Quantifying Confidence*

George-Marios Angeletos
MIT

Fabrice Collard
University of Bern

Harris Dellas
University of Bern

December 15, 2016

Abstract

We develop a tractable method for augmenting macroeconomic models with rich dynamics in higher-order beliefs. We use it to accommodate a certain type of waves of optimism and pessimism about the short-term economic outlook, which we interpret as variation in "confidence". 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.

Keywords: Business cycles, strategic uncertainty, higher-order beliefs, confidence, aggregate demand, coordination failure, DSGE models.

*We are grateful to the editor, Lars Hansen, and to three anonymous referees for extensive feedback; and to Stephen Morris, Kristopher Nimark, Morten Ravn, and Juan Rubio-Ramirez for conference discussions of our paper. We have also benefited from the comments of many colleagues and from presentations at Bocconi, BU, Harvard, LSE, MIT, UCL, the University of Zurich, the University of Oslo, the Norwegian Business School, the Minneapolis FRB, the 2013 Hydra Workshop, CRETE 2013, the 2013 Bank of England/CFM conference on Macroeconomics, the 2014 NBER Summer Institute, the 2014 Minnesota Workshop in Macroeconomic Theory, the LSE Conference in Honor of Robert Barro, the 10th Dynare conference, the 2nd Rational Inattention Conference in Oxford, the 11th DSGE Conference in Chicago, and the 8th Banco de Portugal Conference on Monetary Economics. Angeletos would also like to express his gratitude to the University of Zurich, the UBS Center for Economics in Society, and the Excellence Foundation Zurich for their hospitality and support during the 2013-14 academic year, when a significant part of this project was carried out.

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.

There is no consensus on what explains these phenomena: macroeconomists disagree on the significance of technology, nominal rigidities, financial frictions, policy uncertainty, and a variety of other factors. Yet, there is a certain "orthodoxy" regarding the type of theoretical frameworks used to interpret the macroeconomic data and to guide policy: the economy is modeled, in effect, as a game in which all players share a common prior, have complete (homogeneous) information, reach identical beliefs about the current state and the future prospects of the economy, and can perfectly coordinate their current actions and their plans for the future.

In this paper, we depart from this orthodoxy by introducing a friction in the manner in which economic agents form beliefs about the actions of others, thereby letting expectations play an autonomous role in driving business cycles. This departure generates novel insights into the causes of business cycles and leads to novel structural interpretations of the data.

Two contributions. We make two contributions, one methodological and one applied. We first develop a general method for enriching macroeconomic models with a tractable form of aggregate variation in higher-order beliefs.[1] We then use this method to explore the macroeconomic implications of a certain type of waves of optimism and pessimism, which—unlike the one captured by the literature on news and noise shocks—regards the short-term prospects of the economy.

We interpret these waves as variation in "confidence". We quantify their importance within RBC and NK models of either the textbook or the medium-scale DSGE variety.[2] We show that they offer a parsimonious yet potent explanation of multiple salient features of the macroeconomic data. We finally argue that they capture a form of "demand-driven fluctuations" that does not rely on nominal rigidities and can thus help bypass the inflation puzzles of the NK framework.

Background and methodological contribution. We build heavily upon the macroeconomic literature on incomplete information and higher-order uncertainty.[3] This literature goes back at least to Phelps (1971) and Townsend (1983) and has been revived recently by the influential contributions of Morris and Shin (2001, 2002) and Woodford (2002); see Angeletos and Lian (2016a) for a survey. 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".

[1] By "higher-order beliefs" we refer to the beliefs that agents hold about the beliefs of others.

[2] RBC, NK and DSGE are acronyms for, respectively, Real Business Cycles, New Keynesian, and Dynamic Stochastic General Equilibrium.

[3] By "incomplete information" we mean noisy private information about the underlying state of Nature.

We borrow from this literature the key insight that higher-order beliefs can deviate from first-order beliefs and that this deviation can in turn help accommodate autonomous variation in equilibrium expectations. We nevertheless abstract from noisy private information and instead use a certain type of heterogeneous priors to engineer the desired gap between first- and higher-order beliefs. This 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.

To illustrate, consider the textbook 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.[5] 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. Moreover, the new policy rule can be obtained in a relatively straightforward manner.

This gain in tractability is not limited to the textbook RBC model. For an 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.[6]

Applied contribution. Consider the aforementioned variant of the RBC model, which adds our "unorthodox" belief shock (denoted by ξ_t) to the conventional technology shock (denoted by A_t). By construction, this shock introduces transitory variation in the gap between first- and higher-order beliefs. In equilibrium, it gives rise to waves of optimism and pessimism about the short-run 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 employment and income over the same horizon, without any change in expectations of either the exogenous fundamentals or of the endogenous outcomes in the medium to long run.

[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 a detailed exposition, see Nimark (2011) and Huo and Takayama (2015b).

[5] In fact, for a plausible specification of the available exogenous and endogenous signals, Huo and Takayama (2015b) prove that the true state space is infinite.

[6] 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 a persistence of higher-order uncertainty. We use an example in Subsection 6.2 to elucidate this issue, as well as to argue that this issue may not matter for the applied contribution of our paper.

We interpret this kind of belief waves as variation in "confidence" and, accordingly, we refer to ξ_t as the "confidence shock". The applied contribution of our paper centers on the characterization of the observable implications of these waves; on contrasting them to those of other structural shocks and mechanisms employed in the macroeconomic literature; and on offering a novel, parsimonious, yet potent structural interpretation of the business-cycle data.

In spite of its parsimony, the model has no difficulty in matching the key business-cycle moments in the US data. This success is not trivial: the match between the model and the data deteriorates significantly if we replace the confidence shock with any from a variety of other structural shocks that have been employed in the literature as proxies for shifts in either "supply" or "demand", such as technology, news and noise, investment-specific, and discount-rate shocks.

To understand what underlies this success, consider a negative innovation in ξ_t. As already noted, this triggers pessimism about the short-term outlook, without affecting the medium- to long-term prospects. Because firms expect the demand for their products to be relatively weak over the next few quarters, they find it optimal to lower their own demand for labor and capital. As a consequence, households expect to experience 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 co-movement between employment, output, consumption, and investment at the business-cycle frequency, without commensurate movements in labor productivity, TFP, and inflation at *any* frequency.

It is precisely these patterns that best characterize the US data and that competing structural mechanisms—including those found in the literature on news and noise shocks—have difficulty matching.[7] Importantly, our success derives, not from the specific modeling devices we use in order to create waves of beliefs, but rather from the property that these waves regard the short-run outlook as opposed to prospects in the medium to long run. We hope that this finding will guide future work on the causes and the propagation of business cycle.

Seen through the lenses of business-cycle accounting (Chari, Kehoe, and McGrattan, 2007), the confidence shock registers as joint variation in the labor and the capital wedge.[8] These wedges are the observable counterparts of the gaps between the equilibrium expectations predicted by our model and those predicted by the standard RBC model—or, equivalently, of the underlying gap between first- and higher-order beliefs. The cyclical properties of the predicted wedges align well with the measured wedges in the US data, a fact that provides additional support for our model.

