our model by the confidence shock alone; and those observed in the data after filtering them from the effects of an empirical proxy of the technology shock. We view this comparison of conditional moments as a crucial test of any parsimonious theory that aspires to improve upon the baseline RBC model: if such a theory fails to account for the TFP-filtered "residuals" of the data, then it fails to achieve this objective.

We obtain the relevant component of the data in one of two ways. In the one, we regress each variable of interest on the current level and the four lags of TFP, as measured by Fernald (2014), and extract the residuals. In the other, we include all the variables in a SVAR; identify the technology shock as in Galí (1999), that is, as the only shock that exerts an effect on labor productivity in the long-run; and then take the residuals from the projection of the data on the identified technology shock. Although none of these approaches offers a bullet-proof identification of the technology shock, they generate macroeconomic variables that can be used to test parsimonious theories that seek to explain the business cycle with a single shock besides the standard technology shock.

The first two columns of Table 1 report the relevant moments in the data, namely, the business-cycle correlations and the relative volatilities of the aforementioned residuals, under the two specifications described above.[18] The third column reports the relevant moments in our model, namely, the correlations and relative volatilities induced by the confidence shock. The information contained in

---

is not sufficiently strong to offset the negative comovement of the underlying flexible-price quantities once investment is added to that model, unless additional "bells and whistles" are added to the model.

[18]The moments have been computed on bandpass-filtered series at frequencies corresponding to 6–32 quarters. This filter is preferable to the simpler HP filter because it removes not only low-frequency trends but also high frequency "noise" such as seasonal fluctuations and measurement error; see Stock and Watson (1999). Note, though, that the picture that emerges from Table 1 is not sensitive to the choice of the filter.

Table 1: Conditional Comovements (6-32 quarters)

	Filtered Data		Our Theory	Alternative Theories			
	(a)	(b)	ξ shock	I shock	C shock	News shock	E shock
σ_n/σ_y	1.05	1.29	1.43	1.44	1.44	0.74	0.59
σ_c/σ_y	0.63	0.41	0.25	1.21	1.19	0.35	0.19
σ_i/σ_y	3.35	4.04	3.92	8.93	8.93	5.06	4.26
$\sigma_{y/h}/\sigma_y$	0.45	0.63	0.44	0.51	0.49	0.37	0.43
$corr(c,y)$	0.86	0.80	0.85	-0.92	-0.94	-0.17	0.65
$corr(i,y)$	0.94	0.95	0.99	0.98	0.99	0.97	0.99
$corr(n,y)$	0.91	0.88	0.99	0.98	0.98	0.95	0.99
$corr(c,n)$	0.80	0.93	0.81	-0.98	-0.99	-0.46	0.51
$corr(i,n)$	0.86	0.82	0.99	0.99	0.99	0.99	0.99
$corr(c,i)$	0.73	0.76	0.78	-0.98	-0.98	-0.40	0.55
$corr(y,y/n)$	0.12	-0.23	-0.96	-0.79	-0.84	0.78	0.97
$corr(n,y/n)$	-0.31	-0.66	-0.98	-0.91	-0.92	0.56	0.92

<u>Note:</u> Columns (a) and (b) refer to the residuals that obtain, respectively, from the projection of the data on current and past TFP and from the removal of the technology shock identified in the same was as in Galí (1999). All other columns refer to theoretical predictions.

this column is, of course, the same as the one contained in the IRFs of the confidence shock: the shock causes hours, output, consumption and investment to comove, without commensurate comovement in labor productivity. The next three columns report the corresponding moments for three other candidate shocks, which are often used in the literature as proxies for demand shocks: a discount-rate, or consumption-specific, shock; an investment-specific shock, and a news shocks.[19] The last column considers a transitory shock to the efficiency wedge; this can be thought of as a proxy for the supply-side effects of financial or uncertainty shocks.[20]

The main lesson that emerges from inspection of Table 1 is that the confidence shock does a good job in matching the conditional patterns in the data both absolutely and relatively to the other shocks. This is because none of the aforementioned demand shocks is able to generated positive comovement between hours, consumption, and investment within the baseline RBC model; and the efficiency-wedge shock can generate such comovement only by predicting a positive comovement between hours and labor productivity, which is exactly the opposite of what is observed in the data.

As shown in Online Appendix O.1, the same picture emerges if we consider a New Keynesian

---

[19]To obtain the predictions of each of these alternative shocks, we maintain the parameterization of preferences and technologies and merely replace the confidence shock with the considered alternative.

[20]Such a shock is not removed by the specification used by Galí (1999), because that approach identifies only permanent technology shocks. It may also not be removed by our specification based on regressing the macroeconomic quantities on current and past TFP to the extent that there is measurement error in the available TFP measure.

variant that adds sticky prices and lets monetary policy follow a realistic Taylor rule. In principle, these modifications help improve the empirical performance of the aforementioned kind of demand shocks by letting these shocks induce procyclical output gaps, that is, by letting output increase relative to its flexible-price counterpart in response to positive demand shock. Yet, unless one adds various bells and whistles, the predicted output gaps are not large enough to undo the counterfactual comovement properties of the underlying flexible-price allocations.

Of course, these findings do not mean that no other model can match the conditional moments reported in the first two columns. For instance, DSGE models such as Smets and Wouters (2007) are able to do so by attributing the aforementioned residuals to the joint contribution of several shocks, despite the fact that none of these shocks can by itself generate the right comovement patterns. Nevertheless, these findings illustrate in a simple and transparent manner that our theory does better relative to a number of comparable, parsimonious formalizations of either demand- or supply-driven fluctuations—a property that we view as valuable.

Additional support is provided by the evidence in a companion paper (Angeletos, Collard, and Dellas, 2017), where we use a SVAR approach to document that the bulk of the business-cycle volatility in output, hours, investment and consumption in US data can be accounted for by a single shock whose IRFs look very much like those seen in Figure 1. A similar picture is also painted in Section 5, where the confidence shock emerges as the main driver of the business cycle within medium-scale DSGE models that contain multiple other shocks.[21]

4.4 Wedges, Output Gaps, and Aggregate Demand

We conclude this section by offering two additional perspectives on the empirical performance of our theory and its interpretation.

Suppose first that one approaches the data generated by our model through the lenses of the RBC model augmented with various wedges, as suggested by Chari, Kehoe, and McGrattan (2007). In

---

[21]Throughout this section, we have focused attention on comparing features of the data to theoretical counterparts that do not require us to parameterize the standard deviation of either the confidence shock or the technology shock: the IRFs seen in Figure 1, and the conditional correlations and relative volatilities reported in Table 1 are invariant to the choice of σ_ξ and σ_a. But what about the ability of our baseline model to capture the unconditional moments of the data? Clearly, this depends on the choice of σ_ξ and σ_a. Suppose we pick σ_ξ and σ_a so as to minimize the distance between the unconditional volatilities of hours, output, consumption and investment predicted by our baseline model from those found in the data. This exercise yields $\sigma_a = 0.79$ and $\sigma_\xi = 5.77$; it also attributes almost all of the volatility of hours to the confidence shock. We find these properties of our baseline model problematic for two reasons. First, σ_ξ is too large compared to σ_a, a property that questions the plausibility of the interpretation of ξ_t as a bias in the signals of aggregate TFP. (We thank a referee for pointing this out.) Second, our prior is that coordination frictions cannot possibly explain so much of the business cycle. In Section 5, we alleviate the first concern, not only by allowing for other shocks to absorb part of the volatility in the data, but also by modifying the degree of strategic complementarity. This concern can also be alleviated by re-interpreting ξ_t as a shock to higher-order beliefs of idiosyncratic fundamentals, and thereby to first-order beliefs of idiosyncratic terms of trade, along the lines discussed in Angeletos and La'O (2013), Huo and Takayama (2015), and Subsection 3.3 of our paper. Regarding the second concern, we are open to the idea that our mechanism is proxying for other forces, whose effects are similar to those of the confidence shock but whose micro-foundations remain to be discovered.

Online Appendix O.1, we show that the confidence shock manifests itself as a combination of wedges in the equilibrium conditions that characterize the behavior of the households and of the firms. This is true whether we consider the total wedges between the marginal rates of substitution and the corresponding marginal rates of transformation, or their household-side and firm-side components. What is more, the predicted wedges are consistent with those estimated in the data.

These findings speak to our theory's ability to capture the "residuals" between the data and the predictions of the baseline RBC model. More generally, they illustrate how higher-order uncertainty offers a theory of beliefs-driven wedges. The wedges emerge because, and only because, the agents in the model use a distorted expectations operator relative to the complete-information, common-prior, fully-rational benchmark. The magnitude and correlation structure of these wedges is tied to the underlying structure of the market interactions and the degree of strategic complementarity. For instance, were we to shut down trade across islands in our own model, strategic complementarity and wedges would vanish.[22]

Suppose next that one tries to interpret the data generated by our model through the lenses of the New Keynesian framework. In our setting, prices are flexible. Yet, because firms make their input choices prior to observing the demand for their products, a drop in confidence can manifest itself as an increase in the realized markup. Furthermore, the resulting recession will register as a negative output gap insofar as the latter is measured relative to the frictionless RBC benchmark, a property clearly illustrated by the example in Subsection 3.3. Consequently, an adverse confidence shock in our setting looks like a negative demand shock in the New Keynesian model.

Nonetheless, there is an important difference: in our setting, fluctuations in this output gap can arise without any variation in inflation. This is because our mechanism does not need to satisfy the restriction between the output gap and inflation imposed by the New Keynesian Philips Curve, or its ancestors. We view this as an advantage of our theory because the aforementioned restriction receives little support from the data, as the empirical Philips-curve literature has demonstrated; see Mavroeidis, Plagborg-Møller, and Stock (2014) for a review. The evidence provided in a companion paper (Angeletos, Collard, and Dellas, 2017) also speaks against this restriction and in favor of a mechanism like the one accommodated here: in that paper, we use a SVAR approach to document that the bulk of the business-cycle variation in output can be explained by a structural shock that can be thought of as an "non-inflationary demand shock" in the sense that it triggers strong comovement between employment, output, consumption and investment at the business-cycle frequencies without commensurate comovements in either TFP and labor productivity or inflation at any frequency.[23] Last but not least, the inflation-output implications of Phillips curves are hard to reconcile with the Great Recession, where the severe contraction in output and employment was not accompanied by severe deflation.

---

[22] These points indicate a close relation between our paper and recent work that considers other forms of belief distortions, such as Ilut and Saijo (2016), Bhandari, Borovicka, and Ho (2016), and Pei (2017).

[23] The empirical IRFs of the shock identified in our companion paper are actually quite similar to the theoretical IRFs of the confidence shock in the present paper.

With this backdrop, we like to interpret our confidence shock as a form of demand shock that does not hinge either on nominal rigidity or on the inflation-output nexus of the Keynesian paradigm.[24] One may, however, object to this interpretation on the following grounds. In our model, employment and output are fixed in the morning of each period, whereas consumption and investment are determined in the afternoon. In this sense, supply is determined first and prices adjust to make sure that demand meets supply. By contrast, the Keynesian paradigm assumes that prices are determined first and supply adjusts to meet demand.

We now show that changing the timing of decisions in our model so that demand is determined first does not change the nature of the business-cycle fluctuations generated by the confidence shock. We establish this by requiring that consumption and investment be fixed in the morning of each period and letting employment and output adjust in the afternoon. This timing protocol makes our model seem more in the spirit of the Keynesian view that "demand drives supply". Yet, as is evident in Figure 4, the observable implications differ very little across the two protocols.
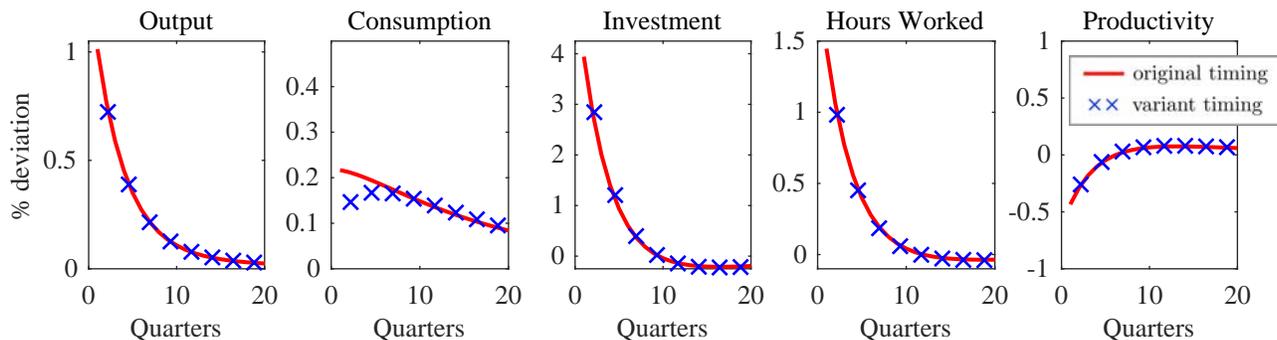
Figure 4: Impulse Responses to a Positive Confidence Shock, under Different Timing Protocols

The solid red lines in Figure 4 repeat the IRFs of the baseline model (previously reported in Figure 1). The blue crosses report the IRFs of the variant in which consumption and investment are determined first. With the exception of consumption, where there is only a modest difference, the IRFs of the two models line up almost perfectly on top of each other. Not surprisingly, this similarity extends to the kind of business-cycle moments we reported earlier in Table 1.

Let us explain why. In the baseline, supply-first version of our model, a positive confidence shock causes employment and output to increase in the morning. The overall spending therefore *has* to increase in the afternoon. Its composition, however, is free to adjust. The only reason that consumption and investment co-move at that point is that the optimism applies only to the short run—which is also the reason why employment increases in the first place during the morning. In the demand-first variant, consumption and investment are determined first. The only reason that they both increase in response to a positive confidence shock is, once again, that the shock causes the agents to become

---

[24]In this regard, our work is related to that of Beaudry and Portier (2013), which offers a different theory of non-inflationary, demand shocks but does not explore its quantitative potential.

optimistic about the short run. If, instead, the shock caused the agents to become optimistic about income in the medium to long run, the agents would like to borrow against their future income, so consumption and investment would move in the opposite direction.

We conclude that the predictions of our theory are not unduly sensitive to the timing protocol and, in this sense, to whether output is supply- or demand-determined. Each protocol, however, is useful for different purposes. On the one hand, the protocol used in our baseline model is the same as the one assumed in the related works of Angeletos and La'O (2013), Benhabib, Wang, and Wen (2015), Huo and Takayama (2015b), and Ilut and Saijo (2016). On the other hand, the variant introduced here better captures the Keynesian spirit of demand-driven fluctuations; it seems more consistent with the idea of sluggish adjustment in aggregate demand;[25] and it boosts the degree of strategic complementarity, helping generate larger fluctuations in the macroeconomic quantities out of the same volatility in higher-order beliefs.[26] For all of these reasons, we opt for the new protocol in the quantitative explorations conducted in the next section.

5   Extension and Estimation

In this section, we apply our method to two medium-scale DSGE models, which are estimated using US data. This requires the introduction of various bells and whistles, which are standard in the DSGE literature but are at odds with our desire for parsimony. The main goal of this section is therefore, not to write and estimate our preferred models, but rather to illustrate the robustness of our theoretical mechanism as we move from the baseline RBC model to richer DSGE models, and as we switch on and off the role of nominal rigidities and monetary policy.

5.1   Two medium-scale models

We start with a brief description of the main features of the two models. A more detailed description and the relevant equations can be found in the Appendix.

In order to accommodate price-setting behavior, we let each island contain a large number of monopolistic firms, each of which produces a differentiated commodity. These commodities are combined through a CES aggregator into an island-specific composite good, which in turn enters the production of the final good in the mainland through another CES aggregator. The elasticity parameter

---

[25]Such sluggishness is captured in the DSGE literature with habit persistence in consumption and adjustment costs in investment of the type considered in the next section.

[26]To understand why, consider the knife-edge case in which the income and the substitution effects of higher terms of trade on labor supply cancel each other out. This eliminates the macroeconomic effects of the confidence shock in the supply-first version of the model: as the income and substitution effects of the confidence shock offset each other, labor and production do not move in the morning, implying that consumption and investment also do not move during the afternoon. In the demand-first version, however, the confidence shock continues to generate a realistic business cycle: the expectation of better terms causes consumption and investment to increase in the morning, requiring higher employment and production in the afternoon.

in the first aggregator is denoted by η and pins down the monopoly markup; the one in the second aggregator is denoted by ϱ and controls, in conjuction with all the other preference and technology parameters, the degree of strategic complementarity across the islands.[27]

In one of the two models, firms are free to adjust their price in each and every period, after observing the realized demand for their product (the flexible-price model). In the other, firms can only adjust prices infrequently, in the familiar Calvo fashion (the sticky-price model). The latter models also contains a conventional Taylor rule for monetary policy.

In order to let other business cycle drivers compete with our mechanism we include several, additional shocks: a permanent and a transitory TFP shock; a permanent and a transitory investment-specific shock; a news shock regarding future productivity; a transitory discount-rate shock; a government-spending shock; and, in the sticky-price model, a monetary shock.[28]

We finally introduce adjustment costs in investment and habit persistence in consumption, of the type assumed in Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2007). Although these modeling features lack compelling micro-foundations, they have become standard in the literature because they serve, not only as sources of persistence, but also as mechanisms that help improve the empirical performance of certain shocks, including monetary, investment-specific, discount-rate, and news shocks. Their inclusion make our results more easily comparable to those in the literature and also and gives these competing shocks a better chance to outperform the confidence shock.

5.2 Estimation

We estimate our models using Bayesian maximum likelihood in the frequency domain, focusing on business-cycle frequencies. The method is described in the Appendix. Here, we discuss briefly the rational behind this empirical strategy, the data used, and the priors and the posteriors.

Rationale. The models described above—like other business-cycle models—cater to business-cycle phenomena and therefore omit shocks and mechanisms that may account for medium- to long-run phenomena, such as trends in demographics and labor-market participation, structural transformation, regime changes in productivity growth or inflation, and so on. Because of this omission, estimating our models by simple maximum likelihood is likely to lead to erroneous inferences about their business-cycle properties. This is because the estimation will guide the parameters of the model

---

[27]The baseline model is nested with $\eta = 0$ and $\varrho = 1$. Letting $\eta > 0$ accommodates monopoly power. Letting $\varrho \neq 1$ helps parameterize the degree of strategic complementarity.