[7]The literature on news and noise shocks was spurred by the influential contribution of Beaudry and Portier (2006) and includes, inter alia, Christiano, Ilut, Motto, and Rostagno (2008), Jaimovich and Rebelo (2009), Lorenzoni (2009) and Barsky and Sims (2012). This literature formalizes a type of optimism and pessimism which derives from beliefs about future technology and which, in sharp contrast to the type accommodated in our paper, concerns primarily prospects in the medium to long run.

[8]The former is the wedge between the marginal product of labor and the marginal rate of intra-temporal substitution between consumption and leisure; the latter is the wedge between the marginal product of capital and the marginal rate of inter-temporal substitution in consumption.

Our mechanism also helps capture the notion that recessions are periods of "weak aggregate demand" without the need for either nominal rigidities or frictions in the conduct of monetary policy. Seen through the lenses of the New-Keynesian framework, a drop in confidence indeed registers as an increase in measured markups and in the "output gap" relative to the underlying RBC benchmark. Yet, there are no nominal rigidities at work, no constraints on monetary policy, and no commensurate drop in prices. Unlike what is required in the New-Keynesian model, demand-driven recessions do not have to be deflationary episodes.[9]

The ability of our theory to provide a parsimonious structural interpretation of the data is, at least in our eyes, a desirable quality that is not shared by other theories. Nevertheless, it is also important to assess whether the mechanism remains quantitatively potent once it has been embedded in richer, estimated, DSGE models that contain multiple competing structural shocks. We address this question in the context of two "medium-scale" DSGE models, one with sticky and the other with flexible prices. Both models include permanent and transitory TFP shocks, news shocks, investment- and consumption-specific shocks, fiscal shocks, and monetary shocks. To minimize the risk that our estimates of the relevant business-cycle effects are contaminated by the omission, or mis-specification, of mechanisms that operate primarily in the medium- to long-run, such as demographics, we carry out the estimation with maximum likelihood in the frequency domain, focusing on the band of frequencies that is customarily associated with business-cycle phenomena (6-32 quarters).[10]

In spite of the presence of multiple competing forces, the estimated confidence shock accounts for about one half of GDP volatility at business-cycle frequencies. Furthermore, its observable properties are similar across the two estimated models, and to those in our baseline model. This underscores the robustness of our mechanism across RBC and NK settings, of either the textbook or the medium-scale variety, a property that is not shared by other mechanisms. Last but not least, the model-based confidence shock tracks well empirical counterparts such as the University of Michigan Index of Consumer Sentiment, a fact that lends additional support to our theory.

Layout. The rest of the paper is organized as follows. Section 2 expands on the relation of our paper to the literature. Section 3 sets up the baseline model. Section 4 explains the recursive formulation of the equilibrium and our solution method. Section 5 explores the quantitative performance of the baseline model. Section 7 extends the quantitative analysis to two richer, estimated, models. Section 8 concludes. The Appendix contains the details of the solution method and a number of auxiliary results.

---

[9] The standard structural account of the Great Recession attributes all of the "deficiency in aggregate demand" to the presence of nominal rigidity. This deficiency was allegedly not rectified by monetary policy due to the zero-lower bound. A plausible alternative, which however we will not pursue in this paper, is that a large part of this deficiency is due to a coordination failure of the type formalized in this paper.

[10] The method was initially proposed by Hansen and Sargent (1993) and Sims (1993) in the context of seasonal adjustment. It has been applied to business-cycle research by, inter alia, Christiano and Vigfusson (2003) and Sala (2015).

2 Related literature

The macroeconomic literature on informational frictions has many strands. One strand focuses on the decision-theoretic aspect of how agents collect information or pay attention (Sims, 2003, Reis, 2006, Alvarez, Lippi, and Paciello, 2011). Another strand embeds such aspects in general-equilibrium settings (Mankiw and Reis, 2002, Mackowiak and Wiederholt, 2009). A third strand emphasizes the roles of coordination and higher-order uncertainty (Morris and Shin, 1998, 2002, Woodford, 2002, Angeletos and Pavan, 2007, Angeletos and La'O, 2013). A fourth strand focuses on signal-extraction problems (Amador and Weill, 2010, Benhabib, Wang, and Wen, 2015, Chahrour and Gaballo, 2015). Our paper builds on the third line of work but uses heterogeneous priors as a short-cut to engineer the desired variation in higher-order beliefs.[11]

Prior attempts to advance the quantitative frontier of this literature include Lorenzoni (2009), Nimark (2011), Rondina and Walker (2014), Melosi (2014), Maćkowiak and Wiederholt (2015), David, Hopenhayn, and Venkateswaran (2014), Schaal and Taschereau-Dumouchel (2015), and, especially, Huo and Takayama (2015a,b).[12] We view our methodological contribution as complementary to all these attempts. As already noted, the main advantage of our approach is its flexibility and its straightforward applicability to macroeconomic models. A potential cost is that it bypasses the restrictions that the common-prior assumption, in combination with auxiliary assumptions and/or appropriate evidence, can impose on the size and dynamics of higher-order uncertainty.

By emphasizing the role of coordination, we connect to the literature on coordination failures and sunspot fluctuations (Benhabib and Farmer, 1994, Cass and Shell, 1983, Diamond, 1982, Farmer, 2012). We view our contribution, along with those of Angeletos and La'O (2013), Benhabib, Wang, and Wen (2015), and Huo and Takayama (2015b), as a bridge that extends the *spirit* of that literature to the class of unique-equilibrium DSGE models that dominate modern research.

At a broader level, our paper relates to work that relaxes the rational-expectations equilibrium concept, such as that on robustness and ambiguity (Hansen and Sargent, 2007, Ilut and Schneider, 2014), non-Bayesian learning (Evans and Honkapohja, 2001, Eusepi and Preston, 2011), and eductive stability (Guesnerie, 1992); see also Woodford (2013) and the references therein. Although the mechanism we study and the methods we develop are distinct, we share with these works the desire to enrich the belief dynamics of macroeconomic models. Complementary in this respect is also the recent paper by Ilut and Saijo (2016), which builds, in effect, a bridge between our work and the work on ambiguity aversion.

[11] Other applications of the same short-cut include Allen, Morris, and Postlewaite (1993), Scheinkman and Xiong (2003), Izmalkov and Yildiz (2010), and Section 7 of Angeletos and La'O (2013).

[12] Huo and Takayama (2015a,b) represents a breakthrough in this literature. It shows that a finite-state-space solution is possible for a large class of linear models insofar as the observed signals and the exogenous shocks follow finite ARMA processes. However, finding the right state space can be non-trivial, as it requires the solution of a fixed-point problem in the frequency domain. Moreover, an exact finite-state-space solution is typically impossible when agents observe signals of endogenous outcomes such as output. Their method therefore remains less tractable and less flexible than our approach, especially with regard to the estimation of medium-scale DSGE models.

3 An RBC Prototype with a tractable form of 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 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.[13] 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.

[13] For the economy to be a consistent with a balanced-growth path, it must be that $\gamma = 1$, a restriction that we impose in all our quantitative exercises. We allow $\gamma \neq 1$ only to accommodate an analytical example that helps illustrate certain insights in Sections 4 and 6.

Final-good sector. The final good is produced with a Cobb-Douglas technology, so that $\log Y_t = \int_0^1 \log y_{it}\, \mathrm{d}i$. 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 $P_t = 1$.[14]

Technology shocks. 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 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 models with $\sigma \approx 0$, we only study the model with $\sigma = 0$. This guarantees that 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.

We close the model by letting 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 specification is in line with standard DSGE practice. It also helps mimic, within our heterogeneous-prior framework, the property that the fluctuations sustained by higher-order uncertainty in common-prior settings are necessarily stationary. This is because the common-prior assumption guarantees that the gap between first- and higher-order beliefs must vanish as more information arrives with time.[15]

Remarks and Interpretation. We conclude this section with some remarks on our methodological approach and the interpretation of the ξ_t shock.

1. 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

[14] This applies to the present model. In the monetary variants studied later, P_t be endogenous.

[15] See Subsection 6.2 for an example that illustrates this point.

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 (2011), 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 6.2 by establishing an exact observational equivalence, from the point of view of aggregate data, between a simplified version of our model and a common-prior variant that builds on Angeletos and La'O (2013) and Huo and Takayama (2015b). An exact observational equivalence does not apply to the richer models we seek to quantify in this paper, for a good reason: incomplete-information, common-prior, extensions of these models tend to have large—possibly infinite—state spaces, whereas our heterogeneous-prior formulation features a small state space by design. The example, nevertheless, represents a useful illustration of the key qualitative similarities of the two modeling approaches.

2. The aforementioned example helps illustrate also an important difference between the two approaches. The common-prior assumption imposes certain bounds on the magnitude and the dynamics of higher-order uncertainty as functions of the underlying first-order uncertainty. By contrast, our heterogeneous-prior approach does not impose any such bounds. This indicates a potential trade off between the computational simplicity afforded by our approach and the extra discipline imposed by common-prior specifications.

This trade off, however, need not be significant in practice. We use a back-of-the-envelop exercise in the context of the aforementioned example to explain why. The recent work of Huo and Takayama (2015b) further corroborates this point: it reaches similar quantitative conclusions as our paper while working with common-prior variants of our baseline model. We thus encourage the reader to adopt a flexible interpretation of our modeling approach: the belief waves we accommodate in this paper can be thought either as the product of relaxing the rational-expectation hypothesis, or as a proxy for incomplete information.[16]

3. Notwithstanding the above methodological points, the applied contribution of our paper is ultimately found in the property that the variation in ξ_t maps, in equilibrium, to variation in the expectations that firms and consumers form about the *short run*, as opposed to the medium or long run. This property, which will become evident only after we have characterized the equilibrium of

[16] This point echoes the broader observations made in Angeletos and La'O (2013) and Angeletos and Lian (2016a,b) about the modeling role of incomplete information in macroeconomic models: by relaxing the tight connection between expectations of fundamentals and expectations of outcomes, higher-order uncertainty plays a similar role as the the one played by relaxations of the rational-expectations solution concept.

our model, clarifies the particular notion of "confidence" that we have in mind: it pertains to beliefs about where the economy is expected to be during the following few quarters. As we elaborate in Sections 5–7, the quantitative performance of our mechanism derives entirely from this property.

4. Although we model ξ_t as an exogenous shock, we think of it as a proxy of a mechanism whose deeper micro-foundations we abstract from in order to make progress in understanding its consequences. As already noted, this mechanism regards the uncertainty that firms and consumers form about the short-term outlook of the economy, due to frictions in coordination. In richer settings, the nature of these frictions and the resulting uncertainty can be endogenous to the network of markets or to other social interactions; to the details of how agents collect, digest, and exchange information; and to a variety of forces that influence the formation of expectations.

By the same token, we do not exclude the possibility that a drop in confidence is itself triggered by, say, an adverse financial shock or some other economic event. By treating ξ_t both as exogenous and as independent of any other shock, we only wish to "orthogonalize" its contribution from those of other familiar structural mechanisms. Future work might endogenize the relevant distortions in expectations of the short run, thus recasting our "confidence shocks" as a propagation mechanisms. For a complementary attempt in this direction, see Ilut and Saijo (2016).

4 Equilibrium characterization and solution method

In this section, we characterize the equilibrium of the model and present our solution method. First, we develop a recursive representation, which helps clarify how the ξ_t shock enters the formation of the equilibrium expectations in the model. Next, we consider the log-linear approximation of the equilibrium and review how our solution method works. Finally, we illustrate the method with an example that can be solved by hand and use this example to anticipate certain properties that emerge in the subsequent quantitative analysis. Throughout, we maintain the literal interpretation of ξ_t as a systematic bias in beliefs; appealing re-interpretations are discussed in Section 6.

4.1 Recursive equilibrium

As behavior is forward looking, the optimal choices of any agent (or island) at any given point of time depend on her beliefs, not only about the concurrent behavior of others, but also about their behavior in all future periods. This suggests a high-dimensional fixed-point relation between actual behavior and the expectations that agents form at any given point of 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 rich manner that a low-dimensional recursive representation is not feasible. 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:

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

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 represents the type of aggregate-demand externalities and other general-equilibrium effects that are at the core of DSGE models.

To make his 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) in 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-equation restrictions that discipline the exercises we conduct in this paper: if expectations were "completely" irrational, the beliefs of different endogenous objects would not have to 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 co-movement 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 $\mathbf{P}$, $\mathbf{G}$, V_1, and V_2, such that the following is true:

- $\mathbf{P}(z,\xi,K)$ captures the price (or the terms of trade, or 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:

$$\left.\begin{array}{rcll} V_1(k;z,\xi,K) & = & \max\limits_{n} & V_2(\hat{m};z,\xi,K) - \frac{1}{1+\nu}n^{1+\nu} \\ & & s.t. & \hat{m} = \hat{p}\hat{y} + (1-\delta)k \\ & & & \hat{y} = zk^{\alpha}n^{1-\alpha} \\ & & & \hat{p} = \mathbf{P}(z,\xi,K) \end{array}\right\} \tag{3}$$

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

- $\mathbf{P}$ 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 of each period; m denotes the quantity of the final good held by the island 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 all islands make the same choices in equilibrium and that the function $\mathbf{y}$ 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 $\mathbf{g}$ is the optimal savings rule, this gives condition (6).

Summing up, an equilibrium is given by a fixed point between 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 steady state. This makes it possible to obtain the equilibrium dynamics of the belief-augmented model as a tractable transformation of the equilibrium dynamics of the canonical RBC model.

4.2 Log-linear Solution

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

The terms of trade faced by island i are given by $p_{it} = Y_t - y_{it}$. The associated marginal revenue products of labor and capital are given by, respectively,

$$MRPL_{it} \equiv p_{it} + y_{it} - n_{it} \qquad \text{and} \qquad MRPK_{it} \equiv p_{it} + y_{it} - k_{it}$$

The optimal behavior of island i is thus characterized by the following system:

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

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

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

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

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

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.