[28]The motivation for the inclusion of these particular shocks is as follows. First, previous research has argued that investment-specific technology shocks are at least as important as neutral, TFP shocks (Fisher, 2006). Second, monetary, fiscal, and transitory discount-rate or investment-specific shocks, as well as news shocks, have been proposed as formalizations of the notion of "aggregate demand shocks" within the NK framework. Third, transitory TFP, investment-specific, or discount-rate shocks are often used as proxies for financial frictions that lead to, respectively, misallocation, a wedge in the firm's investment decisions, or a wedge in the consumer's saving decisions; see Christiano, Eichenbaum, and Trabandt (2015) for a recent example of these short-cuts. Fourth, the introduction of multiple transitory shocks, whatever their interpretation, increases the chance that these shocks will pick up the transitory fluctuations in the data.

towards matching all the frequencies of the data, as opposed to only those that pertain to business-cycle phenomena. In a nutshell, there is a risk of contamination of the estimates of a model by frequencies that the model was not designed to capture.

This problem was first discussed by Hansen and Sargent (1993) and Sims (1993) in the context of seasonal adjustment, but the logic applies more generally. Sala (2015) has recently documented the relevance of this problem for standard DSGE practice: estimating the model of Smets and Wouters (2007) over different frequency bands leads to different estimates of the model's impulse responses and of the underlying parameters, a fact that underscores the importance of making a judicious selection of the band of frequencies used to estimate the model.

Figure 5 indicates that this concern may be particularly relevant in the context of the exercise carried out in this section. This figure inspects the spectral density of hours.[29] The red line corresponds to the raw data; the blue line results from application of a bandpass filter that keeps only the business-cycle frequencies, namely those ranging from 6 to 32 quarters. The figure reveals substantial movements at the medium and long-run frequencies. Such movements may originate from changes in demographics or in the labor-market participation of women, structural transformation, and other mechanisms which our models have neither hope nor ambition to capture.

Raw Data, ——— Bandpass Filtered Data, 6-32 Quarters.
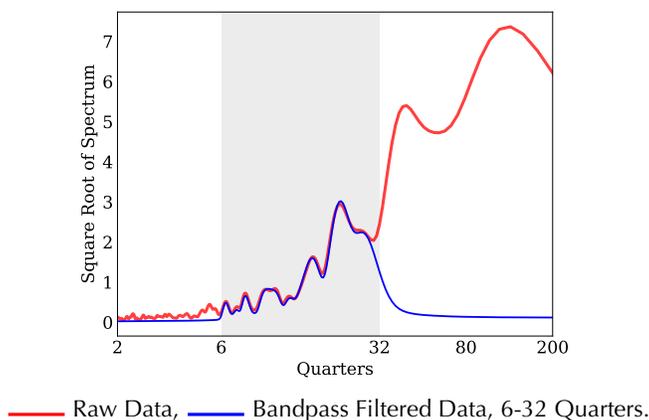
Figure 5: Spectral Density of Hours, 1960Q1-2007Q4.

There are two possible ways to try to mitigate the problem. One is to add the missing mechanisms that would enable the model(s) to account for all the frequencies at once. Another is to estimate the model(s) on the basis of only the business-cycle frequencies. We follow the latter route because of two reasons. First, while we believe that our mechanism and the models considered in this paper are useful for understanding business-cycle phenomena, we are relatively less confident about the "right" choice of mechanisms that can account for the medium- to long-term phenomena; adding the

---

[29]The spectrum is computed as the smoothed periodogram, a Hamming window with a bandwidth parameter of 15 is used, and the x-axis is represented in periods rather than frequencies to ease interpretation. A similar figure appears in Beaudry, Galizia, and Portier (2015), although that paper uses it towards a different goal: to motivate a model that actually connects the short to the medium rum.

"wrong" mechanisms could aggravate the mis-specification problem. Second, we believe that low frequencies of the data contain relative little information about the business-cycle properties of the model, especially those that regard the confidence shock or any other transitory shock; inclusion of the low frequencies is therefore more likely to contaminate, than to improve, the estimation of the business-cycle properties.

Data. The data used in the estimation includes GDP, consumption, investment, hours worked, the inflation rate, and the federal fund rate for the period 1960Q1 to 2007Q4; a detailed description is in Online Appendix O.2. The first four variables are in logs and linearly de-trended; the remaining two are in percentage points.Our sticky-price model is estimated on the basis of all these six variables while flexible-price model is estimated on the basis of real quantities only (GDP, consumption, investment, and hours). The rationale is that the latter model is not designed to capture the properties of nominal data.

Remark on ϱ and σ_ξ. A challenge faced in the estimation of the two models is the following. Consider the parameter ϱ. Holding constant all the other parameters, this parameter governs the degree of strategic complementarity across the islands. In so doing, this parameter also governs the magnitude of the response of the macroeconomic quantities to the confidence shock, without however affecting their covariation structure. It follows that this parameter cannot be identified separately from σ_ξ, the standard deviation of the confidence shock, on the basis of the macroeconomic times-series alone.

For our main estimation exercise, we fix ϱ exogenously at 0.75; this yields an estimate for σ_ξ that is lower than the estimated volatility in aggregate TFP. In Online Appendix O.4, we motivate this choice with an exercise that tries to identify both parameters jointly by combining the macroeconomic times series with the time series of the University of Michigan Index of Consumer Sentiment and by making an assumption about how to extract the expectations that are relevant for our theory from that index. This leads to an estimate of ϱ that is in the neighborhood of 0.75 and to results that are similar to those reported below.

However, we do not wish to push this exercise too far, because it hinges on delicate assumptions about the mapping between that index and our theory. We thus invite the reader to adopt a broader perspective in thinking about what the estimation results mean for our theory. Namely, that they illustrate that the considered models can match the data with plausible assumptions about the magnitude of the underlying higher-order uncertainty, but leave unanswered the delicate question of whether and how additional discipline in the estimation of the confidence shock could be provided from sources outside the standard macroeconomic time series.

Priors and posteriors. The priors and the estimated values of all the parameters are reported in Table 8 in Online Appendix O.3 and broadly in line with the literature. Posterior distributions were obtained with the MCMC algorithm. The estimated values of the preference, technology, and monetary parameters are similar to those found in the extant literature, an indication that the only essential difference from the state of the art is the accommodation of the confidence shock.

27

5.3 Results

Here we review the main findings. Online Appendix O.3 contains additional results.

The confidence shock. Figure 6 reports the estimated IRFs to a positive confidence shock. The solid blue lines correspond to the flexible-price model, the red dashed lines to the sticky-price model. As far as real quantities are concerned, the IRFs are similar across the two models, as well as similar to those in our baseline model. The introduction of investment-adjustment costs and consumption habit adds a hump but does not alter the comovement patterns found in the baseline model. This underscores the robustness of the key positive implications of our mechanism as we move across RBC and NK settings, or, as we add various bells and whistles.
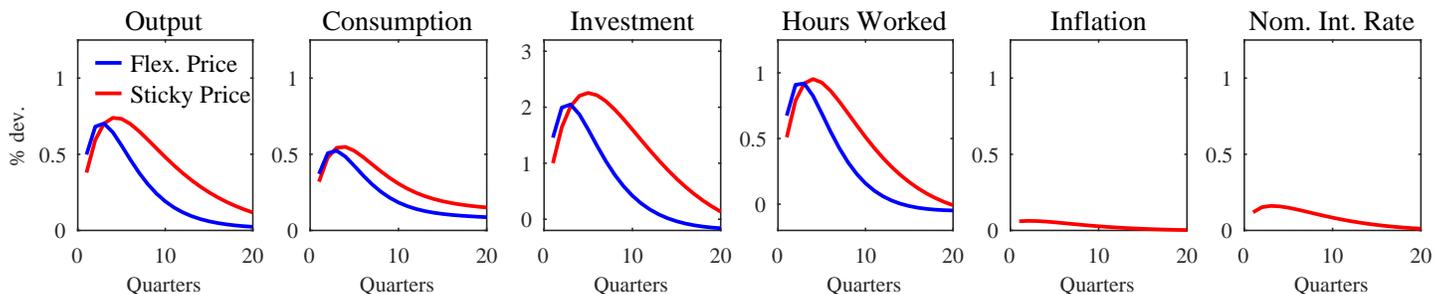
Figure 6: Theoretical IRFs to Confidence Shock

The top half of Table 2 reports the estimated contribution of the confidence shock to the volatility of the key macroeconomic variables at business-cycle frequencies (6–32 quarters). Despite all the competing shocks, the confidence shock emerges as the single most important source of volatility in real quantities. For example, the confidence shock accounts for 55% of the business-cycle volatility in output in the flexible-price model, and for 51% in the sticky-price model.

Table 2: Contribution of Confidence Shock (6–32 Quarters)

Variances	Y	C	I	N	π	R
Flexible Prices	54.72	70.21	41.60	68.32	–	–
Sticky Prices	51.28	61.95	38.50	64.15	11.64	40.84
Covariances	(Y, N)	(Y, I)	(Y, C)	(N, I)	(N, C)	(I, C)
Flexible Prices	74.88	53.74	78.49	66.41	105.08	94.73
Sticky Prices	68.10	50.83	70.95	58.29	104.26	94.89

The bottom half of Table 2 completes the picture by reporting the estimated contribution of the confidence shock to the *covariances* of output, hours, investment, and consumption. The confidence

shock is, by a significant margin, the main driving force behind the comovement of all these vari-ables, underscoring once again the ability of our theory to capture this comovement. In particular, confidence shock explains *more* that one hundred percent of the covariance between hours and consumption, precisely because, as anticipated in the previous section, many of the other structural shocks tend to generate the opposite comovement than the one seen in the data.

Are these findings too good to be true? It depends on how one reads them. In our eyes, they do not mean that our theory is the "true" explanation of the business cycle. They nevertheless reinforce the lessons of the previous section: not only is our theory consistent with salient features of the data, but it is also more potent than other, more familiar, structural interpretations of the data.

Table 3: Moments (6-32 quarters)

	Data	FP	SP	SW	Data	FP	SP	SW
	Standard Deviations				*Correlations with Output*			
Y	1.41	1.28	1.36	1.42				
I	5.12	4.46	4.88	4.86	0.94	0.88	0.86	0.74
N	1.56	1.59	1.66	0.97	0.87	0.82	0.83	0.81
C	0.76	0.82	0.91	1.11	0.85	0.78	0.77	0.67
Y/N	0.76	0.91	0.90	0.84	0.07	-0.02	-0.03	0.74
π	0.23	–	0.25	0.34	0.21	–	0.37	0.13
R	0.35	–	0.34	0.35	0.33	–	0.54	0.06
	Correlations with Investment				*Correlations with Hours*			
N	0.82	0.79	0.82	0.67				
C	0.73	0.56	0.47	0.30	0.83	0.65	0.58	0.59
Y/N	0.07	-0.14	-0.21	0.47	-0.43	-0.58	-0.56	0.22
π	0.09	–	0.41	0.18	0.44	–	0.48	0.23
R	0.23	–	0.60	0.23	0.61	–	0.70	0.21

Note: FP and SP stand for our estimated flexible- and sticky-price models, respectively. SW stands for the model in Smets and Wouters (2007).

Business-cycle moments. Table 3 reports some key moments of the data (first column); those predicted by our estimated models (second and third column); and, for comparison purposes, those predicted by the model in Smets and Wouters (2007) (fourth column). Inspection of this table leads to the following conclusions. First, both of our models do a good job on the real side of the economy. Second, our sticky-price model does a good job in matching also the nominal side of the data. Finally, our sticky-price model appears to outperform the model of Smets and Wouters (2007) in terms of matching the moments of the real quantities as well as the correlations of the nominal variables with output and hours. Of course, this does not mean that our model is as good as theirs in, say, matching the responses to identified monetary shocks or in out-of-sample forecasting. It nevertheless indicates

that the inclusion of our mechanism in New Keynesian models does not interfere with their ability to capture the nominal side of the data and that our mechanism itself is robust to the introduction of realistic nominal rigidities.

5.4 On Demand-Driven Business Cycles

We conclude this section by exploring how our mechanism, viewed as a formalization of demand-driven business cycles, compares to that of the New Keynesian model.

To this goal, Table 4 reports the posterior odds of four models, starting from a uniform prior and estimating them on the real data only. The models differ on whether they assume flexible or sticky prices, and on whether they contain the confidence shock or not. We concentrate on the real data, not only because the flexible-price models are not designed to capture the nominal variables, but also because we wish to evaluate both kinds of models on the basis of the comovements of the real quantities. Once we drop the nominal data for this exercise, the nominal parameters of the sticky-price models are not well identified. We have thus chosen to fix these parameters at the values that obtained when the models were estimated on both real and nominal data. We nevertheless re-estimate the preference and technology parameters and the shock processes in order to give each model a fair chance to match the data on the real quantities.

Table 4: Posterior Odds of Model A vs Model B

| Model A → | sticky prices | |
Model B ↓	without	with
flex prices, without confidence	1.00	1.00
flex prices, with confidence	0.36	0.84
sticky prices, without confidence	–	0.90

Consider first the pair of models that abstract from the confidence shock. In this case, the sticky-price model wins: the posterior odds that the data are generated by that model are nearly 100%. But once the flexible-price model is augmented with the confidence shock, the odds of the sticky-price model fall below 50%, to 36%. By this metric, our mechanism appears to be more potent than the NK mechanism when the two are viewed in isolation. Finally, the sticky-price model that contains the confidence shock wins 90-10 over the the sticky-price model that excludes it. By this metric, the inclusion of our mechanism improves significantly the empirical performance of the NK model.

We interpret these results as follows. Insofar as we abstract from monetary phenomena, our approach emerges as a potent substitute for the NK formalization of demand-driven fluctuations. Perhaps more fruitfully, our approach can complement the NK framework by offering what, in our view, is a more appealing structural interpretation of the observed business cycles—one that attributes the "deficiency in aggregate demand" during a recession in part to a coordination failure and to lack of confidence.

6 Conclusion

By relying on the rational-expectations solution concept together with the auxiliary assumption that all agents share the same information about the aggregate state of the economy, standard macroeconomic models impose a rigid structure on how agents form beliefs about endogenous economic outcomes and how they coordinate their behavior. In this paper, we propose a certain relaxation of this structure and explore its quantitative implications.

In particular, we develop a method for augmenting macroeconomic models with a tractable form of higher-order belief dynamics. We argue that this method helps proxy for the effects of incomplete information and frictional coordination and can be used to accommodate a certain kind of waves of optimism and pessimism about the the short-term outlook of the economy. We document the quantitative importance of such waves within the context of RBC and New Keynesian models of both the textbook and the medium-scale variety.

We believe that our paper adds to the understanding of business-cycle phenomena along the following dimensions:

- It highlights the distinct role played by expectations of the short-run prospects of the economy, as opposed to expectations of productivity and growth in the medium to long run.

- It offers a parsimonious explanation of salient features of the macroeconomic data and does so in a manner that appears to outperform alternative narratives found in the literature.

- It offers a formalization of the notion of demand-driven fluctuations that is both conceptually and empirically distinct from the one found in the New Keynesian paradigm.

- It leads to a structural interpretation of the observed recessions that attributes a significant role to "coordination failures," "lack of confidence," or "market sentiment."

These findings naturally raise the question of where the variation in confidence comes from. Having attributed this variation to a confidence shock that is both exogenous to economic activity and orthogonal to other structural shocks, we can not offer a meaningful answer to this question.[30] Nevertheless, our analysis has revealed the potential importance of two previously overlooked forces, namely frictional coordination and belief waves regarding the short-term economic outlook, and so it can provide the impetus for future research on these subjects.

There is an emerging literature in this area. Ilut and Saijo (2016) and Angeletos and Lian (2016b) consider models that feature a similar kind of belief-driven wedges as the one found here, except that these wedges are allowed to covary with conventional structural shocks; this has the interesting implication that a drop in confidence may be triggered by an adverse financial shock, while a boost

---

[30]This limitation is not specific to this paper: any formal model must ultimately attribute the business cycle to some exogenous trigger, whether this is a technology shock, an uncertainty shock, or a sunpot.

in confidence may be accomplished by a fiscal stimulus. Huo and Takayama (2015b) obtain quantitative findings that are broadly consistent with ours while maintaining the common-prior assumption. Angeletos, Collard, and Dellas (2017) provide VAR-based evidence that the business cycle in the US data can be explained by a shock that has similar properties to the one we have accommodated in our theory. Levchenko and Pandalai-Nayar (2015) provide additional corroborating evidence in an international context.

Finally, it is worth iterating how the belief waves formalized and quantified in this paper compare to those found in the existing literature on news and noise shocks. Our confidence shock resembles the noise shocks of that literature in that both types of shocks are transitory. Yet, our mechanism captures a very different type of beliefs. In that literature, recessions are periods in which the agents expect the economy to do badly for a long time, and more so in the long run than in the short run; in our paper, they are periods in which the agents expect the economy to recover after a few years. Future work could shed further light on which formalization, and accompanying narrative, is more relevant empirically.

Appendix A. Estimated Models

In this appendix we fill in the details of the two models considered in Section 5; we next describe the estimation method, the assumed priors, and the obtained posteriors; we finally review a few additional findings that were omitted from the main text.

The details of the two models. As mentioned in the main text, the two models share the same backbone as our baseline model, but add a number of structural shocks along with certain forms of habit persistent in consumption and adjustment costs in investment, as in Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2007). To accommodate monopoly power and sticky prices, we also introduce product differentiation within each island. We finally assume that there exists a lump sum transfer that eliminates the effects of the markup rate in steady state.

Fix an island $i \in [0, 1]$. Index the firms in this island by $j \in [0, 1]$ and let y_{ijt} denote the output produced by firm j in period t. The composite output of the island is given by

$$y_{it} = \left(\int_0^1 y_{ijt}^{\frac{1}{1+\eta}} dj \right)^{1+\eta},$$

where $\eta > 0$ is a parameter that pins down the monopoly power. The aggregate quantity of the final good, on the other hand, is given by

$$Y_t = \left(\int_0^1 y_{it}^{1-\varrho} di \right)^{\frac{1}{1-\varrho}},$$

where $\varrho > 0$ is a parameter that ultimately governs the degree of strategic complementarity.

The technology is the same as before, so that the output of firm j in island i is

$$y_{ijt} = \exp(\zeta_t^a)(u_{ijt}k_{ijt})^\alpha n_{ijt}^{1-\alpha};$$

but now TFP is given by the sum of a permanent and a transitory component. More specifically,

$$\zeta_t^a = a_t^\tau + a_t^p,$$

where a_t^τ is the transitory component, modeled as an AR(1), and a_t^p is given by

$$a_t^p = a_{t-1}^p + a_{t-1}^n + \varepsilon_t^p$$

where ε_t^p is the unanticipated innovation and a_{t-1}^n captures all the TFP changes that agents anticipated in earlier periods. The latter is given by an AR(1) process of the form

$$a_t^n = \rho_n a_{t-1}^n + \varepsilon_t^n$$

where ε_t^n is the innovation to the anticipated component of TFP.[31] In line with our baseline model, the confidence shock is now modeled as a shock to higher-order beliefs of a_t^p.