Conditions (7)-(11) determine the employment, production, consumption, investment and next-period capital stock of each island as functions of the local expectations of the terms of trade. The following properties of this mapping are worth noting. First, the optimal behavior of an island depends on its beliefs of the behavior of other islands—and thereby on higher-order beliefs—only through its expectations of its terms of trade. Second, because the latter enter the optimality conditions (7) and (8) through the marginal revenue products of labor and capital, varying the expectations of the terms of trade has similar incentive effects as varying the expectations of local TFP. Third, because of the forward-looking nature of the problem, behavior in any given period depends on expectations of the terms of trade, not only in the same period, but also in all future periods; in particular, as we elaborate later, the optimal response of an island to optimism or pessimism about its terms of trade depends crucially on whether the improvement in terms of trade is expected to last for ever or only for few quarters. Finally, the expectations that an island forms about its terms of trade are the same as the expectations that firms and consumers form about income and profitability; the first perspective is useful for understanding the mechanics of the model, the second is preferable for mapping the theory to the real world.

We later use these observations to shed light on the observable implications of the ξ_t shock and to interpret it as a form of optimism and pessimism about the short-run economic outlook. For now, we continue with the description of the solution method.

We propose the following policy rules for employment, consumption, and investment at the island level:

$$n_{it} = \Lambda_k^n (k_{it} - K_t) + \Lambda_K^n K_t + \Lambda_z^n z_{it} + \Lambda_\xi^n \xi_t \tag{12}$$

$$c_{it} = \Gamma_k^c (k_{it} - K_t) + \Gamma_K^c K_t + \Gamma_z^c z_{it} + \Gamma_{\bar{z}}^c \bar{z}_t + \Gamma_a^c A_t + \Gamma_\xi^c \xi_t \tag{13}$$

$$k_{it+1} = \Omega_k^k (k_{it} - K_t) + \Omega_K^k K_t + \Omega_z^k z_{it} + \Omega_{\bar{z}}^k \bar{z}_t + \Omega_a^k A_t + \Omega_\xi^k \xi_t \tag{14}$$

where Λ^n, Γ^c, and Ω^k are coefficients that remain to be determined and where $z_{it}, \bar{z}_t$, and A_t denote, respectively, the individual's signal, the average signal in the population, and the true productivity shock.[17] Here, it is important to note the following. To generate data from the model, we impose $z_{it} = \bar{z}_t = A_t$. To solve for the equilibrium policy rules (or strategies), however, we must treat z_{it}, $\bar{z}_t$ and A_t as distinct objects. Indeed, whereas A_t identifies the true aggregate fundamental, z_{it} and $\bar{z}_t$ are signals that pin down the first and higher-order beliefs of the fundamental; and whereas each agent believes that his own signal coincides with the true value of A_t, he also believes that the average signal of others is perturbed away from A_t by an amount equal to ξ_t. Furthermore, note that the stage-2 policy rules are allowed to depend on the triplet $(z_{it}, \bar{z}_t, A_t)$, because the realizations of these objects are publicly revealed at that point. The stage-1 policy rules, though, are restricted to depend only on the local signal z_{it}, because $\bar{z}_t$ and A_t are not known yet.[18] Finally, note that the individual policy rules are functions, not only of K_t (the aggregate endogenous state variable), but also of k_{it} (the individual endogenous state variable); in equilibrium, however, the realized values of k_{it} and of all other island-specific variables end up coinciding with the realized values of the corresponding aggregates, thanks to the assumed symmetry across the islands.

Substituting (12)-(14) into (7)-(11) results in a system of equations involving the Λ, Γ, and Ω coefficients. The solution to this system gives the equilibrium policy rules at the individual level. Aggregating these rules and letting $\bar{z}_t = A_t$ leads to the following law of motions for aggregate employment, consumption, and capital:

$$N_t = \Lambda_K^n K_t + \Lambda_A^n A_t + \Lambda_\xi^n \xi_t \tag{15}$$

$$C_t = \Gamma_K^c K_t + \Gamma_A^c A_t + \Gamma_\xi^c \xi_t \tag{16}$$

$$K_{t+1} = \Omega_K^k K_t + \Omega_A^k A_t + \Omega_\xi^k \xi_t \tag{17}$$

where $\Lambda_A^n \equiv \Lambda_z^n$, $\Gamma_A^c \equiv \Gamma_z^c + \Gamma_{\bar{z}}^c + \Gamma_a^c$ and $\Omega_A^k \equiv \Omega_z^k + \Omega_{\bar{z}}^k + \Omega_a^k$. Similar expressions can be obtained for any other endogenous variable.[19]

[17] To relate the above policy rules to the general solution developed in Appendix E, note that labor, consumption and capital are elements of the vectors denoted by, respectively, y_{it}, x_{it}^f and x_{it}^b, and similarly Λ^n, Γ^c, and Ω^k are elements of the matrices denoted by, respectively, Λ, Γ and Ω.

[18] 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 (8) and (11), *only* through the realized values of the stage-1 outcomes.

[19] This includes forecast variables, such as the average forecast at date t of aggregate output at dates $\tau \geq t$.

In Appendix E, we show that the equilibrium values of the coefficients $(\Lambda_K^n, \Gamma_K^c, \Omega_K^k)$ and $(\Lambda_A^n, \Gamma_A^c, \Omega_A^k)$ are the same as those in the standard RBC model. This is due to the facts that our model reduces to the standard RBC model when the ξ_t shock is shut down and the ξ_t shock enters in an additive way in the equilibrium system. It follows that any difference in the empirical properties between our and the RBC model is due to the coefficients $(\Lambda_\xi^n, \Gamma_\xi^c, \Omega_\xi^k)$. These coefficients can in turn be solved as functions of the coefficients that characterize the aggregate-level policy rules of the standard model; and certain coefficients that appear in the individual-level equilibrium conditions and that measure, in effect, the strategic complementarity in the economy.

Together with $(\Lambda_K^n, \Gamma_K^c, \Omega_K^k)$, the aforementioned coefficients characterize the impulse responses of the macroeconomic quantities to the confidence shocks and epitomize the "cross-equations restrictions" in the model. This raises the following questions: What are these restrictions? How do they differ from those imposed by the technology or other structural shocks? And are they in line with important regularities in the data?

We answer these questions, first with the help of a simplified version of the model that admits a closed-form solution (Subsection 4.3) and then with a calibrated version that is evaluated quantitatively (Section 5). For now, we wish to emphasize that our solution method is not limited to the simple model we have studied so far: it extends to essentially any linear DSGE model. The details are spelled out in Appendix E. The bottom line is that we obtain the solution of the belief-augmented extension of any linear DSGE model as a tractable transformation of the solution of the original model. The belief-augmented model can thus be simulated and estimated with the same ease as the original model.

4.3 A Closed-form Example

To gain further insight into how our solution method works and how ξ_t matters, we now consider an example that can be solved by hand: we let the utility be linear in consumption ($\gamma = 0$).

In this case, condition (7), which describes optimal labor supply, reduces to the following:

$$\nu n_{it} = \mathbb{E}_{it}[p_{it} + y_{it}] - n_{it}$$

Using $p_{it} = Y_t - y_{it}$ and $Y_t = A_t + \alpha K_t + (1-\alpha) N_t$, we can restate the above as follows:

$$n_{it} = \mathbb{E}_{it}[\theta_t + \omega N_t], \tag{18}$$

where $\theta_t \equiv \frac{1}{1+\nu}(A_t + \alpha K_t)$ and $\omega \equiv \frac{1-\alpha}{1+\nu} \in (0,1)$. Equilibrium employment can therefore be understood as the solution to a static beauty-contest game, of the type found in Morris and Shin (2002), Angeletos and Pavan (2007), and Bergemann and Morris (2013). In this game, a player is an island, her action is local employment, the fundamental is θ_t, and the degree of strategic complementarity is ω. 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).

Solving (18) is straightforward. Start by guessing a policy rule as in (12). Aggregation gives $N_t = \Lambda_K^n K_t + \Lambda_z^n \bar{z}_t + \Lambda_\xi^n \xi_t$. 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 $\mathbb{E}_{it}[N_t] = \Lambda_K^n K_t + \Lambda_z^n \bar{z}_t + (\Lambda_\xi^n + \Lambda_z^n)\xi_t$. Using this fact in (18), we infer that, whenever i expects the others to play according to the rule given by (12), his best response is to set

$$n_{it} = \left(\tfrac{1}{1+\nu} + \tfrac{1-\alpha}{1+\nu}\Lambda_z^n\right) z_{it} + \left(\tfrac{\alpha}{1+\nu} + \tfrac{1-\alpha}{1+\nu}\Lambda_K^n\right) K_{it} + \tfrac{1-\alpha}{1+\nu}(\Lambda_\xi^n + \Lambda_z^n)\xi_t.$$

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:

$$\Lambda_k^n = 0, \quad \Lambda_K^n = \frac{\alpha}{\alpha+\nu}, \quad \Lambda_z^n = \frac{1}{\alpha+\nu}, \quad \text{and} \quad \Lambda_\xi^n = \frac{1-\alpha}{(\alpha+\nu)^2}. \tag{19}$$

The equilibrium policy rule for capital is, in turn, obtained using the Euler condition, namely condition (8). With $\gamma = 0$, this condition reduces to

$$k_{i,t+1} = \mathbb{E}'_{it} Y_{t+1} = \mathbb{E}'_{it}[A_{t+1} + \alpha K_{t+1} + (1-\alpha)N_{t+1}].$$

Following similar steps as above, and using the obtained policy rule for employment, we can verify that the policy rule for capital is given by (14) with

$$\Omega_k^k = \Omega_K^k = \Omega_z^k = \Omega_{\bar{z}}^k = 0, \quad \Omega_a^k = \frac{1+\nu}{\nu(1-\alpha)} \quad \text{and} \quad \Omega_\xi^k = \frac{(1+\nu)}{\nu(\alpha+\nu)}\rho. \tag{20}$$

Note that Ω_ξ^k is positive whenever $\rho > 0$, which underscores how a persistent confidence shock can drive investment by shaping expectations of future aggregate demand and future terms of trade.

This completes the characterization of the equilibrium policy rules. The solution strategy followed above mirrors the one followed in Appendix E for the more general case. One important difference, however, is that the relevant "beauty contest" is dynamic rather than static. Accordingly, whereas in the present example the Λ coefficients could be solved separately from the Ω coefficients, in the general case all the relevant coefficients have to be solved simultaneously.

Let us now use the present example to anticipate some properties that emerge also in our subsequent quantitative explorations.[20]

1. The coefficient Λ_ξ^n, which measures the impact effect of the confidence shock on employment, is tied to ω, the degree of strategic complementarity. The latter increases with the Frisch elasticity and with the elasticity of production with respect to labor. This is because these parameters determine the degree of strategic complementarity, which in turn controls the significance of higher-order beliefs.

[20] Note that, to obtain artificial data from the model, we merely need to aggregate the policy rules and impose $\bar{z}_t = A_t$. This gives aggregate employment and output as in conditions (15) and (17), with $\Lambda_A^n = \Lambda_z^n$, $\Omega_A^k = \Omega_a^k$, and $(\Lambda_K^n, \Lambda_z^n, \Lambda_\xi^n, \Omega_K^k, \Omega_a^k, \Omega_\xi^k)$ determined as in (19) and (20). The solution for aggregate consumption then follows directly from the resource constraint and the production function.

2. The coefficient Ω_ξ^k, which measures the impact effect of the confidence shock on investment, is positive if and only $\rho > 0$. This is because investment is driven by expectations of next period's return to capital, which in turn are driven by expectations of next period's aggregate economic activity, which in turn vary with "confidence" if and only if the latter is persistent.

3. For any given $\rho > 0$, the comparative statics of Ω_ξ^k with respect to ν and α have the same sign as those of Λ_ξ^n.

4. The coefficient Γ_ξ^c, which measures the impact effect of the confidence shock on consumption, is decreasing in ρ and is strictly positive if and only if $\rho < \bar{\rho}$, for some $\bar{\rho} \in (0,1)$.[21] This is because a higher ρ increases the belief-driven volatility in investment without increasing the corresponding volatility of employment and output.

5. A direct implication of the above properties is that $\rho \in (0, \bar{\rho})$ is necessary and sufficient for the confidence shock to trigger positive comovement in all the macroeconomic quantities.

These properties illustrate the so-called cross-equation restrictions that the theory imposes on the joint response of the different endogenous variables to the confidence shock. Once we let $\gamma > 0$, it becomes desirable to smooth consumption over time. As a result, labor supply becomes intertwined with consumption-saving behavior. This changes the aforementioned cross-equation restrictions, thus also changing the observable implications of the model. Nevertheless, as shown in the next section, the qualitative properties described above survive subject to the following modifications. First, investment responds to the confidence shock even if $\rho = 0$, because the desire to smooth consumption justifies positive saving out of any transient movement in income, including those triggered by i.i.d confidence shocks. And second, the bound $\bar{\rho}$ is increased, because the desire to smooth consumption dampens the response of investment to persistent confidence shocks.

5 Quantitative Evaluation: A First Pass

Our applied contribution centers on the assessment of the observable implications of the confidence shock and on their comparison to those of other structural shocks that have been proposed in the literature. We start this assessment by studying a calibrated version of our baseline model.

5.1 Calibration

Table 1 reports the values of the parameters of the model. The preference and technology parameters take conventional values: the discount factor is 0.99; the elasticity of intertemporal substitution

[21] Given the solution for employment and capital, it is straightforward to verify that aggregate consumption satisfies (16) with $\Gamma_\xi^c = \frac{\delta\nu(1-\alpha)^2 - s(1+\nu)(\alpha+\nu)\rho}{\delta\nu(1-s)(\alpha+\nu)^2}$, where $s \equiv \frac{\alpha\beta\delta}{1-\beta(1-\delta)} \in (0,1)$ is the steady-state ratio of investment to income. It follows that $\bar{\rho} \equiv \frac{\delta\nu(1-\alpha)^2}{s(1+\nu)(\alpha+\nu)} \in (0,1)$.

is 1; the Frisch elasticity of labor supply is 2; the capital share in production is 0.3; and the depreciation rate is 0.015. These values ensure that our quantitative exercise is directly comparable to those in the extant literature and also that the steady state values of our model are consistent with the long-run patterns in the data.