To accommodate for a form of habit in consumption as well as discount-rate shocks, we let the per-period utility be as follows:

$$u(c_{it}, n_{it}; \zeta_t^c, C_{t-1}) = \exp(\zeta_t^c)\left(\log(c_{it} - bC_{t-1}) - \theta\frac{n_{it}^{1+\nu}}{1+\nu}\right)$$

where ζ_t^c is a transitory preference shock, modeled as an AR(1), $b \in (0,1)$ is a parameter that controls for the degree of habit persistence, and C_{t-1} denotes the aggregate consumption in the last period.[32]

To accommodate permanent shocks to the relative price of investment as well as transitory shocks to government spoending, we let the resource constraint of the island be given by the following:

$$c_{it} + \exp(\zeta_t^{ip})i_{it} + G_t + \exp(\zeta_t^{ip})\Psi(u_{it})k_{it} = p_{it}y_{it}$$

where ζ_t^{ip} measures the cost of investment, G_t is government spending, and $\exp(\zeta_t^{ip})\Psi(u_{it})$ is the cost of utilization per unit of capital. The latter is scaled by $\exp(\zeta_t^{ip})$ in order to transform the units of capital to units of the final good, and thereby also guaranteed a balanced-growth path. ζ_t^{ip} is modeled as a random walk: $\zeta_t^{ip} = \zeta_{t-1}^{ip} + \varepsilon_t^{ip}$. Literally taken, this represents an investment-specific *technology* shock. But since our estimations do not include data on the relative price of invest, this shock can readily be re-interpreted as a demand-side shock. The utilization-cost function satisfies $u\Psi''(u)/\Psi'(u) = \frac{\psi}{1-\psi}$, with $\psi \in (0,1)$. and government spending is given by $G_t = \bar{G}\exp(\tilde{G}_t)$, where

---

[31]We have experimented with alternative forms of diffusion, as well as with specifications such as $\zeta_t^n = \varepsilon_{t-4}^n$, and we have found very similar results.

[32]Note that we are assuming that habit is external. We experimented with internal habit, as in Christiano, Eichenbaum, and Evans (2005), and the results were virtually unaffected.

$\bar{G}$ is a constant and $\tilde{G}_t = \zeta_t^g + \frac{1}{1-\alpha}a_t^p - \frac{\alpha}{1-\alpha}\zeta_t^{ip}$. In this equation, ζ_t^g denotes a transitory shock, modeled as an AR(1), and the other terms are present in order to guarantee a balanced-growth path.

Finally, to accommodate adjustment costs to investment as well as transitory investment-specific shocks, we let the law of motion of capital on island i take the following form:

$$k_{it+1} = \exp(\zeta_t^{it})i_{it}\left(1 - \Phi\left(\frac{i_{it}}{i_{it-1}}\right)\right) + (1-\delta)k_{it}$$

We impose $\Phi'(\cdot) > 0$, $\Phi''(\cdot) > 0$, $\Phi(1) = \Phi'(1) = 0$, and $\Phi''(1) = \varphi$, so that φ parameterizes the curvature of the adjustment cost to investment. ζ_t^{it} is a temporary shock, modeled as an AR(1) and shifting the demand for investment, as in Justiniano, Primiceri, and Tambalotti (2010).

The above description completes the specification of the flexible-price model of Section 5. The sticky-price model is then obtained by embedding the Calvo friction and a Taylor rule form monetary policy. In particular, the probability that any given firm resets its price in any given period is given by $1 - \chi$, with $\chi \in (0, 1)$. As for the Taylor rule, the reaction to inflation is given by $\kappa_\pi > 1$, the reaction to the output gap is given by $\kappa_y > 0$, and the parameter that controls the degree of interest-rate smoothing is given by $\kappa_R \in (0, 1)$; see condition (31) below.

In the sticky-price model, the log-linear version of the set of the equations characterizing the general equilibrium of the economy is thus given by the following:

$$\mathbb{E}_{it}\left[\zeta_t^c + \nu\tilde{n}_{it}\right] = \zeta_t^c - \frac{1}{1-b}\tilde{c}_{it} + \frac{b}{1-b}\tilde{C}_{t-1} + \mathbb{E}_{it}\left[\tilde{s}_{it} + \varrho\tilde{Y}_t + (1-\varrho)\tilde{y}_{it} - \tilde{n}_{it}\right] \tag{21}$$

$$\mathbb{E}_{it}\left[\tilde{\lambda}_{it} + \tilde{q}_{it}\right] = \mathbb{E}_{it}\left[\begin{array}{l}\tilde{\lambda}_{it+1} + \beta(1-\delta)\tilde{q}_{it+1}+ \\ +(1-\beta(1-\delta))\left(\tilde{s}_{it+1} + \varrho\tilde{Y}_{t+1} + (1-\varrho)\tilde{y}_{it+1} - \tilde{u}_{it+1} - \tilde{k}_{it+1}\right)\end{array}\right] \tag{22}$$

$$\tilde{y}_{it} = a_t + \alpha(\tilde{u}_{it} + \tilde{k}_{it}) + (1-\alpha)\tilde{n}_{it} \tag{23}$$

$$Z_t + \frac{1}{1-\psi}\tilde{u}_{it} = \tilde{s}_{it} + \varrho\tilde{Y}_t + (1-\varrho)\tilde{y}_{it} - \tilde{k}_{it} \tag{24}$$

$$\varrho\tilde{Y}_t + (1-\varrho)\tilde{y}_{it} = s_c\tilde{c}_{it} + (1 - s_c - s_g)(\zeta_t^{ip} + \tilde{\imath}_{it}) + s_g\tilde{G}_t + \alpha\tilde{u}_{it} \tag{25}$$

$$\tilde{k}_{it+1} = \delta(\zeta_t^{it} + \tilde{\imath}_{it}) + (1-\delta)\tilde{k}_{it} \tag{26}$$

$$\tilde{q}_{it} = (1+\beta)\varphi\tilde{\imath}_{it} - \varphi\tilde{\imath}_{t-1} - \beta\varphi\mathbb{E}'_{it}\tilde{\imath}_{it+1} + \zeta_t^{ip} - \zeta_t^{it} \tag{27}$$

$$\tilde{\lambda}_{it} = \zeta_t^c - \frac{1}{1-b}\tilde{c}_{it} + \frac{b}{1-b}\tilde{C}_{t-1} \tag{28}$$

$$\tilde{R}_t = \zeta_t^c - (1+\nu)\tilde{n}_{it} - \tilde{s}_{it} - \varrho Y_t - (1-\varrho)y_{it} - \mathbb{E}'_{it}\left[\tilde{\lambda}_{it+1} - \tilde{\pi}_{it+1}\right] \tag{29}$$

$$\tilde{x}_{it} = s_c\tilde{c}_{it} + (1 - s_c - s_g)(\zeta_t^{ip} + \tilde{\imath}_{it}) + s_g\tilde{G}_t \tag{30}$$

$$\tilde{R}_t = \kappa_R\tilde{R}_{t-1} + (1-\kappa_R)\left(\kappa_\pi\tilde{\pi}_{it} + \kappa_y(\tilde{x}_{it} - \tilde{x}_{it}^F)\right) + \zeta_t^m \tag{31}$$

$$\chi(1 + \chi(1-\beta))\tilde{\pi}_{it} = (1-\chi)(1-\beta\chi)\tilde{s}_{it} + \beta\chi(1-\chi)\tilde{\Pi}_t + \beta\chi\mathbb{E}'_{it}\tilde{\pi}_{it+1} \tag{32}$$

where uppercases stand for aggregate variables, λ_{it} and s_{it} denote, respectively, the marginal utility of consumption and the realized markup in island i, $\tilde{\pi}_{it} \equiv \tilde{p}_{it} - \tilde{p}_{it-1}$ and $\tilde{\Pi}_t \equiv \tilde{P}_t - \tilde{P}_{t-1}$ denote, respectively, the local and the aggregate inflation rate, x_{it} denotes the measured of GDP on island i, X_{it}^F denotes the GDP that would be attained in a flexible price allocation, and s_c and s_g denote the steady-state ratios of consumption and government spending to output.

The interpretation of the above system is straightforward. Conditions (21) and (22) give, respectively, the consumption and investment decisions. Conditions (23) and (24) characterizes the equilibrium employment and utilization levels. Condition (25) gives the local resource constraint. Conditions (26) and (27) give the local law of motion of capital and the equilibrium price of capital. Condition (28) and (29) give the marginal utility of consumption and the optimal bond holdings decision. Condition (30) gives the measured aggregate GDP. Conditions (31) gives the Taylor rule for monetary policy. Finally, condition (32) gives the inflation rate in each island; aggregating this condition across islands gives our model's New Keynesian Phillips Curve. The only essential novelty in all the above is the presence of the subjective expectation operators in the conditions characterizing the local equilibrium outcomes of each island.

Finally, the flexible-price allocations are obtained by the same set of equations, modulo the following changes: we set $s_{it} = 0$, meaning that the realized markup is always equal to the optimal markup; we restate the Euler condition (29) in terms of the real interest rate; and we drop the nominal side of this system, namely conditions (31) and (32).

Estimation. As mentioned in the main text, we follow Christiano and Vigfusson (2003) and Sala (2015) and estimate the model using a Bayesian maximum likelihood technique in the frequency domain. This method amounts to maximizing the following posterior likelihood function:

$$\mathcal{L}(\theta|\mathcal{Y}_T) \propto f(\theta) \times L(\theta|\mathcal{Y}_T)$$

where $\mathcal{Y}_T$ denotes the set of data (for $t = 1 \ldots T$) used for estimation, θ is the vector of structural parameters to be estimated, $f(\theta)$ is the joint prior distribution of the structural parameters, and $L(\theta|\mathcal{Y}_t)$ is the likelihood of the model expressed in the frequency domain. Note that the log-linear solution of the model admits a state-space representation of the following form:

$$Y_t = M_y(\theta)X_t$$
$$X_{t+1} = M_x(\theta)X_t + M_e \varepsilon_{t+1}$$

Here, Y_t and X_t denote, respectively, the vector of observed variables and the underlying state vector of the model; ε is the vector of the exogenous structural shocks, drawn from a Normal distribution with mean zero and variance-covariance matrix $\Sigma(\theta)$; $M_y(\theta)$ and $M_x(\theta)$ are matrices whose elements are (non-linear) functions of the underlying structural parameters θ; and finally M_e is a selection matrix that describes how each of the structural shocks impacts on the state vector. As shown in Whittle (1951), Hannan (1970) and Harvey (1991), the likelihood function is asymptotically given by

$$\log(L(\theta|\mathcal{Y}_T)) \propto -\frac{1}{2}\sum_{j=1}^{T} \gamma_j \left(\log(\det S_Y(\omega_j, \theta) + \text{tr}\left(S_Y(\omega_j, \theta)^{-1}I_Y(\omega_j)\right)\right)$$

where $\omega_j = 2\pi j/T$, $j = 1 \ldots T$ and where $I_Y(\omega_j)$ denotes the periodogram of $\mathcal{Y}_T$ evaluated at frequency ω_j. $S_Y(\omega, \theta)$ is the model spectral density of the vector Y_t, given by

$$S_Y(\omega, \theta) = \frac{1}{2\pi}M_y(\theta)(I - M_x(\theta)e^{-i\omega})^{-1}M_e\Sigma(\theta)M_e'(I - M_x(\theta)'e^{i\omega})^{-1}M_y(\theta)''$$

Following Christiano and Vigfusson (2003) and Sala (2015), we include a weight γ_j in the computation of the likelihood in order to select the desirable frequencies: this weight is 1 when the frequency falls between 6 and 32 quarters, and 0 otherwise.

Priors. The following parameters are estimated in both models: the inverse labor supply elasticity, ν; the capital share, α; the utilization elasticity parameter, ψ; the habit persistence parameter, b; the parameter governing the size of investment adjustment costs, φ; and the standard deviations and persistences of all the structural shocks. In the sticky-price model, the Calvo parameter, χ, and parameters of the Taylor rule, κ_R, κ_π, and κ_y, are also estimated. The priors used for all these parameters are reported in Table 8 in Online Appendix O.3 and are broadly consistent with those used in the DSGE literature. The prior for the confidence shock was set in line with the other shocks. Finally, the following parameters are fixed: the discount factor, β, is 0.99; the depreciation rate, δ, is 0.025; the parameter, η, is such that the monopoly markup is 15%; and the parameter ϱ is 0.75 for the reasons explained in the main text.

Posteriors. Posterior distributions were obtained with the MCMC algorithm, with an acceptance rate of 37%. We generated 2 chains of 200,000 observations each. The posteriors for all the parameters are reported in the last four columns of Table 8. The posteriors for the preference, technology, and monetary parameters are broadly consistent with other estimates in the literature.

IRFs and Variance/Covariance Decompositions. The IRFs of our estimated models with respect to all the structural shocks are delegated to Online Appendix O.3: see Figures 8–9 therein. With the exception of the confidence shock, which is novel, the IRFs to all the other shocks are comparable to those found in the literature.

The estimated contribution of the shocks to, respectively, the variances and the co-variances of the key variables at business-cycle frequencies is reported in the same appendix, in Tables 9 and 10. For comparison purposes, we also include the estimated contributions that obtain in the variants of the models that remove the confidence shock. Three findings are worth mentioning.

First, unlike the case of the confidence shock, the variance/covariance contributions of some of the other shocks changes significantly as we move from the flexible-price to the sticky-price model.

Second, in the models that assume away the confidence shocks, the combination of permanent and transitory investment shocks emerge as the main driver of the business cycle. This is consistent with existing findings in the DSGE literature (e.g., Justiniano, Primiceri, and Tambalotti, 2010) and confirms that, apart from the inclusion of the confidence shock, our exercises are quite typical.

Finally, in all models, neither the investment-specific shocks, nor the news or discount-rate shocks are able to contribute to a positive covariation between all of the key real quantities (output, consumption, investment, hours) at the same time. This illustrates, once again, the superior ability of our mechanism to generate the right kind of comovement patterns.

References

Abraham, Arpad and Kirk White. 2006. "The Dynamics of Plant-Level Productivity in U.S. Manufacturing." Working Paper 06-20, Center for Economic Studies, U.S. Census Bureau.

Akerlof, Georges and Robert Shiller. 2009. *Animal Spirits: How Human Psychology Drives the Economy, and Why it Matters for Global Capitalism*. Princeton, NJ: Princeton University Press.

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas. 2017. "An Anatomy of the Business Cycle." miméo, MIT.

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

Angeletos, George-Marios and Chen Lian. 2016a. "Incomplete Information in Macroeconomics: Accommodating Frictions in Coordination." In *Handbook of Macroeconomics, Vol. 2*, vol. 2, edited by John B. Taylor and Harald Uhlig, chap. 14. 1065–1240.

———. 2016b. "A (Real) Theory of Keynesian Multipliers." *MIT mimeo* .

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

Barsky, Robert B. and Eric R. Sims. 2011. "News Shocks and Business Cycles." *Journal of Monetary Economics* 58 (3):273–289.

———. 2012. "Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence." *American Economic Review* 102 (4):1343–77.

Beaudry, Paul, Dana Galizia, and Franck Portier. 2015. "Reviving the Limit Cycle View of Macroeconomic Fluctuations." Working Paper 21241, National Bureau of Economic Research.

Beaudry, Paul and Franck Portier. 2006. "Stock Prices, News, and Economic Fluctuations." *American Economic Review* 96 (4):1293–1307.

———. 2013. "Understanding Noninflationary Demand-Driven Business Cycles." In *NBER Macroeconomics Annual 2013, Volume 28*. 69–130.

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

Bhandari, Anmol, Jaroslav Borovicka, and Paul Ho. 2016. "Identifying Ambiguity Shocks in Business Cycle Models Using Survey Data." Working Paper 22225, National Bureau of Economic Research.

Blanchard, Olivier Jean and Danny Quah. 1989. "The Dynamic Effects of Aggregate Demand and Supply Disturbances." *American Economic Review* 79 (4):655–73.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo. 2016. "Understanding Booms and Busts in Housing Markets." *Journal of Political Economy* 124 (4):1088–1147.

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 Charles L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113 (1):1–45.

Christiano, Lawrence J., Martin Eichenbaum, and Mathias Trabandt. 2015. "Understanding the Great Recession." *American Economic Journal: Macroeconomics* 7 (1):110–67.

Christiano, Lawrence J. and Robert J. Vigfusson. 2003. "Maximum likelihood in the Frequency Domain: the Importance of Time-to-Plan." *Journal of Monetary Economics* 50 (4):789–815.

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

Fernald, John G. 2014. "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity." Working Paper Series 2012-19, Federal Reserve Bank of San Francisco.

Fisher, Jonas D. M. 2006. "The Dynamic Effects of Neutral and Investment-Specific Technology Shocks." *Journal of Political Economy* 114 (3):413–451.

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

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

Hannan, Edward J. 1970. *Multiple Time Series*. Wiley Series in Probability and Statistics. Wiley, John & Sons.

Hansen, Lars Peter and Thomas J. Sargent. 1993. "Seasonality and Approximation Errors in Rational Expectations Models." *Journal of Econometrics* 55 (1-2):21–55.

———. 2007. *Robustness*. Princeton, NJ: Princeton University Press.

———. 2012. "Three types of ambiguity." *Journal of Monetary Economics* 59 (5):422–445.

Harvey, Andrew C. 1991. *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge, UK: Cambridge University Press.

Huo, Zen and Naoki Takayama. 2015a. "Rational Expectations Models with Higher Order Beliefs." miméo, Yale University.

———. 2015b. "Higher Order Beliefs, Confidence, and Business Cycles." miméo, Yale University.

Ilut, Cosmin and Hikaru Saijo. 2016. "Learning, Confidence and Business Cycle." Working Paper 22958, National Bureau of Economic Research.

Jaimovich, Nir and Sergio Rebelo. 2009. "Can News about the Future Drive the Business Cycle?" *American Economic Review* 99 (4):1097–1118.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti. 2010. "Investment Shocks and Business Cycles." *Journal of Monetary Economics* 57 (2):132–145.

Levchenko, Andrei A. and Nitya Pandalai-Nayar. 2015. "TFP, News, and "Sentiments": The International Transmission of Business Cycles." Working Paper 21010, National Bureau of Economic Research.

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

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

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

Morris, Stephen and Hyun Song Shin. 2001. "Rethinking Multiple Equilibria in Macroeconomic Modeling." In *NBER Macroeconomics Annual 2000, Volume 15*. National Bureau of Economic Research, Inc, 139–182.

———. 2002. "Social Value of Public Information." *American Economic Review* 92 (5):1521–1534.

Nimark, Kristoffer. 2017. "Dynamic Higher Order Expectations." miméo, Cornell University.

Pei, Guangyu. 2017. "Ambiguity, Pessimism and Economic Fluctuations." mimeo, University of Zurich.

Phelps, Edmund. 1971. "Inflation Expectations and Economic Theory." In *Inflation and the Canadian Experience*. Kingston, Ontario: Industrial Relation Centre, Queen's University, 31–47.

Rondina, Giacomo and Todd B. Walker. 2014. "Dispersed Information and Confounding Dynamics." miméo, University of Colorado, Boulder.

Sala, Luca. 2015. "DSGE Models in the Frequency Domains." *Journal of Applied Econometrics* 30 (2):219–240.

Sims, Christopher A. 1993. "Rational Expectations Modeling with Seasonally Adjusted Data." *Journal of Econometrics* 55 (1-2):9–19.

———. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50 (3):665–690.

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

Townsend, Robert M. 1983. "Forecasting the Forecasts of Others." *Journal of Political Economy* 91 (4):546–587.

Whittle, Peter. 1951. *Hypothesis Testing in Time Series Analysis*. Almqvist & Wiksells.

Woodford, Michael. 2002. "Imperfect Common Knowledge and the Effects of Monetary Policy." In *Knowledge, Information, and Expectations in Modern Macroeconomics: in Honor of Edmund S. Phelps*. Princeton University Press.

———. 2010. "Robustly Optimal Monetary Policy with Near-Rational Expectations." *American Economic Review* 100 (1):274–303.

Quantifying Confidence:
Online Appendices

George-Marios Angeletos   Fabrice Collard   Harris Dellas

December 3, 2017

O.1   Additional Material for Sections 3 and 4

In this appendix, we provide few results that complement the analysis in Section 3 and 4.

Heterogeneous vs Common Priors (continued)

In Subsection 3.3 of the main text, we established an observational equivalence between a special case of our heterogeneous-prior model and a common-prior variant featuring idiosyncratic uncertainty. We now elaborate on the bounds that this mapping can impose on the magnitude and the persistence of the ξ_t shock in our setting.

Suppose that we had data that allowed the estimation of the AR(1) process described in condition (18). Suppose next that we possessed information on the value for $\tilde{\sigma}_a$, perhaps from micro-economic observations. This information could be used in Proposition 2 to derive the bounds on (φ, ψ) and then, using Corollary 1, to get bounds on (ρ_ξ, σ_ξ).

Figure 7 depicts these bounds. To construct this figure, we let $\nu = 0.5$ and $\tilde{\sigma}_a = 0.2$. The latter value is based on the observation that $\tilde{\sigma}_a$ determines the uncertainty that islands face about their terms of trade (demand for their products), and may thus be proxied by the idiosyncratic risk that the typical firm faces about its productivity and sales.[33] In the left panel of the figure, we plot the set of the (φ, ψ) pairs that satisfy the bounds in Proposition 2, under the assumed value for $\tilde{\sigma}_a$. Using Corollary 1, we can translate this set into corresponding values for (ρ_ξ, σ_ξ).

In the right panel of the figure, we plot a more useful transformation of this set: instead of measuring σ_ξ on the vertical axis, we measure the corresponding value of σ_y, where σ_y henceforth stands for the standard deviation of the business-cycle component of output (i.e., of output bandpass filtered over 6-32 quarters) that is accounted by the confidence shock. Finally, the dot indicates the values of φ (in the left panel) and of σ_y (in the right panel) that obtain when we fix $\rho_\xi = 0.75$ and calibrate the volatilities of the confidence shock and the technology shock in the model so as to match the volatilities of aggregate output and employment in the data. The figure shows that under a plausible value for $\tilde{\sigma}_a$, the range of values for σ_ξ that would be consistent with the restrictions imposed by a common-prior specification is very large.

---

[33] Empirical estimates of the volatility of firm-level productivity suggest setting $\tilde{\sigma}_a$ between 0.2 to 0.43 (Abraham and White, 2006, Foster, Haltiwanger, and Syverson, 2008). In a similar setting as ours, Huo and Takayama (2015b) use a value of 0.14.
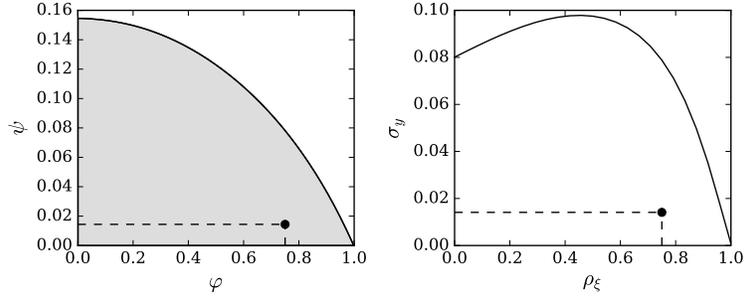
Figure 7: The bounds on persistence and volatility

Note that the relevant bound remains large even for lower values of $\tilde{\sigma}_a$, say 4%. Such a value would not appear implausibly large even if we confined first-order uncertainty to concern aggregate fundamentals. For instance, this value is only about twice as large as the standard deviation of the quarterly innovations in the aggregate Solow residual. Furthermore, as the behavior in the richer models used in the quantitative exercises in this paper is forward-looking, it seems more appropriate to think about a present-value measure of the uncertainty in fundamentals, as opposed to merely the quarter-by-quarter changes. Therefore, even though we can not extend the results of this subsection to such richer models, we feel confident that our quantitative findings are consistent with realistic common-prior models. The recent work of Huo and Takayama (2015b) seems to corroborate this conjecture. That said, there is no reason to view our approach *exclusively* as a proxy for incomplete information and rational confusion.

The Confidence Shock in the Baseline New Keynesian Model

In Section 4 of the main text, we compared the comovement patterns generated by the confidence shock to those of a few alternative shocks within the context fo the baseline RBC model. We now extend the comparison to the baseline New Keynesian model. The latter is obtained from the former by adding monopoly power, sticky prices, and a Taylor rule for monetary policy.

Table 5 revisits the exercise conducted in Table 1. The preferences, the technology and the confidence shock remain as before; the monopoly distortion is offset by a subsidy; the Calvo parameter is set to 0.75; and the Taylor rule is specified as $R_t = \phi_\pi \pi_t$ with $\phi_\pi = 1.5$.

The following key findings emerge. First, the good and superior to other shocks empirical performance of the confidence shock survives as we move from the RBC model to the New Keynesian model. Second, with the exception of the monetary shock, none of the competing shocks is able to generate realistic comovements patterns in the relevant quantities. Finally, the similarity between the real effects of the confidence shock and those of the monetary shock provide further justification for our claim that the confidence plays similar role in the RBC framework as demand shocks

2

Table 5: Conditional Comovements (6-32 quarters)

	Filtering		Our Mechanism		Alternative Mechanisms				
	(a)	(b)	(c)	(d)	I shock	C shock	News	E shock	M shock
σ_n/σ_y	0.87	1.07	1.43	1.43	1.44	1.44	0.87	0.70	1.44
σ_c/σ_y	0.56	0.52	0.25	0.22	0.44	1.01	0.16	0.19	0.13
σ_i/σ_y	3.54	3.65	3.92	4.10	6.09	8.26	4.80	4.30	4.52
$\sigma_{y/n}/\sigma_y$	0.40	0.63	0.44	0.44	0.47	0.49	0.18	0.33	0.45
$corr(c,y)$	0.86	0.85	0.85	0.76	-0.83	-0.94	-0.16	0.60	0.38
$corr(i,y)$	0.94	0.95	0.99	0.99	0.99	0.99	0.99	0.99	0.99
$corr(n,y)$	0.91	0.82	0.99	1.00	0.99	0.98	0.99	0.99	1.00
$corr(c,n)$	0.86	0.75	0.81	0.70	-0.90	-0.99	-0.25	0.47	0.31
$corr(i,n)$	0.85	0.81	0.99	1.00	1.00	1.00	1.00	1.00	1.00
$corr(c,i)$	0.75	0.79	0.78	0.67	-0.90	-0.98	-0.27	0.50	0.29
$corr(y,y/n)$	-0.03	0.21	-0.96	-0.96	-0.90	-0.86	0.76	0.94	-0.97
$corr(n,y/n)$	-0.43	-0.40	-0.98	-0.98	-0.95	-0.93	0.67	0.87	-0.99
σ_π/σ_y	0.16	0.42	–	0.07	0.22	0.04	0.07	0.03	0.10
σ_R/σ_y	0.24	0.72	–	0.10	0.34	0.06	0.10	0.05	0.02
$corr(y,\pi)$	0.21	-0.90	–	0.96	0.99	0.42	0.37	0.84	0.99
$corr(y,R)$	0.38	-0.81	–	0.96	0.99	0.42	0.37	0.84	-0.38

Note: Columns (a) and (b) refer to the residuals that obtain, respectively, from the projection of the data on current and past TFP and from the removal of the technology shock identified in the same was as in Galí (1999). Column (c) refers to the predictions of our baseline model and column (d) to those of its NK variant. All other columns refer to alternative NK models.

Belief-Driven Wedges

In this section we derive the predictions of our theories about the wedges. We consider both the overall wedges between the marginal rates of substitution and the corresponding marginal rates of transformation, and their decomposition in household- and firm-side wedges.

Let us fill in the details. First, denote with $MRSN_t \equiv \nu N_t + \gamma C_t$ the measured marginal rate of intra-temporal substitution between leisure and consumption; with $MRSC_{t,t+1} \equiv \gamma(C_{t+1} - C_t)$ the measured marginal rate of inter-temporal substitution in consumption; with $MPL_t \equiv Y_t - N_t$ the measured marginal product of labor; and with $MPK_t = Y_t - K_t$ the measured marginal product of capital. Next, define the wedges τ_t^{nh}, τ_t^{kh}, τ_t^{nf}, and τ_t^{kf} so that the following conditions hold:

$$MRSN_t = w_t - \tau_t^{nh} \qquad \mathbb{E}_t[MRSC_{t,t+1}] = (1 - \beta(1-\delta))(R_t - \tau_t^{kh}) \qquad (33)$$

$$MPL_t = w_t + \tau_t^{nf} \qquad \mathbb{E}_t[MPK_{t+1}] = R_t + \tau_t^{kf}. \qquad (34)$$

This means that τ_t^{nh} and τ_t^{kh} can be interpreted as taxes payed by the household on labor income and on the return to savings, while τ_t^{nf} and τ_t^{kf} can be interpreted as taxes payed by the firm on the use of labor and capital. We finally measure the total labor wedge by $\tau^n \equiv \tau_t^{nh} + \tau^{nf}$ and the total capital wedge by $\tau^k \equiv \tau_t^{kh} + \tau^{kf}$.

When the data is generated by the plain-vanilla RBC model, all the wedges are zero. At the other extreme, the wedges can be arbitrary stochastic processes if the data is generated by a medium-scale model that lets each of the optimality conditions of the RBC model be perturbed by a different shock. Our model is in between these two extremes, arguably closer to the plain-vanilla RBC model than to DSGE models such as Smets and Wouters (2007): the wedges differ from zero but they are all linear functions of the underlying confidence shock.

Furthermore, as shown next, $\xi_t > 0$ maps to $\tau_t^{nh} > 0$, $\tau_t^{kh} > 0$, $\tau_t^{nf} < 0$, and $\tau_t^{kf} < 0$. That is, whenever there is a boost in confidence, it is *as if* the household faces a positive tax on its supply of labor and savings, while the firm faces a positive subsidy on its use of labor and capital services. The first property reflects the excessive optimism that the households have about their income during a confidence-driven boom; the second property reflects the excessive optimism that the firms have about the demand for their product and their terms of trade. Finally, the combination of these forces gives rise to a procylical labor wedge and a countercylical capital wedge, in line with the US data.

The labor wedge on the household side. Consider first τ_t^{nh}, which is is defined as the equivalent of a labor-income tax faced by the as-if representative household:

$$\tau_t^{nh} \equiv w_t - MRSN_t = w_t - (C_t + \nu N_t)$$

In the equilibrium of our model, the household of every island i equates the local wage to the local expectation of its marginal rate of substitution between consumption and leisure:

$$w_{i,t} = \mathbb{E}_{it}[c_{i,t}] - \nu n_{i,t}.$$

In addition, the realized outcomes satisfy $w_{i,t} = w_t$, $n_{i,t} = N_t$, and $c_{i,t} = C_t$ for all i. It follows that

$$\tau_t^{nh} = \mathbb{E}_{it}[c_{i,t}] - C_t = \mathbb{E}_{it}[c_{i,t}] - c_{it} \quad \forall i,$$

which reveals that τ_t^{nh} captures the excessive optimism (during a boom) or pessimism (during a recession) of the households about their own consumption. Condition (13) in the main text, together with the fact that $k_{it} = K_t$ for all i, implies that $c_{it} = \Gamma_K^c K_t + \Gamma_z^c z_{it} + \Gamma_{\bar{z}}^c \bar{z}_t + \Gamma_a^c A_t + \Gamma_\xi^c \xi_t$ and therefore $\mathbb{E}_{it}[c_{it}] = \Gamma_K^c K_t + (\Gamma_z^c + \Gamma_{\bar{z}}^c + \Gamma_a^c) A_t + (\Gamma_\xi^c + \Gamma_{\bar{z}}^c) \xi_t$. Realized consumption, on the other hand, is given by $C_t = \Gamma_K^c K_t + (\Gamma_z^c + \Gamma_{\bar{z}}^c + \Gamma_a^c) A_t + \Gamma_\xi^c \xi_t$. Combining, we infer that

$$\tau_t^{nh} = \Gamma_{\bar{z}}^c \xi_t$$

The adopted parameterization implies $\tau_t^{nh} = 0.0152 \xi_t$.

The labor wedge on the firm side. Consider next τ_t^{nf}, which is defined as the equivalent of a payroll tax faced by the as-if representative firm:

$$\tau_t^{nf} \equiv MPL_t - w_t = (Y_t - N_t) - w_t$$

In the equilibrium of our model, the firm of every island i equates the local wage to the local expectation of the marginal revenue product of labor:

$$w_{i,t} = \mathbb{E}_{it}[MRPL_{it}] = \mathbb{E}_{it}[p_{it} + y_{it} - n_{it}] = \mathbb{E}_{it}[Y_t] - n_{it} \quad \forall i$$

In addition, the realized outcomes satisfy $w_{it} = w_t$ and $n_{it} = N_t$ for all i. It follows that

$$\tau_t^{nf} = Y_t - \mathbb{E}_{it}[Y_t] \quad \forall i,$$

which reveals that τ_t^{nf} captures the excessive optimism or pessimism of the firms about aggregate income and the resulting demand for the local good. Using conditions (12) and (15) from the main text along with the production function, we have that $Y_t = \Gamma_K^y K_t + \Gamma_a^y A_t + \Gamma_{\bar{z}}^y \bar{z}_t + \Gamma_{\xi}^y \xi_t$ and therefore $\mathbb{E}_{it}[Y_t] = \Gamma_K^y K_t + (\Gamma_a^y + \Gamma_{\bar{z}}^y)A_t + (\Gamma_{\bar{z}}^y + \Gamma_{\xi}^y)\xi_t$. It follows that

$$\tau_t^{nf} \equiv -\Gamma_{\bar{z}}^y \xi_t$$

For our parameterization, we have $\tau_t^{nf} = -0.2548\xi_t$.

The capital wedge on the firm side. Consider now τ_t^{kf}, which is defined as the equivalent of an investment tax faced by the as-if representative firm:

$$\tau_t^{kf} \equiv \mathbb{E}_t[MPK_{t+1}] - R_t = \mathbb{E}_t[Y_{t+1}] - K_{t+1} - R_t,$$

where $\mathbb{E}_t$ is the rational (or objective) expectation operator. In the equilibrium of our model,

$$R_t = \mathbb{E}'_{it}[MRPK_{i,t+1}] = \mathbb{E}'_{it}[p_{i,t+1} + y_{i,t+1} - k_{i,t+1}] = \mathbb{E}'_{it}[Y_{t+1}] - k_{i,t+1} \quad \forall i,$$

where $\mathbb{E}_{i,t}$ is the subjective expectation operator in the morning of period t. It follows that

$$\tau_t^{kf} = \mathbb{E}_t[Y_{t+1}] - \mathbb{E}'_{it}[Y_{t+1}] \quad \forall i,$$

which reveals that τ_t^{kf} captures the excessive optimism or pessimism of the firms about aggregate income and demand next period. Using similar steps as before, we can show that

$$\tau_t^{kf} = -\Gamma_{\bar{z}}^y \rho \xi_t$$

where $\Gamma_{\bar{z}}^y$ is the elasticity of the realized income of each island with respect to the realized average signal. For our calibration, we have $\tau_t^{kf} = -0.1911\xi_t$.

The savings wedge on the household side. Consider τ_t^{kh}, which is defined as the tax on the returns to savings faced by the as-if representative household:

$$\tau_t^{kh} \equiv R_t - \frac{1}{1-\beta(1-\delta)}\mathbb{E}_t[MRSC_{t,t+1}] = R_t + \frac{\gamma}{1-\beta(1-\delta)}\mathbb{E}_t[C_{t+1} - C_t]$$

In the equilibrium of our model,

$$R_t = \frac{1}{1-\beta(1-\delta)}\mathbb{E}'_{it}[MRSC_{i,t,t+1}] = \frac{\gamma}{1-\beta(1-\delta)}\mathbb{E}'_{it}[c_{i,t+1} - c_t],$$

where $\mathbb{E}'_{it}$ is the subjective operator in the afternoon of period t. It follows that

$$\tau_t^{kh} = \frac{\gamma}{1-\beta(1-\delta)}\left(\mathbb{E}'_{it}[c_{it+1}] - \mathbb{E}_t[C_{t+1}]\right) = \frac{\gamma}{1-\beta(1-\delta)}\left(\mathbb{E}'_{it}[c_{it+1}] - \mathbb{E}_t[c_{i,t+1}]\right),$$

which reveals that τ_t^{kh} captures the excessive optimism or pessimism of the households about their future consumption. From the policy rules for individual and aggregate consumption:

$$\mathbb{E}_{it}'[c_{it+1}] = \Gamma_K^c K_{t+1} + (\Gamma_z^c + \Gamma_{\bar{z}}^c + \Gamma_a^c)A_t + (\Gamma_{\bar{z}}^c + \Gamma_{\xi}^c)\rho\xi_t$$

$$\mathbb{E}_t[C_{t+1}] = \Gamma_K^c K_{t+1} + (\Gamma_z^c + \Gamma_{\bar{z}}^c + \Gamma_a^c)A_t + \Gamma_{\xi}^c\rho\xi_t$$

Combining, we infer that

$$\tau_t^{kh} = \frac{\Gamma_{\bar{z}}^c\rho}{1 - \beta(1-\delta)}\xi_t$$

For our parameterization, we have $\tau_t^{kh} = 0.3277\xi_t$.