Table 1: Parameters

Parameter	Role	Value
β	Discount Factor	0.99
γ	Inverse Elasticity of Intertemporal Substitution	1.00
ν	Inverse Elasticity of Labor Supply	0.50
α	Capital Share in Production	0.30
δ	Depreciation Rate	0.025
σ_a	St.Dev. of Technology Shock	0.79
σ_ξ	St.Dev. of Confidence Shock	5.72
ρ	Persistence of Confidence Shock	0.75

Since the technology shock follows a random walk, three parameters remain to be chosen in order to complete the parametrization of the model: the standard deviations of the two shocks and the persistence of the confidence shock. For the latter, we set $\rho = 0.75$. This choice is somewhat arbitrary, but it is motivated by the following considerations. In our setting, ρ pins down the persistence of the deviations between first- and higher-order beliefs. In common-prior settings[22], such deviations may be quite persistent if learning is slow but cannot last for ever. $\rho = 0.75$ implies a half-life of these deviations less than 2.5 quarters.[23]

The standard deviations of the two shocks, σ_a and σ_ξ, are set so as to minimize the distance between the volatilities of output, consumption, investment and hours generated by the model and their counterparts in the data.[24] This yields $\sigma_a = 0.788$ and $\sigma_\xi = 5.721$. We do not take a position now on whether this value for σ_ξ is plausible or not and how it compares to that in the common-prior version of the model. These issues are discussed in Section 7 and in Appendix D. In any case, our calibration strategy is both standard and useful as it facilitates the comparison of our mechanism to competing structural mechanisms in the literature. Furthermore, although the predicted *magnitudes* are of course sensitive to the chosen values for σ_a and σ_ξ, the shapes of the IRFs and hence the predicted *co-movement* patterns are invariant to these values.

[22]We discuss the mapping between the heterogeneous and the common prior specifications in detail in section 6.2.

[23]This is broadly consistent with the estimates obtained in Coibion and Gorodnichenko (2012) of the persistence of average forecast errors in surveys of economic forecasts. Furthermore, 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 (2016), our parametrization of ρ is consistent with the evidence in those papers as well.

[24]By "volatilities" we always refer to the Bandpass-filtered variances at frequencies corresponding to 6-32 quarters. Also, in the minimization objective, each of the model-based volatilities is weighted by the precision of its estimator.

5.2 Business-Cycle Moments

We now evaluate the empirical fit of the calibrated model using standard business-cycle statics. Table 2 reports key moments in the US data (column 1), in our model (column 2), and in four competing models (columns 3-6). Each of the competing models replaces the confidence shock with one of the following, popular structural shocks: a news shock; a discount-rate shock; an investment-specific shock; and a transitory TFP shock. To ensure a fair horserace, the parametrization of the competing models is done similarly to our model; see Appendix C for details. Finally, all the moments are computed on bandpass-filtered series, at frequencies 6-32 quarters.[25]

Table 2: Bandpass-filtered Moments

			Alternative Two-Shock Models			
	Data	Our Model	TFP	Invt	Disc	News
Standard deviations						
stddev(y)	1.41	1.52	1.66	1.21	1.21	1.52
stddev(h)	1.55	1.52	0.84	0.98	0.98	0.97
stddev(c)	0.76	0.69	0.62	0.99	0.98	0.66
stddev(i)	5.12	5.00	6.25	6.36	6.39	6.75
stddev(y/h)	0.76	0.91	0.91	0.81	0.81	0.76
Correlations						
corr(c, y)	0.85	0.89	0.75	0.12	0.11	0.30
corr(i, y)	0.94	0.97	0.97	0.82	0.82	0.94
corr(h, y)	0.87	0.82	0.95	0.74	0.75	0.90
corr(c, h)	0.83	0.49	0.50	-0.57	-0.57	-0.13
corr(i, h)	0.82	0.92	0.99	0.99	0.99	0.99
corr(c, i)	0.73	0.76	0.57	-0.46	-0.46	-0.04
Correlations with productivity						
corr($y, y/h$)	0.07	0.29	0.95	0.59	0.59	0.83
corr($h, y/h$)	-0.43	-0.30	0.80	-0.09	-0.09	0.52
corr(y, sr)	0.82	0.73	0.99	0.84	0.84	0.93
corr(h, sr)	0.47	0.22	0.90	0.29	0.29	0.69

Note: The first column reports moments in the US data, bandpass-filtered, over the 1960-2007 period. The second column reports the moments in our model. The rest of the table reports the moments in the four competing models discussed in the text.

Our model does a very good job in matching the relevant moments in the data. The only notable shortcoming is that it underestimates the correlation of consumption with hours. As explained in detail in Appendix B, the overall fit owes to a delicate balance between the contributions of the

[25]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 (Stock and Watson, 1999). That said, the picture that emerges from Table 2 is invariant to the choice of filter.

technology and the confidence shock. If we shut either shock down, the model fails to match many moments. But once the two shocks are combined, all the moments fall in place.[26]

This success is not a trivial consequence of adding a second shock to the RBC prototype: none of the competing two-shock models is able to replicate the empirical fit of our model. The main reason is that all these competing shocks induce negative co-movement between consumption and investment and/or between consumption and employment. By trying to attribute the variation in the data to structural shocks that do not exhibit the appropriate co-movement patterns, these competing models end up missing, not only certain key correlations, but also the relative volatilities of the macroeconomic quantities.

Appendix C shows that a similar property characterizes baseline versions of the NK framework. Unless one augments that framework with appropriate bells and whistles, the aforementioned shocks fail to generate the co-movement patterns seen in the data, and the associated narratives—such as that a drop in consumer spending can trigger a recession—do not go through. By contrast, our mechanism has no difficulty in generating all the salient features of the data, namely the strong pro-cyclical movements in employment, investment, and consumption without commensurate procyclical movements in TFP and labor productivity. We discuss the sources of this success in the next subsection.[27]

5.3 The Confidence Shock: Observable Implications and Mechanism

Figure 1 reveals the co-movement patterns associated with our mechanism by reporting the impulse response functions (IRFs) of the model's key variables to a positive innovation in ξ_t. The IRFs to the technology shock are the same as in the standard RBC model and are thus omitted.

As is evident from this figure, a positive ξ_t shock triggers a transitory boom in output, consumption, hours, and investment, without a commensurate increase in labor productivity. These co-movement patterns encapsulate the key testable implications of our theory. As already indicated by Table 2, these patterns are not easily generated by other structural mechanisms in the literature. Moreover, they do not hinge on the *magnitude* of higher-order uncertainty: the shape of the IRFs does not depend on the chosen values for σ_a and σ_ξ.

[26]The literature sometimes specifies the technology shock as an AR(1) process with an autocorrelation coefficient around 0.95, instead of a random walk. Adopting this alternative value has only negligible effects in the moments reported in Table 2. In fact, it slightly improves the moments that relate to labor productivity.

[27]The horserace conducted in Table 2 does not include uncertainty shocks, because such shocks require richer, heterogenous-agent, models that are beyond the scope of our paper. Nevertheless, it is worth noting that Bloom (2009) and Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2012) face essentially the same challenge as the one faced by the standard RBC model. Namely, they can generate realistic co-movement patterns between employment, consumption and investment only by letting uncertainty shocks induce strong pro-cyclical movements in aggregate TFP. But this is at odds with the low-to-negative correlations of hours with output and labor productivity observed in the US data (see Table 2 here), as well as with the complementary evidence we provide in a companion paper (Angeletos, Collard, and Dellas, 2016). An interesting alternative, which however we do not pursue here, is that uncertainty shocks matter by inducing variation in confidence as opposed to TFP.