The total wedges in the model. Combining the above results, we conclude that the total labor wedge in the calibrated version of our baseline model is given by $\tau_t^n = \tau_t^{nh} + \tau_t^{nf} = -0.2396\xi_t$, whereas the total capital wedge is given by $\tau_t^k = \tau_t^{kh} + \tau_t^{kf} = 0.1366\xi_t$. That is, the labor wedge is negatively correlated with the confidence shock, and therefore countercyclical, while the capital wedge is positively correlated with the confidence shock, and therefore procyclical.

In the main text we claimed that both of these predictions are driven by the fact that the ξ_t shock shifts the perceptions of short-run returns without moving much the perceptions of permanent income. Let us now explain why this is the case. As noted above, our model predicts that the wedges for firms and households move in opposite directions. Furthermore, the procyclicality of τ_t^{nh} is tied to the effect of the confidence shock on perceived permanent income, while the countercyclicality of $\tau_t^n f$ is tied to the effect on the perceived marginal return to labor. For the reasons already explained, the latter effect dominates the former. Consequently, the overall labor wedge, τ_t^n, is predicted to be countercyclical. The opposite is true for the capital wedge, τ_t^k. To see why, note first that the Euler condition equates expected consumption growth with a quantity that is equal to unity plus the expected return to capital. Note next that, while the variation in τ_t^{kf} is of similar magnitude to the variation in τ_t^{nf}, it represents a small component in the aforementioned quantity, and is therefore overwhelmed by the variation in τ_t^{kh}, which captures the household's optimism and pessimism about future consumption. It follows τ_t^k shares the cyclical properties of τ_t^{kh}, that is, the total capital wedge is procyclical.

Estimation of wedges in the data. We now turn attention to the estimation of the wedges US data. This is done in a similar fashion as in Chari, Kehoe, and McGrattan (2007).

The estimation is based on the baseline RBC model, augmented with ad hoc stochastic processes for the following four wedges: an efficiency wedge, τ_t^e, a labor wedge, τ_t^n, a capital wedge, τ_t^k, and a resource wedge τ_t^g. Accordingly, the system to be estimated is the following:

$$\nu N_t + C_t = Y_t - N_t - \tau_t^n \tag{35}$$

$$\mathbb{E}_t[C_{t+1}] - C_t = (1 - \beta(1-\delta))(\mathbb{E}_t[Y_{t+1} - K_{t+1}] - \tau_t^k) \tag{36}$$

$$Y_t + (1-\delta)K_t = C_t + K_{t+1} + \tau_t^g \tag{37}$$

$$Y_t = \tau_t^e + \alpha K_t + (1-\alpha)N_t \tag{38}$$

We set the structural parameters ν, α, β and δ to the values chosen in our baseline calibration. As in Chari, Kehoe, and McGrattan (2007), we assume that the vector $T_t = (\tau_t^e, \tau_t^n, \tau_t^k, \tau_t^g)'$ follows a VAR(1) process of the form

$$T_t = \Phi T_{t-1} + \mathcal{E}_t$$

where Φ is a matrix, $\mathcal{E}_t = (\varepsilon_t^e, \varepsilon_t^n, \varepsilon_t^k, \varepsilon_t^g)'$ is normally distributed with $\mathbb{E}(\mathcal{E}_t) = 0$ and $\mathbb{E}(\mathcal{E}_t \mathcal{E}_t') = \Omega\Omega'$, and Ω is a lower-triangular matrix. We finally estimate the matrices Φ and Ω using data on GDP, investment, hours, and the difference between GDP and the sum of investment and consumption, over the period 1960Q1-2007Q4. The estimation yields

$$\Phi = \begin{pmatrix} 0.6537 & 0.1184 & 0.2268 & 0.0049 \\ -0.2487 & 1.0716 & 0.1605 & 0.0089 \\ -0.2808 & 0.0883 & 1.1620 & 0.0068 \\ 0.2017 & -0.1390 & -0.1741 & 0.9829 \end{pmatrix} \quad \text{and} \quad \Omega = \begin{pmatrix} 0.6148 & 0.0000 & 0.0000 & 0.0000 \\ 0.2580 & 0.8828 & 0.0000 & 0.0000 \\ 0.6261 & -0.3505 & 0.1793 & 0.0000 \\ 0.2492 & 0.2278 & 0.4964 & 1.5210 \end{pmatrix},$$

and results to the moments reported in Table 6. We thus see that the labor wedge is countercyclical and the capital wedge procyclical, just as predicted by our theory.

Table 6: Wedges in the Data

	Efficiency	Labor	Capital
Standard Deviation	0.86	1.40	1.04
Correlation with Output	0.78	-0.57	0.91

O.2 Data

In this Appendix we describe the data we use in this paper to obtain the various business-cycle moments and to estimate the models considered in Section 5.

Table 7 summarizes the data, all of which is from FRED, the Economic Database of the Federal Reserve Bank of Saint-Louis. GDP, Y, is measured by the seasonally adjusted GDP. Consumption, C, is measured by the sum of personal consumption expenditures in nondurables goods (CND) and services (CS). Investment, I, is measured by the sum of personal consumption expenditures on durables goods (CD), fixed private investment (FPI) and changes in inventories (DI). Government Spending, G, is measured by government consumption expenditures (GCE). Hours worked, N, are measured by hours of all persons in the non-farm business sector. Labor productivity, Y/N, is measured by real output per hour of all persons in the non-farm business sector. The inflation rate, π, is the log-change in the implicit GDP deflator. The nominal interest rate, R, is the effective federal funds rate measured on a quarterly basis. Given that the effective federal funds rate is available at the monthly frequency, we use the average over the quarter (denoted FEDFUNDS). Finally, when relevant, Total Factor Productivity (TFP) is measured as in Fernald (2014), which adjusts for utilization.

The sample ranges from the first quarter of 1960 to the last quarter of 2007. We dropped the post-2007 data because the models we study are not to designed to deal with the financial phenomena that appear to have play a more crucial role in the recent recession as opposed to earlier times. All quantities are expressed in real, per capita terms—that is, deflated by the implicit GDP deflator (GDPDEF) and by the civilian non-institutional population (CNP16OV). Because the latter is reported monthly, we used the last month of each quarter as the quarterly observation.

Table 7: Description of the Data

Description of the Data

Data	Formula
GDP	`Y=GDP/(GDPDEF×CNP16OV)`
Consumption	`C=(CND+CS)/(GDPDEF×CNP16OV)`
Investment	`I=(CD+FPI+DI)/(GDPDEF×CNP16OV)`
Government Spending	`G=GCE/(GDPDEF×CNP16OV)`
Hours Worked	`H=HOANBS/CNP16OV`
Labor Productivity	`GDP/H`
Inflation Rate	π=`log(GDPDEF)-log(GDPDEF)`$_{-1}$
Nominal Interest Rate	`R=FEDFUNDS/4`

Mnemonic	Source
GDP	`http://research.stlouisfed.org/fred2/series/GDP`
CND	`http://research.stlouisfed.org/fred2/series/PCND`
CD	`http://research.stlouisfed.org/fred2/series/PCEDG`
CS	`http://research.stlouisfed.org/fred2/series/PCESV`
FPI	`http://research.stlouisfed.org/fred2/series/FPI`
DI	`http://research.stlouisfed.org/fred2/series/CBI`
GCE	`http://research.stlouisfed.org/fred2/series/GCE`
HOANBS	`http://research.stlouisfed.org/fred2/series/HOANBS`
GDPDEF	`http://research.stlouisfed.org/fred2/series/GDPDEF`
FEDFUNDS	`http://research.stlouisfed.org/fred2/series/FEDFUNDS`
CNP16OV	`http://research.stlouisfed.org/fred2/series/CNP16OV`

O.3 Additional Material for Section 5

This appendix contains additional material regarding the two estimated models in Section 5. Table 8 reports the priors and the posteriors of the estimated parameters. Figures 8–9 report the IRFs of our estimated models with respect to all the structural shocks. Tables 9 and 10 report and the estimated contribution of the shocks to, respectively, the variances and the co-variances of the key variables at business-cycle frequencies. The confidence shock is omitted here, because its contributions were reported in the main text.

Table 8: Estimated Parameters

| | Priors | | | Posteriors | | | |
| | | | | Flexible Price Model | | Sticky Price Model | |
	Shape	Mean	Std. Dev.	Median	90%HPDI	Median	90%HPDI
ν	G	0.500	0.200	0.456	[0.226,0.814]	0.282	[0.161,0.429]
α	B	0.300	0.150	0.261	[0.234,0.286]	0.255	[0.229,0.280]
ψ	B	0.500	0.200	0.576	[0.255,0.856]	0.500	[0.315,0.708]
φ	G	2.000	1.000	3.370	[2.026,5.346]	3.312	[1.917,5.394]
b	B	0.500	0.200	0.860	[0.809,0.899]	0.758	[0.649,0.836]
χ	B	0.660	0.100	–	–	0.732	[0.673,0.782]
κ_R	B	0.600	0.200	–	–	0.198	[0.072,0.371]
κ_π	N	1.700	0.300	–	–	2.271	[1.901,2.660]
κ_y	N	0.125	0.050	–	–	0.121	[0.052,0.199]
ρ_a	B	0.500	0.200	0.394	[0.126,0.747]	0.412	[0.115,0.846]
ρ_n	B	0.500	0.200	0.309	[0.113,0.545]	0.224	[0.075,0.428]
ρ_i	B	0.500	0.200	0.365	[0.136,0.626]	0.374	[0.155,0.604]
ρ_c	B	0.500	0.200	0.477	[0.175,0.786]	0.888	[0.802,0.964]
ρ_g	B	0.500	0.200	0.787	[0.588,0.921]	0.786	[0.632,0.902]
ρ_m	B	0.500	0.200	–	–	0.647	[0.471,0.753]
ρ_ξ	B	0.500	0.200	0.620	[0.369,0.804]	0.833	[0.717,0.911]
σ_a^{p}	IG	1.000	4.000	0.396	[0.270,0.565]	0.406	[0.278,0.569]
σ_a^{t}	IG	1.000	4.000	0.338	[0.239,0.489]	0.347	[0.244,0.498]
σ_n	IG	1.000	4.000	0.376	[0.266,0.521]	0.378	[0.263,0.520]
σ_i^{p}	IG	1.000	4.000	0.845	[0.358,2.252]	0.610	[0.321,1.306]
σ_i^{t}	IG	1.000	4.000	5.961	[2.046,11.657]	5.805	[2.839,11.029]
σ_c	IG	1.000	4.000	0.658	[0.327,2.676]	0.357	[0.244,0.564]
σ_g	IG	1.000	4.000	1.675	[1.387,2.072]	1.705	[1.431,2.076]
σ_m	IG	1.000	4.000	–	–	0.313	[0.256,0.388]
σ_ξ	IG	1.000	4.000	1.798	[1.208,2.839]	0.613	[0.348,1.194]

Note: B, G, IG, N stand respectively for Beta, Gamma, Inverse Gamma and Normal distribution.
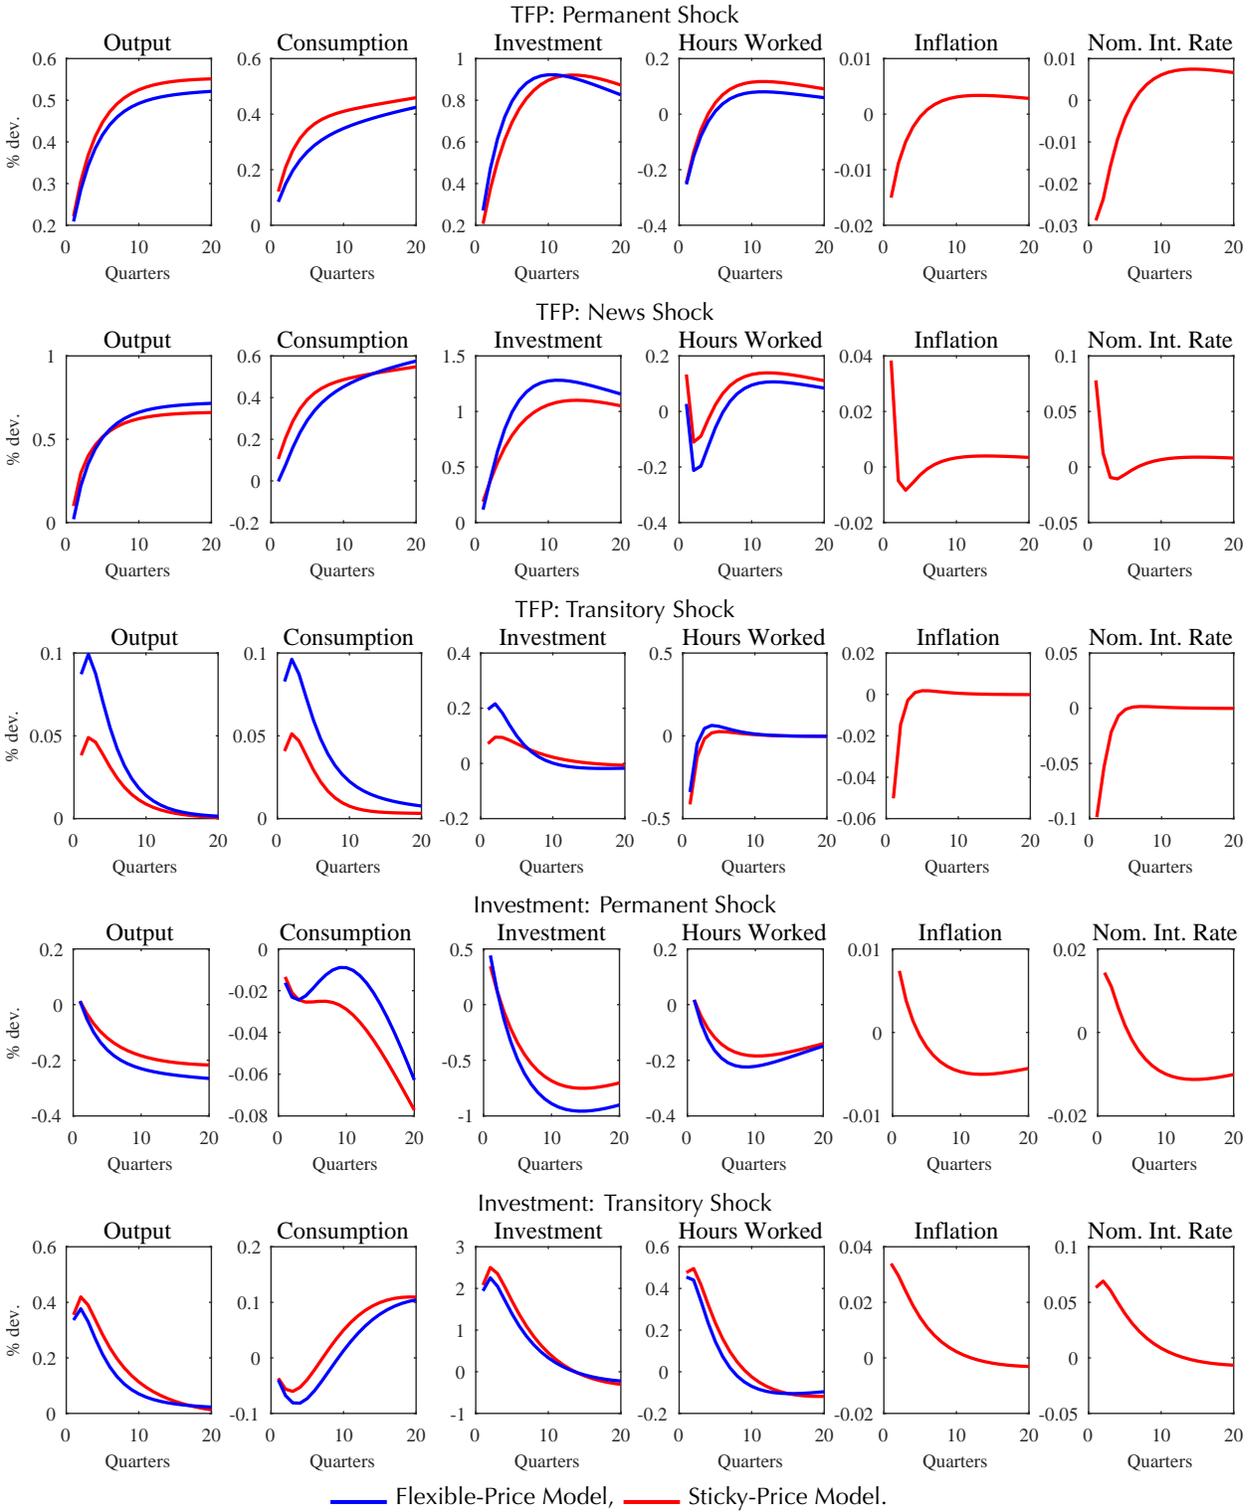
Figure 8: Theoretical IRFs, Part I
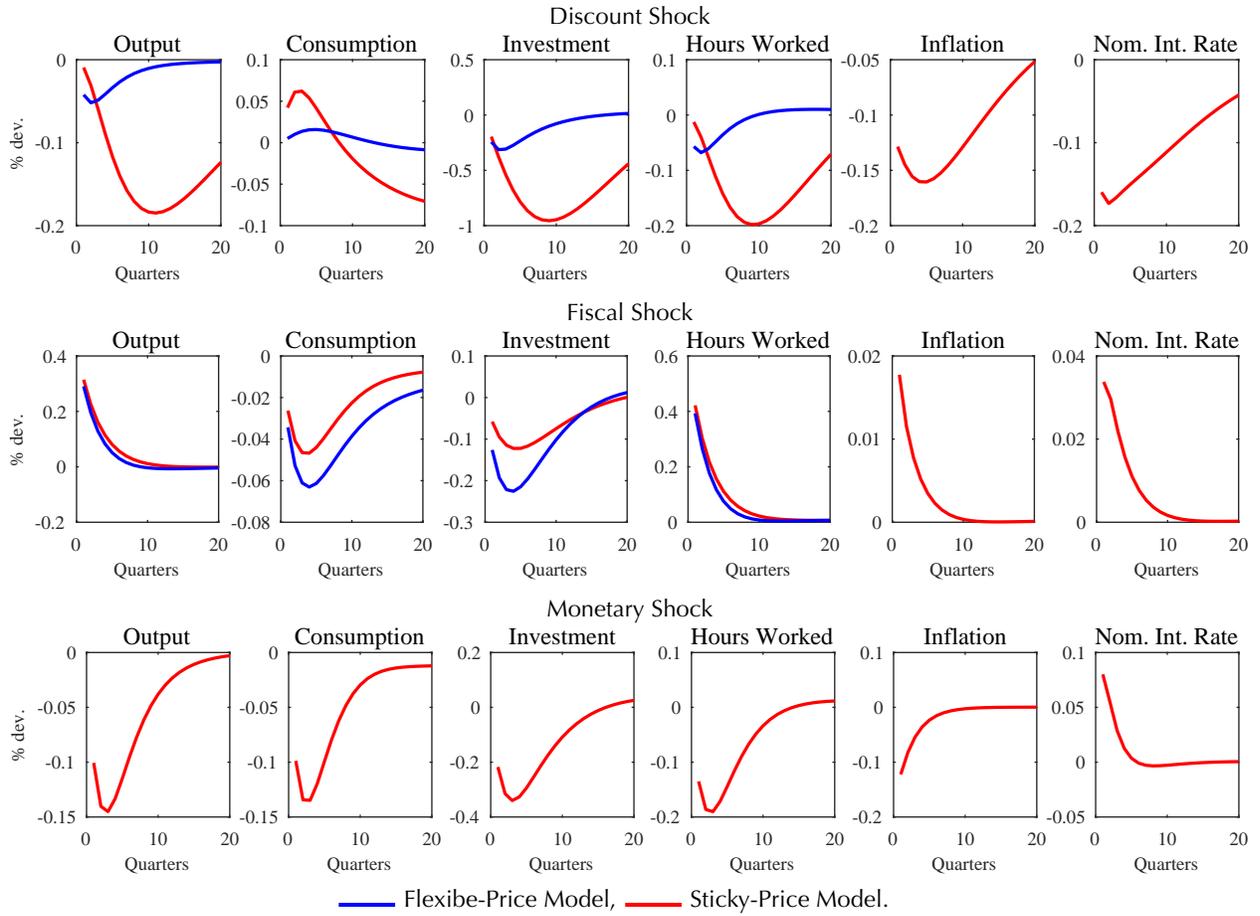
Figure 9: Theoretical IRFs, Part II

Table 9: Contribution of Shocks to Volatilities (6–32 Quarters)

	Y	C	I	N	π	R
Permanent TFP Shock						
Flexible Price	9.05	8.25	3.13	2.91	–	–
Sticky Price	9.66	12.50	1.99	2.94	0.34	1.02
In the absence of belief shock						
Flexible Price	16.45	28.30	4.50	1.30	–	–
Sticky Price	19.84	23.23	6.05	5.20	0.11	0.16
Transitory TFP Shock						
Flexible Price	0.95	2.14	0.40	2.05	–	–
Sticky Price	0.22	0.52	0.08	2.86	1.69	5.26
In the absence of belief shock						
Flexible Price	12.98	3.75	24.49	5.10	–	–
Sticky Price	0.17	0.39	0.07	7.38	0.03	0.79
Permanent Investment Shock						
Flexible Price	1.81	0.30	4.30	2.15	–	–
Sticky Price	0.98	0.09	2.34	1.13	0.13	0.38
In the absence of belief shock						
Flexible Price	49.70	39.12	26.04	78.87	–	–
Sticky Price	1.17	0.44	4.05	2.48	0.12	0.12
Transitory Investment Shock						
Flexible Price	12.94	3.90	43.52	15.16	–	–
Sticky Price	15.48	2.43	49.34	17.20	2.42	6.93
In the absence of belief shock						
Flexible Price	2.71	2.58	13.80	4.50	–	–
Sticky Price	18.58	1.38	62.23	40.70	0.16	1.59

	Y	C	I	N	π	R
News Shock						
Flexible Price	15.92	14.13	5.67	3.89	–	–
Sticky Price	13.72	16.51	2.67	1.40	0.64	1.90
In the absence of belief shock						
Flexible Price	13.37	22.39	5.52	3.06	–	–
Sticky Price	47.24	67.06	14.49	12.17	3.16	12.27
Discount Shock						
Flexible Price	0.27	0.10	0.89	0.35	–	–
Sticky Price	1.52	1.40	3.97	1.71	63.25	37.75
In the absence of belief shock						
Flexible Price	2.53	2.41	12.83	4.21	–	–
Sticky Price	8.28	3.45	12.84	20.72	96.38	55.48
Fiscal Shock						
Flexible Price	4.34	0.98	0.50	5.17	–	–
Sticky Price	4.98	0.48	0.12	5.94	0.40	1.19
In the absence of belief shock						
Flexible Price	2.26	1.45	12.81	2.97	–	–
Sticky Price	3.61	0.59	0.07	8.62	0.03	0.22
Monetary Policy Shock						
Flexible Price	–	–	–	–	–	–
Sticky Price	2.16	4.10	0.99	2.68	19.50	4.73
In the absence of belief shock						
Flexible Price	–	–	–	–	–	–
Sticky Price	1.10	3.46	0.20	2.73	0.02	29.37

Table 10: Contribution of Shocks to Comovements (6–32 Quarters)

			Covariance			
	(Y,N)	(Y,I)	(Y,C)	(N,I)	(N,C)	(I,C)
Permanent TFP Shock						
Flexible Price	0.21	5.91	10.62	0.74	1.99	8.86
Sticky Price	0.86	4.82	13.99	1.24	2.46	10.03
In the absence of belief shock						
Flexible Price	4.60	11.52	26.05	3.04	9.24	23.50
Sticky Price	5.36	11.87	27.87	5.63	8.14	31.12
Transitory TFP Shock						
Flexible Price	-0.39	0.69	1.80	-0.34	-0.56	1.61
Sticky Price	-0.42	0.15	0.44	-0.19	-1.02	0.40
In the absence of belief shock						
Flexible Price	8.97	24.10	8.36	16.84	6.45	20.31
Sticky Price	-1.14	0.12	0.33	-0.25	-5.15	0.42
Permanent Investment Shock						
Flexible Price	2.38	2.75	-0.17	2.98	-0.02	-1.18
Sticky Price	1.23	1.57	-0.00	1.60	0.06	-0.39
In the absence of belief shock						
Flexible Price	70.47	43.08	52.57	57.90	88.89	52.91
Sticky Price	2.58	2.76	-0.39	4.10	-2.16	-1.82
Transitory Investment Shock						
Flexible Price	16.64	26.60	-7.33	30.25	-7.99	-19.91
Sticky Price	18.91	31.96	-7.28	34.04	-9.72	-21.39
In the absence of belief shock						
Flexible Price	3.92	8.25	-3.04	11.86	-5.02	-12.64
Sticky Price	41.57	43.07	-3.41	66.09	-18.16	-19.19

			Covariance			
	(Y,N)	(Y,I)	(Y,C)	(N,I)	(N,C)	(I,C)
News Shock						
Flexible Price	0.10	10.67	18.07	0.40	3.47	15.08
Sticky Price	0.09	6.80	19.28	0.58	0.46	13.40
In the absence of belief shock						
Flexible Price	5.44	10.27	20.70	6.11	8.03	18.60
Sticky Price	20.27	28.84	68.05	2.85	83.48	65.41
Discount Shock						
Flexible Price	0.37	0.54	-0.14	0.67	-0.17	-0.41
Sticky Price	1.90	2.63	0.16	3.02	-0.11	-1.46
In the absence of belief shock						
Flexible Price	3.68	7.67	-2.82	11.10	-4.68	-11.73
Sticky Price	20.08	13.04	6.75	21.11	30.17	21.36
Fiscal Shock						
Flexible Price	5.81	-0.90	-1.34	-1.10	-1.79	1.23
Sticky Price	6.47	-0.45	-1.34	-0.51	-1.91	0.49
In the absence of belief shock						
Flexible Price	2.92	-4.89	-1.82	-6.85	-2.93	9.05
Sticky Price	8.60	-0.22	-1.74	-0.36	-7.46	0.47
Monetary Policy Shock						
Flexible Price	–	–	–	–	–	–
Sticky Price	2.86	1.68	3.80	1.93	5.52	4.04
In the absence of belief shock						
Flexible Price	–	–	–	–	–	–
Sticky Price	2.67	0.51	2.54	0.83	11.15	2.22

O.4 Estimating ϱ and σ_ξ

In the main text, we noted that the data considered in Section 5 do not allow us to identify separately the standard deviation of the confidence shock and the degree of strategic complementary. Nevertheless, this may be achieved if the data set were augmented to include data on expectations. In this Appendix, we elaborate on these points and describe the "augmented estimation" that motivates the value of ϱ used in Section 5.

To illustrate the main identification issue, consider again the example studied in Subsections 3.2 and 3.3. From conditions (18) and (19), we see that the volatility of the non-fundamental (confidence-driven) innovations in output is given by

$$Var\left(Y_t - Y_t^* \,|\, history\right) = \psi^2 = \frac{\omega^2}{(1-\omega)^4}\sigma_\xi^2, \tag{39}$$

where Y_t^* is the fundamental (TFP-driven) component, σ_ξ is the standard deviation of the confidence shock, and ω is the degree of strategic complementarity. Under the assumption, made in the baseline model, that the CES aggregator across the islands is Cobb-Douglas, ϱ is unity. Relaxing this assumption gives ω as a monotone function of ϱ. From condition (39) it is then evident that exactly the same non-fundamental volatility in output can be accounted for by a continuum of values for the pair (ϱ, σ_ξ). This is the crux of the identification issue faced in Section 5: the models of that section are more complicated, something that hinders analytical results, but the issue remains the same.

To illustrate how data on expectations could possibly aid identification, aggregate condition (12) to obtain the following equation:

$$N_t - N_t^* = \omega \cdot \bar{\mathbb{E}}_t[N_t - N_t^*],$$

where N_t^* denote the fundamental component of employment. This condition reveals how expectations of employment (or some other variable) together with a measure of its "fundamental" component can be used to identify the degree of strategic complementarity, and therefore ϱ.

The procedure, thought, is fraught with difficulties. Unlike the example discussed above, the models of Section 5 have expectations mattering through multiple horizons and multiple channels. It is not clear how to combine these expectations into a single measure, or how to map the theoretical objects to the available empirical measures. For instance, the University of Michigan Index of Consumer Sentiment, which is known to forecast future employment and output, is constructed on the basis of answers to qualitative questions that do not have an immediate counterpart in the theory.

These challenges, in combination with the desire to stay as close as possible to standard practice, account for our choice to estimate the models of Section 5 on the macroeconomic data alone. Note, though, that this choice does not matter for the estimated contribution of the confidence shock to the business cycle. Fixing the value of ϱ or allowing it to be estimated freely makes little difference for the shock's estimated contribution to the variances and covariances of the macroeconomic quantities.

Does the lack of identification of σ_ξ and ϱ pose a problem for our assertion that confidence shocks are a major driver of the business cycle? We think it does not. Not being able to rule out values of

σ_ξ that seem implausibly high relative to the innovations in aggregate TFP and other fundamentals only implies that a narrow interpretation of the confidence shock as capturing mis-coordination and higher-order uncertainty may be tenuous. But it allows our shock to proxy for alternative kinds of waves of optimism and pessimism, for instance, irrational beliefs.

Notwithstanding our preference for a broad interpretation of the confident shock, we now describe an exercise that supports the more narrow interpretation and justifies the value of ϱ used in Section 5. Consider either one of the models of Section 5 and construct an "augmented" model by adding the following equation, for some $k \geq 0$:

$$\text{mcsi}_t = \lambda \bar{\mathbb{E}}_t[N_{t+k}] + \eta_t \tag{40}$$

where N_{t+k} is aggregate employment k periods ahead, λ is a scalar, and η_t is a random variable, that is orthogonal to the confidence shock and other structural shocks, and that follows an AR(1) process $\eta_t = \rho_\eta \eta_{t-1} + \varepsilon_t^\eta$ where $\rho_\eta \in [0, 1)$ and $\varepsilon_t^\eta \rightsquigarrow \mathcal{N}(0, \sigma_\eta^2)$. We take mcsi_t as the theoretical counterpart of the University of Michigan Consumer Sentiment Index; η_t as measurement error, or as a crude proxy for mis-specification in the "true" relation between the theory and the aforementioned index;[34] and λ as a scaling parameter.

Now let θ be the vector that collects all the parameters of the original model, inclusive of σ_ξ and ϱ. The parameters of the augmented model are given by the union of θ and $(\lambda, \sigma_\eta, \rho_\eta)$. Trying to estimate all the parameters jointly creates a new problem that prevents the MCMC from converging properly. This seems to be due to the fact that the same covariation between the sentiment index and the macroeconomic variables can be captured with different combinations of the scaling parameter λ, the volatility of the measurement error, and the degree of strategic complementarity. To cut the Gordian knot, we chose to impose an *ad-hoc* identification restriction that requires the augmented model to produce a particular value for the share of the variance in mcsi_t that is accounted for by the measurement error η_t. This is equivalent to imposing one's prior on the noise-to-signal ratio in the sentiment index.

More specifically, for any $k \geq 0$, there exists a function v_k such that

$$\text{Var}\left(\bar{\mathbb{E}}_t[N_{t+k}]\right) = v_k(\theta).$$

This function is generated by the same system of equations as the one that pins down the equilibrium outcomes and is not affected by the addition of equation (40). It follows that the relative contribution of the measurement error in the theoretical counterpart of the sentiment index is given by

$$\frac{\text{Var}(\eta_t)}{\text{Var}(\text{mcsi}_t)} = \mathcal{M}(\lambda, \theta') \equiv \frac{\sigma_\eta^2}{(1 - \rho_\eta)\lambda^2 v_k(\theta) + \sigma_\eta^2}$$

where $\theta' \equiv (\theta, \sigma_\eta, \rho_\eta)$. For any θ' and any target $me \in (0, 1)$ for the contribution of the measurement error, solving the equation $\mathcal{M}(\lambda, \theta') = me$ gives the value of λ that is consistent with that target.

---

[34]In the data, the correlation of the Consumer Sentiment Index with hours worked attains a maximal value of about 0.7 when the former leads the latter by 3 quarters. To some extent, this corroborates the specification assumed above and suggests $k = 3$ as a possible benchmark.

Fixing a value for me is therefore equivalent to adding an identification restriction on the parameters of the augmented model; in that case the MCMC converges properly and θ' is well identified for any given me. Our strategy is therefore to select various values for me and to estimate θ' on the data used in Section 5 together with the times series of the aforementioned index.

Table 11: Estimating both σ_ξ and ϱ

k	me	ϱ	σ_ξ	σ_ξ^2/σ_a^2	Y	C	I	N
		Estimated Parameters			Variance Contribution of Confidence Shock			
Flexible-Price Model								
0	0.25	0.70	0.64	0.99	54.26	54.45	50.29	84.56
	0.50	0.72	0.57	0.78	54.41	52.35	52.01	85.19
4	0.25	0.70	0.43	0.40	45.80	24.40	47.79	84.01
	0.50	0.70	0.45	0.44	46.26	29.53	49.25	84.78
20	0.25	0.79	0.44	0.41	58.09	46.31	61.04	85.33
	0.50	0.77	0.45	0.47	55.60	44.49	58.58	84.75
Sticky-Price Model								
0	0.25	0.60	0.58	0.27	12.37	13.52	10.04	21.69
	0.50	0.72	0.36	0.32	53.99	50.74	51.06	71.50
4	0.25	0.69	0.36	0.30	45.15	41.93	39.87	64.26
	0.50	0.72	0.36	0.30	46.12	43.28	43.05	67.51
20	0.25	0.65	0.38	0.33	50.80	39.60	46.87	70.36
	0.50	0.69	0.38	0.33	49.45	41.89	46.34	69.84

The results from the "augmented" estimation are reported in Table 11. Let us focus on the flexible-price model and consider two values for me, the share of the measurement error, and three values for k, the horizon of the expectations that show up in condition (40). For each k and me (first two columns), the table reports the estimated values for ϱ and σ_ξ (next two columns), the ratio of the estimated σ_ξ to the estimated σ_a (fifth column), and the estimated contributions of the confidence shock to the business-cycle volatilities of output, consumption, investment and hours (last four columns). The findings suggest a value of ϱ in the neighborhood of 0.75, which in turn motivates the value used in Section 5. Furthermore, the estimated σ_ξ is smaller than the estimated σ_a, allowing for a narrow interpretation of the confidence shock. Finally, the estimated contribution of the confidence shock to the business cycle is of the same magnitude as the one estimated in Section 5. There are, however, two notable differences: the confidence shock now explains a larger share of the volatility in hours and a smaller in consumption.

We find the results of this exercise useful even if they do not constitute proof that ϱ and σ_ξ lie in those ranges.

O.5   Log-Linear Solution

In this appendix we explain how to augment a large class of DSGE models with our proposed type of higher-order belief dynamics and how to obtain the solution of the augmented model as a simple transformation of the solution of the original model.

A Prelude. Before considering the general case it is instructive to review the linearized version of our baseline model.

The log-linearized equilibrium conditions of the model are given by (7)-(11) in Section 3 and have a familiar interpretation. The only novelty is the presence of two distinct expectation operators $\mathbb{E}_{it}$ and $\mathbb{E}'_{it}$, which denote local expectations in stage 1 and stage 2 of period t respectively. The difference between these two expectation operators derives from the fact that islands form beliefs about one another's signals and thereby about Y_t in stage 1 on the basis of their mis-specified priors, but observe the true state of nature and the true realized Y_t in stage 2. Under the supply first timing protocol, the first expectation shows up in the optimality condition for labor, while the second shows up in the optimality condition for consumption/saving.

The following points are worth emphasizing. The aggregate-level variables are, of course, obtained from averaging the individual-level variables across all islands. In equilibrium, the *realized* values of the aggregate variables coincide with the realized values of the corresponding individual variables; e.g., $y_{it} = Y_t$ for all i, all t, and all realizations of uncertainty. This is because all islands receive the same signals and the same fundamentals. However, this does not mean that one can just replace the island-specific variables in the above conditions with the aggregate ones, or vice versa. Even though the "objective truth" is that all islands receive the same signals, in stage 1 of each period each island believes that the signals of other islands can differ from its own signal. Accordingly, each island reasons that y_{it} can differ form Y_t, even when all other islands follow the same strategy as itself and receive the same TFP shock.

Keeping track of this delicate difference between the realizations and the beliefs of different variables is key to obtaining the solution to the model. Our method deals with this delicate matter by (i) using appropriate notation to distinguish the signal received by each agent/island from either the average signal in the population or the true underlying shock to fundamentals; and (ii) choosing appropriate state spaces for both the individual policy rules and the aggregate ones.

In what follows, we first set up the general class of log-linear DSGE models that our solution method handles. We next introduce a class of linear policy rules, which describe the behavior of each agent as a function of his information set. Assuming that all other islands follow such a policy rules, we can use the equilibrium conditions of the model to obtain the policy rules that are optimal for the individual island; that is, we can characterize the best responses of the model. Since the policy rules are linear, they are parameterized by a collection of coefficients (matrices), and the aforementioned best responses reduce to a system of equations in these coefficients. The solution to this system gives the equilibrium of the model.