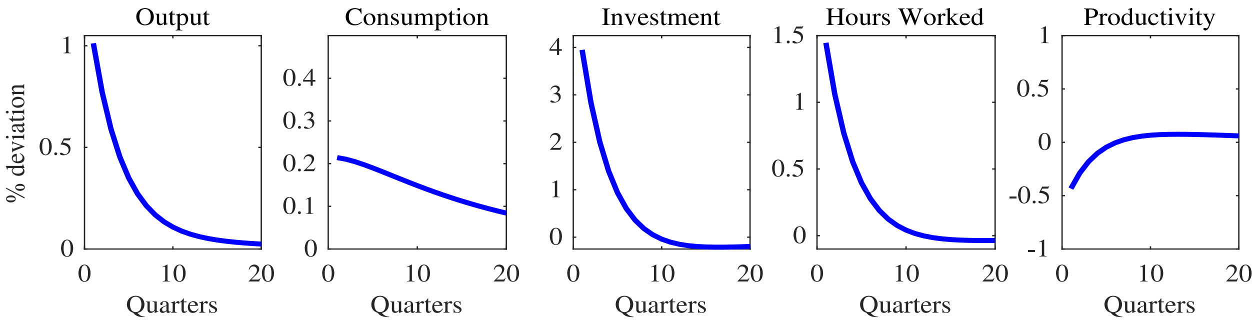

Figure 1: Impulse responses to a positive confidence shock

Let us now explain what drives these patterns. By construction, a positive ξ_t shock causes each island to believe that *other* islands believe that aggregate productivity has gone up. In equilibrium, this induces every island to become optimistic about aggregate economic activity and, thereby, about its own terms of trade. This optimism is short-lived, because ξ_t quickly mean-reverts to zero. Most importantly, this optimism is of a specific type: it causes firms and consumers to upgrade their short-term outlook, *without* modifying their expectations of the medium to long run.

We illustrate this crucial property in Figure 2 by reporting the "term structure of expectations" of the terms of trade at the moment of a positive confidence shock. More specifically, we fix t, shock the economy with a one-standard-deviation increase in ξ_t, and draw the cross-sectional average of $\mathbb{E}_{i\tau}[p_{i,\tau+k}]$ against $k \in \{1, \ldots, 12\}$, for $\tau = t$ (solid line) and for $\tau = t+4$ (dashed line). That is, we draw the forecasts of the terms of trade at different horizons k, both on impact (solid line) and one year after the shock (dashed line). As is evident the figure, the positive innovation in ξ_t causes optimism only about the short-term outlook: every island expects its prospects to brighten for a few quarters and soon after to converge to the steady-state level. As time passes, the optimism fades away and the curve in the figure shifts down. Importantly, however, the curve remains downward-sloping, underscoring once again the fact that the waves of optimism (and pessimism) accommodated in this paper regard exclusively the short-term economic outlook.

Figure 2 focuses on the forecasts of the terms of trade because these are the forecasts that matter for behavior when the economy is represented as a game among the various islands. Nevertheless, the same picture emerges when we consider the forecasts of aggregate output, consumption, investment, and employment: they are all expected to boom but only in the short run.

In the eyes of the firms, a positive ξ shock means an increase in the expected demand for their product—or, equivalently, an improvement in the expected terms of trade. To take advantage of this, the firms raise their demand for both labor and capital, pushing wages and capital returns up. In the eyes of the households, the same shock means a modest increase in their expected permanent income. This helps stimulate consumption while also limiting the wealth effect on labor supply. Importantly, because the households understand that the shifts in the prices of their inputs are

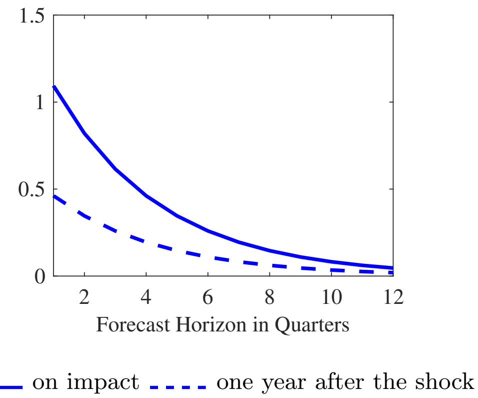

Figure 2: Forecasts of terms of trade at different horizons, following a confidence shock

transitory, the relevant substitution effects dominate the competing wealth effects, guaranteeing that employment and investment increase in equilibrium. All in all, the shock therefore causes employment, output, consumption and investment to increase in tandem—which completes the explanation of the co-movement patterns seen in Figure 1.

5.4 Wedges

Suppose that a macroeconomist is presented with data on aggregate quantities and prices that are generated by the equilibrium of our model. Suppose further that the macroeconomist seeks to interpret these data through the lenses of the standard RBC model, along the lines suggested by Chari, Kehoe, and McGrattan (2007). The confidence shock will manifests itself as a combination of wedges in the optimality conditions of the representative household and the representative firm in the standard RBC model.

Let us elaborate. 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 $MRSC_{t,t+1} \equiv \gamma\left(C_{t+1} - C_t\right)$ 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:

$$MRSN_t = w_t - \tau_t^{nh} \qquad \mathbb{E}_t[MRSC_{t,t+1}] = (1-\beta(1-\delta))(R_t - \tau_t^{kh}) \tag{21}$$

$$MPL_t = w_t + \tau_t^{nf} \qquad \mathbb{E}_t[MPK_{t+1}] = R_t + \tau_t^{kf}. \tag{22}$$

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.

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 plain-vanilla RBC model be perturbed by a different structural. 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.[28]

The predictions of the model about the wedges are related to the IRFs reviewed earlier. In particular, $\xi_t > 0$ 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.

How do these predictions relate to the US data? The answer to this question is complicated because of the difficulty of measuring the relevant prices (the wage and the expected returns to capital). To bypass this complication, we focus on the total labor wedge, defined as the sum $\tau_t^n \equiv \tau_t^{nh} + \tau^{nf}$, and the total capital wedge, defined as the sum $\tau_t^k \equiv \tau_t^{kh} + \tau_t^{kf}$.[29]

Our model predicts a countercylical labor wedge and a procyclical capital wedge. We first explain why this is the case and then compare the model's predictions to the data.

As already noted, 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 τ_t^n 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.

The model's predictions about the total wedges appear to be consistent with the data. This is shown in Table 3, which compares the second moments of the labor and capital wedges predicted by the model to those recovered from the US data using the business-cycle-accounting method of Chari, Kehoe, and McGrattan (2007); see Appendix B for details. For completeness, the table also reports the moments of the efficiency wedge; the latter is the technology shock in the model and, essentially, the Solow residual in the data. The model matches well both the volatilities of all three wedges and their correlations with output.

[28] With our preferred calibration, $\tau_t^{nh} = 0.0152\xi_t$, $\tau_t^{kh} = 0.3277\xi_t$, $\tau_t^{nf} = -0.2548\xi_t$, and $\tau_t^{kf} = -0.1911\xi_t$.

[29] Considering the predictions of our theory for the total wedges seems preferable also for the following reason: these predictions are robust to the timing-protocol issue we discuss in Subsection 6.4, as well as to alternative market structures that implement the same allocations with different prices and therefore also with different splits of the total wages between the firms and the households.