A "generic" DSGE model. We henceforth consider an economy whose equilibrium is represented by the following linear dynamic system:

$$M_{yy}y_{it} = M_{yx}x_{it}^b + M_{yX}X_t^b + M_{yY}\mathbb{E}_{it}Y_t + M_{yf}\mathbb{E}_{it}x_{it}^f + M_{yF}\mathbb{E}_{it}X_t^f + M_{ys}z_{it}$$

$$M_{xx0}x_{it+1}^b = M_{xx1}x_{it}^b + M_{xX1}X_t^b + M_{xy1}y_{it} + M_{xY1}Y_t + M_{xf1}x_{it}^f + M_{xF1}X_t^f + M_{xs1}s_t$$

$$M_{ff0}\mathbb{E}_{it}'x_{it+1}^f = M_{fF0}\mathbb{E}_{it}'X_{t+1}^f + M_{ff1}x_{it}^f + M_{fF1}X_t^f + M_{fx0}x_{it+1}^b + M_{fx1}x_{it}^b + M_{fX1}X_t^b$$

$$+ M_{fy0}\mathbb{E}_{it}'y_{it+1} + M_{fY0}\mathbb{E}_{it}'Y_{t+1} + M_{fy1}y_{it} + M_{fY1}Y_t + M_{fs0}\mathbb{E}_{it}'s_{t+1} + M_{fs1}s_t$$

$$s_t = Rs_{t-1} + \varepsilon_t$$

$$\xi_t = Q\xi_{t-1} + \nu_t$$

This system is a generalization of the one we obtained in our baseline RBC model. Here, x^b, x^f, y, s, and ξ are allowed to be vectors; x^b collects backward-looking variables (such as capital in our model); x^f collects forward-looking variables that are chosen in stage 2 of each period (such as consumption and investment in our model); y collects the variables that are instead chosen in stage 1 (such as employment in our model); s collects the shocks to payoff (such as technology); and finally ξ is meant to capture the shocks to higher-order beliefs. X^b, X^f and Y correspond to the aggregate versions of, respectively, x^b, x^f and y.

Beliefs. We assume that, as of stage 2, the realizations of s_t, of all the signals, and of all the stage-1 choices become commonly known, which implies that $y_{it}, x_{it}^f, x_{it+1}^b$ and Y_t, X_t^f, X_{t+1}^b are also commonly known in equilibrium). Furthermore, the actual realizations of the signals satisfy $z_{it} = s_t$ for all t and all i. However, the agents have misspecified belief in stage 1. In particular, for all i, all $j \neq i$, all t, and all states of nature, agent i's belief during stage 1 satisfy

$$\mathbb{E}_{it}[s_t] = z_{it},$$

$$\mathbb{E}_{it}[\mathbb{E}_{jt}s_t] = \mathbb{E}_{it}[z_{jt}] = z_{it} + \Delta\xi_t,$$

where z_{it} is the signal received by agent i, ξ_t is the higher-order belief shocks, and Δ is a loading matrix. We next let $\bar{z}_t$ denote the average signal in the economy and note that the "truth" is that $z_{it} = \bar{z}_t = s_t$. Yet, this truth is publicly revealed only in stage 2 of period t. In stage 1, instead, each island believes, incorrectly, that

$$\mathbb{E}_{it}\bar{z}_t = z_{it} + \Delta\xi_t.$$

Note next that the stage-1 variables, y_{it}, can depend on the local signal z_{it}, along with the commonly-observed belief shock ξ_t and the backward-looking (predetermined) state variables x_{it}^b and X_t^b, but cannot depend on either the aggregate signal $\bar{z}_t$ or the underlying fundamental s_t, because these variables are not known in stage 1. By contrast, the stage-2 decisions depend on the entire triplet $(z_{it}, \bar{z}_t, s_t)$. As already mentioned, the truth is that these three variables coincide. Nevertheless, the islands believe in stage 1 that the average signal can differ from either their own signal or the actual fundamental. Accordingly, it is important to write stage-2 strategies as functions of the three conceptually distinct objects in $(z_{it}, \bar{z}_t, s_t)$ in order to do specify the appropriate equilibrium beliefs

in stage-1. (Note that this is equivalent to expressing the stage-2 strategies as functions of the realized values of the stage-1 variables y and Y, which is the approach we took in the characterization of the recursive equilibrium in Section 3.) In what follows, we show how this belief structure facilitates a tractable solution of the aforementioned general DSGE model.

Preview of key result. To preview the key result, let us first consider the underlying "belief-free" model, that is, of the complete-information, representative-agent, counterpart of the model we are studying. The equilibrium system is given by the following:

$$Y_t = M_X X_t^b + M_{EY} Y_t + M_F X_t^f + M_s s_t$$

$$X_{t+1}^b = N_X X_t^b + N_Y Y_t + N_F X_t^f + N_s s_t$$

$$(P_{f0} - P_{F0})\mathbb{E}_t X_{t+1}^f = P_{F1} X_t^f + P_{Y0}\mathbb{E}_t Y_{t+1} + P_X X_t^b + P_{Y1} Y_t + P_s s_t$$

$$s_t = R s_{t-1} + \varepsilon_t$$

$$\xi_t = Q \xi_{t-1} + \nu_t$$

(This system can be obtained from the one we introduced before once we impose the restriction that all period-t variables are commonly known in period t, which means that $\mathbb{E}'_{it}[x_t] = \mathbb{E}_{it}[x_t] = x_t$ for any variable x.) It is well known how to obtain the policy rules of such a representative-agent model. Our goal in this appendix is to show how the policy rules of the belief-augmented model that we described above can be obtained as a simple, tractable transformation of the policy rules of the representative-agent benchmark.

In particular, we will show that the policy rules for our general DSGE economy are as follows:

$$\mathbf{X}_t = \Theta_X X_t^b + \Theta_s s_t + \Theta_\xi \xi_t,$$

where $\mathbf{X}_t = (Y_t, X_t^f, X_{t+1}^b)$ collects all the variables, Θ_X and Θ_s are the same matrices as those that appear in the solution of the underlying belief-free model, and Θ_ξ is a new matrix, which encapsulates the effects of higher-order beliefs.

The model, restated. To ease subsequent algebraic manipulations, we henceforth restate the model as follows:

$$y_{it} = M_x(x_{it}^b - X_t^b) + M_X X_t^b + M_{EY}\mathbb{E}_{it} Y_t + M_f \mathbb{E}_{it}(x_{it}^f - X_t^f) + M_F \mathbb{E}_{it} X_t^f + M_s z_{it} \qquad (41)$$

$$x_{it+1}^b = N_x(x_{it}^b - X_t^b) + N_X X_t^b + N_y(y_{it} - Y_t) + N_Y Y_t + N_f(x_{it}^f - X_t^f) + N_F X_t^f + N_s s_t \qquad (42)$$

$$P_{f0}\mathbb{E}'_{it} x_{it+1}^f = P_{f1}(x_{it}^f - X_t^f) + P_{F0}\mathbb{E}'_{it} X_{t+1}^f + P_{F1} X_t^f + P_x(x_{it}^b - X_t^b) + P_X X_t^b +$$
$$+ P_{y0}(\mathbb{E}'_{it} y_{it+1} - \mathbb{E}'_{it} Y_{t+1}) + P_{Y0}\mathbb{E}'_{it} Y_{t+1} + P_{y1}(y_{it} - Y_t) + P_{Y1} Y_t + P_s s_t \qquad (43)$$

where

$$M_x = M_{yy}^{-1} M_{yx}, \ M_X = M_{yy}^{-1}(M_{yx} + M_{yX}), \ M_{EY} = M_{yy}^{-1} M_{yY},$$

$$M_f = M_{yy}^{-1} M_{yf}, \ M_F = M_{yy}^{-1}(M_{yf} + M_{yF}), \ M_s = M_{yy}^{-1} M_{ys}$$

$$N_x = M_{xx0}^{-1} M_{xx1}, \ N_X = M_{xx0}^{-1}(M_{xx1} + M_{xX1}), \ N_y = M_{xx0}^{-1} M_{xy1}, \ N_Y = M_{xx0}^{-1}(M_{xy1} + M_{xY1}),$$

$$N_f = M_{xx0}^{-1} M_{xf1}, \ N_F = M_{xx0}^{-1}(M_{xf1} + M_{xF1}), \ N_s = M_{xx0}^{-1} M_{xs1}$$

$$P_{f0} = M_{ff0}, \ P_{f1} = M_{ff1} + M_{fx0}N_f, \ P_{F0} = M_{fF0}, \ P_{F1} = M_{fF1} + M_{ff1} + M_{fx0}N_F$$

$$P_x = M_{fx1} + M_{fx0}N_x, \ P_X = M_{fX1} + M_{fx1} + M_{fx0}N_X,$$

$$P_{y0} = M_{fy0}, \ P_{Y0} = M_{fY0} + M_{fy0}, \ P_{y1} = M_{fy1} + M_{fx0}N_y, \ P_Y = M_{fy1} + M_{fY1} + M_{fx0}N_Y,$$

$$P_s = M_{fs0}R + M_{fs1} + M_{fx0}N_s$$

Proposed Policy Rules. We propose that the equilibrium policy rules take the following form:

$$y_{it} = \Lambda_x(x_{it}^b - X_t^b) + \Lambda_X X_t^b + \Lambda_z z_{it} + \Lambda_\xi \xi_t \tag{44}$$

$$x_{it}^f = \Gamma_x(x_{it}^b - X_t^b) + \Gamma_X X_t^b + \Gamma_z z_{it} + \Gamma_{\bar{z}} \bar{z}_t + \Gamma_s s_t + \Gamma_\xi \xi_t \tag{45}$$

where the Λ's and Γ's are coefficients (matrices), whose equilibrium values are to be obtained in the sequel. Following our earlier discussion, note that the stage-2 policy rules are allowed to depend on the triplet $(z_{it}, \bar{z}_t, s_t)$, while the stage-1 policy rules are restricted to depend only on the local signal z_{it}. It is also useful to note that we would obtain the same solution if we were to represent the stage-2 policy rules as functions of y_{it} and Y_t in place of, respectively, z_{it} and $\bar{z}_t$: the latter two variables enter the equilibrium conditions that determine the stage-2 decisions, namely conditions (42) and (43), *only* through the realized values of the stage-1 outcomes y_{it} and Y_t.

Obtaining the solution. We obtain the solution in three steps. In step 1, we start by characterizing the equilibrium determination of the stage-1 policy rules, taking as given the stage-2 rules. Formally, we fix an arbitrary rule in (45); we assume that all islands believe that the stage-2 variables are determined according to this rule; and we then look for the particular rule in (44) that solves the fixed-point relation between y_{it} and Y_t described in (41) under this assumption. This step, which we can think of as the "static" component of the equilibrium, gives as a mapping from Γ matrices to the Λ matrices. In step 2, we obtain a converse mapping by characterize the policy rules for the forward-looking variables that solve conditions (42) and (43) under the assumption that the stage-1 outcomes are determined according to an arbitrary rule in (45). We can think of this step as solving for the "dynamic" component of the equilibrium. In step 3, we use the fixed-point between these two mappings to obtain the overall solution to the model.

Step 1. As noted above, we start by studying the equilibrium determination of the stage-1 policy rules, taking as given the stage-2 policy rules.

Thus suppose that all islands follow a policy rule as in (45) and consider the beliefs that a given island i forms, under this assumption, about the stage-2 variables x_{it}^f and X_t^f. From (45), we have

$$x_{it}^f = \Gamma_x(x_{it}^b - X_t^b) + \Gamma_X X_t^b + \Gamma_z z_{it} + \Gamma_{\bar{z}} \bar{z}_t + \Gamma_s s_t + \Gamma_\xi \xi_t$$

$$X_t^f = \Gamma_X X_t^b + (\Gamma_z + \Gamma_{\bar{z}})\bar{z}_t + \Gamma_s s_t + \Gamma_\xi \xi_t$$

Along with the fact that $\mathbb{E}_{it}[s_t] = z_{it}$ and $\mathbb{E}_{it}[\bar{z}_t] = z_{it} + \Delta\xi_t$, the above gives

$$\mathbb{E}_{it}x_{it}^f = \Gamma_x(x_{it}^b - X_t^b) + \Gamma_X X_t^b + (\Gamma_z + \Gamma_{\bar{z}} + \Gamma_s)z_{it} + (\Gamma_\xi + \Gamma_{\bar{z}}\Delta)\xi_t$$

$$\mathbb{E}_{it}X_t^f = \Gamma_X X_t^b + (\Gamma_z + \Gamma_{\bar{z}} + \Gamma_s)z_{it} + (\Gamma_\xi + (\Gamma_z + \Gamma_{\bar{z}})\Delta)\xi_t$$

20

which also implies that

$$x_{it}^f - X_t^f = \Gamma_x(x_{it}^b - X_t^b) + \Gamma_z(z_{it} - \bar{z}_t)$$
$$\mathbb{E}_{it}(x_{it}^f - X_t^f) = \Gamma_x(x_{it}^b - X_t^b) - \Gamma_z\Delta\xi_t$$

Plugging the above in (41), the equilibrium equation for y_{it}, we get

$$
\begin{aligned}
y_{it} &= M_x(x_{it}^b - X_t^b) + M_X X_t^b + M_{EY}\mathbb{E}_{it}Y_t + M_f\mathbb{E}_{it}(x_{it}^f - X_t^f) + M_F\mathbb{E}_{it}X_t^f + M_s z_{it} \\
&= M_x(x_{it}^b - X_t^b) + M_X X_t^b + M_{EY}\mathbb{E}_{it}Y_t + M_f\left[\Gamma_x(x_{it}^b - X_t^b) - \Gamma_z\Delta\xi_t\right] \\
&\quad + M_F\left[\Gamma_X X_t^b + (\Gamma_z + \Gamma_{\bar{z}} + \Gamma_s)z_{it} + (\Gamma_\xi + (\Gamma_z + \Gamma_{\bar{z}})\Delta)\xi_t\right] + M_s z_{it}
\end{aligned}
$$

Equivalently,

$$
\begin{aligned}
y_{it} &= (M_x + M_f\Gamma_x)(x_{it}^b - X_t^b) + (M_X + M_F\Gamma_X)X_t^b + M_{EY}\mathbb{E}_{it}Y_t \\
&\quad + (M_s + M_F(\Gamma_z + \Gamma_{\bar{z}} + \Gamma_s))z_{it} + (M_F\Gamma_\xi + M_F\Gamma_{\bar{z}}\Delta + (M_F - M_f)\Gamma_z\Delta)\xi_t
\end{aligned}
\tag{46}
$$

Note that the above represents a *static* fixed-point relation between y_{it} and Y_t. This relation is itself determined by the Γ matrices (i.e., by the presumed policy rule for the stage-2 variables). Notwithstanding this fact, we now focus on the solution of this static fixed point.

Thus suppose that this solution takes the form of a policy rule as in (44). If all other island follow this rule, then at the aggregate we have

$$Y_t = \Lambda_X X_t^b + \Lambda_z\bar{z}_t + \Lambda_\xi\xi_t$$

and therefore the stage-1 forecast of island i about Y_t is given by

$$\mathbb{E}_{it}Y_t = \Lambda_X X_t^b + \Lambda_z z_{it} + (\Lambda_\xi + \Lambda_z\Delta)\xi_t$$

Plugging this into (46), we obtain the following best response for island i:

$$
\begin{aligned}
y_{it} &= (M_x + M_f\Gamma_x)(x_{it}^b - X_t^b) + (M_X + M_F\Gamma_X)X_t^b + M_{EY}\left(\Lambda_X X_t^b + \Lambda_z z_{it} + (\Lambda_\xi + \Lambda_z\Delta)\xi_t\right) \\
&\quad + (M_s + M_F(\Gamma_z + \Gamma_{\bar{z}} + \Gamma_s))z_{it} + (M_F(\Gamma_\xi + \Gamma_{\bar{z}}\Delta) + (M_F - M_f)\Gamma_z\Delta)\xi_t
\end{aligned}
$$

For this to be consistent with our guess in (44), we must have

$$\Lambda_x = M_x + M_f\Gamma_x \tag{47}$$
$$\Lambda_X = (I - M_{EY})^{-1}(M_X + M_F\Gamma_X) \tag{48}$$
$$\Lambda_z = (I - M_{EY})^{-1}\left[M_s + M_F(\Gamma_z + \Gamma_{\bar{z}} + \Gamma_s)\right] \tag{49}$$
$$\Lambda_\xi = (I - M_{EY})^{-1}\left\{M_F(\Gamma_\xi + \Gamma_{\bar{z}}\Delta) + (M_F - M_f)\Gamma_z\Delta + M_{EY}\Lambda_z\Delta\right\} \tag{50}$$

This completes the first step of our solution strategy: we have characterized the "static" component of the equilibrium and have thus obtained the Λ coefficients as functions of primitives and of the Γ coefficients.

Step 2. We now proceed with the second step, which is to characterize the equilibrium behavior in stage 2, taking as given the behavior in stage 1.

Recall that, once agents enter stage 2, they observe the true current values of the triplet $(z_{it}, \bar{z}_t, s_t)$ along with the realized values of the past stage-1 outcomes, y_{it} and Y_t. Furthermore, in equilibrium this implies common certainly of current choices, namely of the variables x_{it}^f and X_t^f, and thereby also of the variables x_{it+1}^b and X_{t+1}^b. Nevertheless, agents face uncertainty about the next-period realizations of the aforementioned triplet and of the corresponding endogenous variables. In what follows, we thus take special care in characterizing the beliefs that agents form about the relevant future outcomes.

Consider first an agent's beliefs about the aggregate next-period stage-1 variables:

$$Y_{t+1} = \Lambda_X X_{t+1}^b + \Lambda_z \bar{z}_{t+1} + \Lambda_\xi \xi_{t+1}$$

$$\mathbb{E}_{it+1} Y_{t+1} = \Lambda_X X_{t+1}^b + \Lambda_z z_{it+1} + (\Lambda_\xi + \Lambda_z \Delta) \xi_{t+1}$$

$$\mathbb{E}'_{it} Y_{t+1} = \Lambda_X X_{t+1}^b + \Lambda_z R s_t + (\Lambda_\xi + \Lambda_z \Delta) Q \xi_t$$

Consider next his beliefs about his own next-period stage-1 variables:

$$y_{it+1} = \Lambda_x (x_{it+1}^b - X_{t+1}^b) + \Lambda_X X_{t+1}^b + \Lambda_z z_{it+1} + \Lambda_\xi \xi_{t+1}$$

$$\mathbb{E}'_{it} y_{it+1} = \Lambda_x (x_{it+1}^b - X_{t+1}^b) + \Lambda_X X_{t+1}^b + \Lambda_z R s_t + \Lambda_\xi Q \xi_t$$

It follows that

$$\mathbb{E}'_{it} (y_{it+1} - Y_{t+1}) = \Lambda_x (x_{it+1}^b - X_{t+1}^b) - \Lambda_z \Delta Q \xi_t$$

Consider now his beliefs about his own next-period forward variables:

$$x_{it+1}^f = \Gamma_x (x_{it+1}^b - X_{t+1}^b) + \Gamma_X X_{t+1}^b + \Gamma_z z_{it+1} + \Gamma_{\bar{z}} \bar{z}_{t+1} + \Gamma_s s_{t+1} + \Gamma_\xi \xi_{t+1}$$

$$\mathbb{E}_{it+1} x_{it+1}^f = \Gamma_x (x_{it+1}^b - X_{t+1}^b) + \Gamma_X X_{t+1}^b + (\Gamma_z + \Gamma_{\bar{z}} + \Gamma_s) z_{it+1} + (\Gamma_\xi + \Gamma_{\bar{z}} \Delta) \xi_{t+1}$$

$$\mathbb{E}'_{it} x_{it+1}^f = \Gamma_x (x_{it+1}^b - X_{t+1}^b) + \Gamma_X X_{t+1}^b + (\Gamma_z + \Gamma_{\bar{z}} + \Gamma_s) R s_t + (\Gamma_\xi + \Gamma_{\bar{z}} \Delta) Q \xi_t$$

For the aggregate next-period forward variables we have

$$\mathbb{E}_{it+1} X_{t+1}^f = \Gamma_X X_{t+1}^b + (\Gamma_z + \Gamma_{\bar{z}} + \Gamma_s) z_{it+1} + (\Gamma_\xi + (\Gamma_z + \Gamma_{\bar{z}}) \Delta) \xi_{t+1}$$

$$\mathbb{E}'_{it} X_{t+1}^f = \Gamma_X X_{t+1}^b + (\Gamma_z + \Gamma_{\bar{z}} + \Gamma_s) R s_t + (\Gamma_\xi + (\Gamma_z + \Gamma_{\bar{z}}) \Delta) Q \xi_t$$

and therefore

$$\mathbb{E}'_{it} (x_{it+1}^f - X_{t+1}^f) = \Gamma_x (x_{it+1}^b - X_{t+1}^b) - \Gamma_z \Delta Q \xi_t$$

Next, note that our guesses for the policy rules imply the following properties for the current-period variables:

$$y_{it} - Y_t = \Lambda_x (x_{it}^b - X_t^b) + \Lambda_z (z_{it} - \bar{z}_t)$$

$$x_{it}^f - X_t^f = \Gamma_x (x_{it}^b - X_t^b) + \Gamma_z (z_{it} - \bar{z}_t)$$

$$Y_t = \Lambda_X X_t^b + \Lambda_z \bar{z}_t + \Lambda_\xi \xi_t$$

$$X_t^f = \Gamma_X X_t^b + (\Gamma_z + \Gamma_{\bar{z}}) \bar{z}_t + \Gamma_s s_t + \Gamma_\xi \xi_t$$

22

Plugging these results in the law of motion of backward variables, we get

$$
\begin{aligned}
x_{it+1}^b &= N_x(x_{it}^b - X_t^b) + N_X X_t^b + N_y(y_{it} - Y_t) + N_Y Y_t + N_f(x_{it}^f - X_t^f) + N_F X_t^f + N_s s_t \\
&= N_x(x_{it}^b - X_t^b) + N_X X_t^b + N_y\left\{\Lambda_x(x_{it}^b - X_t^b) + \Lambda_z(z_{it} - \bar{z}_t)\right\} + N_Y\left\{\Lambda_X X_t^b + \Lambda_z \bar{z}_t + \Lambda_\xi \xi_t\right\} \\
&\quad + N_f\left\{\Gamma_x(x_{it}^b - X_t^b) + \Gamma_z(z_{it} - \bar{z}_t)\right\} + N_F\left\{\Gamma_X X_t^b + (\Gamma_z + \Gamma_{\bar{z}})\bar{z}_t + \Gamma_s s_t + \Gamma_\xi \xi_t\right\} + N_s s_t
\end{aligned}
$$

Equivalently,

$$
x_{it+1}^b = \Omega_x(x_{it}^b - X_t^b) + \Omega_X X_t^b + \Omega_z z_{it} + \Omega_{\bar{z}}\bar{z} + \Omega_s s_t + \Omega_\xi \xi_t
$$

and hence

$$
X_{t+1}^b = \Omega_X X_t^b + (\Omega_z + \Omega_{\bar{z}})\bar{z}_t + \Omega_s s_t + \Omega_\xi \xi_t
$$
$$
x_{it+1}^b - X_{t+1}^b = \Omega_x(x_{it}^b - X_t^b) + \Omega_z(z_{it} - \bar{z}_t)
$$

where

$$
\begin{aligned}
\Omega_x &= N_x + N_y\Lambda_x + N_f\Gamma_x & \Omega_z &= N_y\Lambda_z + N_f\Gamma_z \\
\Omega_X &= N_X + N_Y\Lambda_X + N_F\Gamma_X & \Omega_{\bar{z}} &= (N_Y - N_y)\Lambda_z + (N_F - N_f)\Gamma_z + N_F\Gamma_{\bar{z}} \\
\Omega_s &= N_s + N_F\Gamma_s & \Omega_\xi &= N_Y\Lambda_\xi + N_F\Gamma_\xi
\end{aligned}
$$

It follows that

$$
\begin{aligned}
\mathbb{E}_{it}' x_{it+1}^f &= \Gamma_x(x_{it+1}^b - X_{t+1}^b) + \Gamma_X X_{t+1}^b + (\Gamma_z + \Gamma_{\bar{z}} + \Gamma_s)R s_t + (\Gamma_\xi + \Gamma_{\bar{z}}\Delta)Q\xi_t \\
&= \Gamma_x\left\{\Omega_x(x_{it}^b - X_t^b) + \Omega_z(z_{it} - \bar{z}_t)\right\} + \Gamma_X\left\{\Omega_X X_t^b + (\Omega_z + \Omega_{\bar{z}})\bar{z}_t + \Omega_s s_t + \Omega_\xi \xi_t\right\} \\
&\quad + (\Gamma_z + \Gamma_{\bar{z}} + \Gamma_s)R s_t + (\Gamma_\xi + \Gamma_{\bar{z}}\Delta)Q\xi_t
\end{aligned}
$$

or equivalently

$$
\mathbb{E}_{it}' x_{it+1}^f = \Phi_x(x_{it}^b - X_t^b) + \Phi_X X_t^b + \Phi_z z_{it} + \Phi_{\bar{z}}\bar{z}_t + \Phi_s s_t + \Phi_\xi \xi_t \tag{51}
$$

where

$$
\begin{aligned}
\Phi_x &= \Gamma_x\Omega_x & \Phi_z &= \Gamma_x\Omega_z & \Phi_s &= \Gamma_X\Omega_s + (\Gamma_z + \Gamma_{\bar{z}} + \Gamma_s)R \\
\Phi_X &= \Gamma_X\Omega_X & \Phi_{\bar{z}} &= (\Gamma_X - \Gamma_x)\Omega_z + \Gamma_X\Omega_{\bar{z}} & \Phi_\xi &= \Gamma_X\Omega_\xi + (\Gamma_\xi + \Gamma_{\bar{z}}\Delta)Q
\end{aligned}
$$

Similarly, the expectation of the corresponding aggregate variable is given by

$$
\mathbb{E}_{it}' X_{t+1}^f = \Phi_X X_t^b + \Phi_z z_{it} + \Phi_{\bar{z}}\bar{z}_t + \Phi_s s_t + (\Phi_\xi + \Gamma_z\Delta Q)\xi_t \tag{52}
$$

With the above steps, we have calculated all the objects that enter the Euler condition (43). We can thus proceed to characterize the fixed-point relation that pins down the solution for the stage-2 policy rule.

To ease the exposition, let us repeat the Euler condition (43) below:

$$
\begin{aligned}
P_{f0}\mathbb{E}_{it}' x_{it+1}^f &= P_{f1}(x_{it}^f - X_t^f) + P_{F0}\mathbb{E}_{it}' X_{t+1}^f + P_{F1}X_t^f + P_x(x_{it}^b - X_t^b) + P_X X_t^b + \\
&\quad + P_{y0}(\mathbb{E}_{it}' y_{it+1} - \mathbb{E}_{it}' Y_{t+1}) + P_{Y0}\mathbb{E}_{it}' Y_{t+1} + P_{y1}(y_{it} - Y_t) + P_{Y1}Y_t + P_s s_t
\end{aligned}
$$

Use now (51) to write the left-hand-side of the Euler condition as

$$P_{f0}\mathbb{E}'_{it}x^f_{it+1} \;=\; P_{f0}\left\{\Phi_x(x^b_{it} - X^b_t) + \Phi_X X^b_t + \Phi_z z_{it} + \Phi_{\bar z}\bar z_t + \Phi_s s_t + \Phi_\xi \xi_t\right\}$$

Next, use our preceding results to replace all the expectations that show up in the right-hand-side of the Euler condition, as well as the stage-1 outcomes. This gives

$$
\begin{aligned}
P_{f0}\mathbb{E}'_{it}x^f_{it+1} \;=\; & P_{f1}\left\{\Gamma_x(x^b_{it} - X^b_t) + \Gamma_z(z_{it} - \bar z_t)\right\} + \\
& + P_{F0}\left\{\Phi_X X^b_t + (\Phi_z + \Phi_{\bar z})\bar z_t + \Phi_s s_t + (\Phi_\xi + \Gamma_z \Delta Q)\xi_t\right\} \\
& + P_{F1}\left\{\Gamma_X X^b_t + (\Gamma_z + \Gamma_{\bar z})\bar z_t + \Gamma_s s_t + \Gamma_\xi \xi_t\right\} \quad + P_x\left\{x^b_{it} - X^b_t\right\} \\
& + P_X X^b_t + P_{y0}\left\{\Lambda_x\left(\Omega_x(x^b_{it} - X^b_t) + \Omega_z(z_{it} - \bar z_t)\right) - \Lambda_z \Delta Q\xi_t\right\} \\
& + P_{Y0}\left\{\Lambda_X\left(\Omega_X X^b_t + (\Omega_z + \Omega_{\bar z})\bar z_t + \Omega_s s_t + \Omega_\xi \xi_t\right) + \Lambda_z R s_t + (\Lambda_\xi + \Lambda_z \Delta)Q\xi_t\right\} \\
& + P_{y1}\left\{\Lambda_x(x^b_{it} - X^b_t) + \Lambda_z(z_{it} - \bar z_t)\right\} + P_{Y1}\left\{\Lambda_X X^b_t + \Lambda_z \bar z_t + \Lambda_\xi \xi_t\right\} + P_s s_t
\end{aligned}
$$

For our guess to be correct, the above two expressions must coincide in all states of nature, and the following must therefore be true:

$$P_{f0}\Phi_x \;=\; P_x + P_{f1}\Gamma_x + P_{y0}\Lambda_x\Omega_x + P_{y1}\Lambda_x \tag{53}$$

$$(P_{f0} - P_{F0})\Phi_X \;=\; P_{F1}\Gamma_X + P_X + P_{Y0}\Lambda_X\Omega_X + P_{Y1}\Lambda_X \tag{54}$$

$$P_{f0}\Phi_z \;=\; P_{f1}\Gamma_z + P_{y0}\Lambda_x\Omega_z + P_{y1}\Lambda_z \tag{55}$$

$$
\begin{aligned}
(P_{f0} - P_{F0})\Phi_{\bar z} \;=\; & P_{F0}\Phi_z + (P_{F1} - P_{f1})\Gamma_z + P_{F1}\Gamma_{\bar z} + P_{Y0}\Lambda_X(\Omega_z + \Omega_{\bar z}) \\
& - P_{y0}\Lambda_x\Omega_z + (P_{Y1} - P_{y1})\Lambda_z
\end{aligned} \tag{56}
$$

$$(P_{f0} - P_{F0})\Phi_s \;=\; P_{F1}\Gamma_s + P_{Y0}(\Lambda_X\Omega_s + \Lambda_z R) + P_s \tag{57}$$

$$(P_{f0} - P_{F0})\Phi_\xi \;=\; P_{F0}\Gamma_z\Delta Q + P_{F1}\Gamma_\xi + P_{Y0}\left\{\Lambda_X\Omega_\xi + \Lambda_\xi Q\right\} + (P_{Y0} - P_{y0})\Lambda_z\Delta Q + P_{Y1}\Lambda_\xi \tag{58}$$

Recall that the Φ and Ω matrices are themselves transformations of the Γ and Λ matrices. Therefore, the above system is effectively a system of equations in Γ and Λ matrices. This completes Step 2.

Step 3. Steps 1 and 2 resulted in two systems of equations in the Λ and Γ matrices, namely system (47)-(50) and system (53)-(58). We now look at the joint solution of these two systems, which completes our guess-and-verify strategy and gives the sought-after equilibrium policy rules.

First, let as write the solution of the underlying representative-agent model as

$$Y_t \;=\; \Lambda^*_X X^b_t + \Lambda^*_s s_t \qquad \text{and} \qquad X^f_t \;=\; \Gamma^*_X X^b_t + \Gamma^*_s s_t$$

It is straightforward to check that the solution to the beliefs-augmentend model satisfies the following:

$$\Lambda_X = \Lambda^*_X, \qquad \Lambda_z = \Lambda^*_s, \qquad \Gamma_X = \Gamma^*_X, \qquad \text{and} \qquad \Gamma_z + \Gamma_{\bar z} + \Gamma_s = \Gamma^*_s.$$

That is, the solution for the matrices Λ_X, Λ_z, and Γ_X, and for the sum $\bar\Gamma_s \equiv \Gamma_z + \Gamma_{\bar z} + \Gamma_s$, can readily be obtained from the solution of the underlying representative-agent model.

24

With the sum $\bar{\Gamma}_s \equiv \Gamma_z + \Gamma_{\bar{z}} + \Gamma_s$ determined as above, we can next obtain each of its three components as follows. First, Γ_s can be obtained from (57):

$$(P_{f0} - P_{F0})\Phi_s = P_{F1}\Gamma_s + P_{Y0}(\Lambda_X\Omega_s + \Lambda_z R) + P_s$$

Plugging the definition of Φ_s and Ω_s in the above, we have

$$-\underbrace{\{((P_{F0} - P_{f0})\Gamma_X + P_{Y0}\Lambda_X)\, N_F + P_{F1}\}}_{A_S}\Gamma_s = \underbrace{P_s + P_{Y0}(\Lambda_z R + \Lambda_X N_s) + (P_{F0} - P_{f0})(\bar{\Gamma}_s R + \Gamma_X N_s)}_{B_S}$$

and therefore $\Gamma_s = A_S^{-1}B_S$. Next, Γ_z can be obtained from (55). Plugging the definition of Φ_z and Ω_z in this condition, we have

$$\underbrace{((P_{f0}\Gamma_x - P_{y0}\Lambda_x)N_f - P_{f1})}_{A_Z}\Gamma_z = \underbrace{P_{y1}\Lambda_z - (P_{f0}\Gamma_x - P_{y0}\Lambda_x)N_y\Lambda_z}_{B_Z}$$

and therefore $\Gamma_z = A_Z^{-1}B_Z$. Finally, we obtain $\Gamma_{\bar{z}}$ simply from the fact that $\Gamma_{\bar{z}} = \bar{\Gamma}_s - \Gamma_z - \Gamma_s$.

Consider now the matrices Λ_x and Γ_x. These are readily obtained from (47) and (53) once we replace the already-obtained results. It is also straightforward to check that these matrices correspond to the solution of the version of the model that shuts down all kinds of uncertainty but allows for heterogeneity in the backward-looking state variables ("wealth").

To complete our solution, what remains is to determine the matrices Γ_ξ and Λ_ξ. These matrices solve conditions (50) and (58), which we repeat below:

$$\Lambda_\xi = (I - M_{EY})^{-1}\left\{M_F(\Gamma_\xi + \Gamma_{\bar{z}}\Delta) + (M_F - M_f)\Gamma_z\Delta + M_{EY}\Lambda_z\Delta\right\}$$

$$(P_{f0} - P_{F0})\Phi_\xi = P_{F0}\Gamma_z\Delta Q + P_{F1}\Gamma_\xi + P_{Y0}\left\{\Lambda_X\Omega_\xi + \Lambda_\xi Q\right\} + (P_{Y0} - P_{y0})\Lambda_z\Delta Q + P_{Y1}\Lambda_\xi$$

Let us use the first condition to substitute away Λ_ξ from the second, and then the facts that

$$\Omega_\xi = N_Y\Lambda_\xi + N_F\Gamma_\xi$$
$$\Phi_\xi = \Gamma_X(N_Y\Lambda_\xi + N_F\Gamma_\xi) + (\Gamma_\xi + \Gamma_{\bar{z}}\Delta)Q$$

to substitute away also Ω_ξ and Φ_ξ. We then obtain a single equation in Γ_ξ, namely

$$B\Gamma_\xi + A\Gamma_\xi Q + C = 0$$

where

$$A \equiv (P_{F0} - P_{f0}) + P_{Y0}(I - M_{EY})^{-1}M_F$$

$$B \equiv ((P_{F0} - P_{f0})\Gamma_X N_Y + P_{Y0}\Lambda_X N_Y + P_{Y1})(I - M_{EY})^{-1}M_F + (P_{F0} - P_{f0})\Gamma_X N_F + P_{F1} + P_{Y0}\Lambda_X N_F$$

$$C \equiv \left(P_{F0}\Gamma_z\Delta Q + (P_{Y0} - P_{y0})\Lambda_z + (P_{F0} - P_{f0})\Gamma_{\bar{z}} + P_{Y0}(I - M_{EY})^{-1}\left[M_F\Gamma_{\bar{z}} + (M_F - M_f)\Gamma_z + M_{EY}\Lambda_z\right]\right)\Delta Q$$
$$+ ((P_{F0} - P_{f0})\Gamma_X N_Y + P_{Y0}\Lambda_X N_Y + P_{Y1})(I - M_{EY})^{-1}\left[M_F\Gamma_{\bar{z}} + (M_F - M_f)\Gamma_z + M_{EY}\Lambda_z\right]\Delta$$

Note that A, B, and C are determined by primitives, plus some of the coefficients that we have also characterized. The above equation therefore gives us the unique solution for the matrix Γ_ξ as a function of the primitives of the model. Λ_ξ is then readily obtained from (50). This completes the solution.