Table 3: Second Moments for Wedges (Bandpass filtered)

	Standard Deviation			Correlation with Output		
	Efficiency	Labor	Capital	Efficiency	Labor	Capital
Data	0.86	1.40	1.04	0.78	-0.57	0.91
Model	0.89	1.37	0.78	0.74	-0.67	0.67

This finding provides additional support for our baseline model. It also illustrates the following points. If we take any standard DSGE model, whether of the RBC or the NK type, and add to it higher-order uncertainty, the latter will naturally manifest itself as a wedge in *all* the optimality conditions of the original model that involve expectations of economic outcomes (or, equivalently, of the actions of others). Importantly, these wedges emerge because, and only because, the agents in the model use a "distorted" expectation operator relative to the complete-information, common-prior, fully-rational benchmark. By the same token, the wedge that appears in one condition is *not* independent from the wedge that appears in another condition. Instead, all the wedges are functions of the same underlying variation in higher-order beliefs. These functions are themselves determined by the equilibrium conditions of the model and are tied to the degree of strategic complementarity that is embedded in the model. For instance, if we shut down the trade across the islands in the baseline model, strategic complementarity vanishes, and so do the wedges. These points highlight the broader sense in which our approach can offer a parsimonious theory of wedges.

We close this section with the following clarification. Insofar as we maintain the standard specification of preferences and technologies, a labor wedge is possible only if there is some asynchronicity in the determination of employment, production, and consumption-investment within each period: without such an asynchronicity, the intra-period optimality conditions of the firms and the households would no more contain an expectation operator, and the related wedges would thus vanish. This explains why we need two stages within each period. However, as explained in Subsection 6.4, our findings do not appear to hinge on whether it is "supply" (employment and production) or "demand" (consumption and investment) that is determined in the first stage. Last but not least, the two-stage specification should not be taken literally: it is only meant to capture the realistic property that labor-market outcomes depend on expectations of the short-term economic outlook.

6 Discussion

In this section, we elaborate on a number of important characteristics of the fluctuations formalized in this paper. More specifically, we argue that they can be thought of as the product of miscoordination; they do not necessarily require a departure from rationality; they capture a type of optimism and pessimism that is distinct from that found in the news-and-noise literature; and they help accommodate the Keynesian view of recessions with nominal rigidity.

6.1 Confusion and Mis-coordination

The mechanism in our paper resembles the one in Lucas (1972) in the following respect. During a recession, each island expects a deterioration in its terms of trade, which motivates the local firms to produce less and the local workers to work less. Yet, as all islands end up reducing their production in unison, the expected deterioration in the terms of trade fails to materialize. In this sense, economic agents suffer from misperceptions of relative prices. In the present paper, these misperceptions reflect systematic biases in beliefs; in Angeletos and La'O (2013) they are the product of rational confusion.

Notwithstanding this similarity, the recession triggered by a negative confidence shock is not merely the product of individual-level confusion. It is also—and indeed to a large extent—the product of mis-coordination. By this we mean the following. If we let one island have perfect knowledge of its actual terms of trade, this island will still find it optimal to produce less in response to a negative ξ shock, simply because it knows that the pessimism of the others reduces the demand for its own good. Consequently, whenever agents expect the economy to be in recession, an actual recession *does* materialize, albeit of a lower size than expected.

The phenomena we have in mind are then *self-reinforcing* at first and *self-correcting* later on. The self-reinforcing part is driven by the strategic complementarity among the islands, a feature that helps capture the interdependence of economic decisions in the real world; the self-correcting part comes from learning, a feature that is proxied here by the assumption that ξ_t is mean-reverting towards zero. We refer the reader to Section 6 of Angeletos and La'O (2013) for a formalization of these ideas within a common-prior setting that allows the belief wave to gain force at first, as the rational confusion propagates from one part of the economy to another, and to be self corrected later, as agents eventually reach common knowledge of one anther's actions.

6.2 Heterogeneous vs Common Priors

We now elaborate on our earlier claim that our heterogeneous-prior specification can also be thought as a proxy for higher-order uncertainty in common-prior settings. To that purpose, we consider two economies. The one a special case of our model. The other is a common-prior variant that builds on Angeletos and La'O (2013) and Huo and Takayama (2015b) and, importantly, can be solved analytically. This permits us to establish two results. The first substantiates the aforementioned claim by showing that the two economies are observationally equivalent, in a sense that is made precise below. The second sheds light on the bounds that the common-prior assumption imposes on the volatility and the persistence of fluctuations in beliefs and economic activity, thus helping evaluate whether our quantitative findings rely on "abusing" the flexibility afforded by our heterogeneous-prior approach.

The heterogeneous-prior economy is the same as the example studied in Subsection 4.3, except that we set $\alpha = 0$, that is, we remove investment. This renders the beliefs formed in the "afternoon" of each period irrelevant, leaving only the beliefs that are formed in the "morning" and that drive employment and output. The common-prior variant, on the other hand, 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.

The details of the common-prior economy are as follows. The TFP in island i is given by $A_{it} = A_t + a_i$, where A_t is the period-t aggregate TFP shock and a_i is an island-specific TFP component. 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.[30] As in Angeletos and La'O (2013), the modeling role of this signal structure is to introduce aggregate variation in higher-order beliefs: the $\tilde{\xi}_t$ shock plays the same role as our confidence shock.

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.[31] Let Y_t^* denote the level of output in that benchmark; it is straightforward to show that $Y_t^* = \frac{1}{1-\omega} A_t$ where $\omega \equiv \frac{1}{1+\nu}$. Let Y_t be the level of output obtained in the presence of the aforementioned shocks to higher-order beliefs. The "output gap" relative to the RBC benchmark is then $Y_t - Y_t^*$. In either economy, this gap is a sufficient statistic for all the aggregate quantities. The two economies are therefore observationally equivalent at the aggregate level if and only if they share the same equilibrium process for the output gap. We next investigate the conditions under which this occurs.

Consider the heterogeneous-prior economy. From the results of Section 4 (and setting $\alpha = 0$), we have that the equilibrium level of output is given by $Y_t = \frac{1}{1-\omega} A_t + \frac{\omega}{(1-\omega)^2}\xi_t = Y_t^* + \frac{\omega}{(1-\omega)^2}\xi_t$.

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

[31]The fact that one economy has cross-sectional heterogeneity in TFP whereas the other does not, is inconsequential for aggregate outcomes.

Since ξ_t is an AR(1) process, we conclude that the output gap is also an AR(1) process:

$$Y_t - Y_t^* = \varphi(Y_{t-1} - Y_{t-1}^*) + \psi\varepsilon_t, \tag{23}$$

where $\varepsilon_t \rightsquigarrow \mathcal{N}(0,1)$ is i.i.d. over time and independent of the technology shock,[32] and where

$$\varphi = \rho \qquad \text{and} \qquad \psi = \frac{\omega\sigma_\xi}{(1-\omega)^2}. \tag{24}$$

Consider now the common-prior economy. Its solution is far from trivial, but can be obtained by adapting Theorem 1 in Huo and Takayama (2015b).[33] We thus have that the output gap in this economy also follows an AR(1) process as in (23), except that now

$$\left.\begin{array}{lcl} \varphi &=& \Phi\left(\tilde{\theta},\omega\right) \equiv \frac{1}{2}\left[\left(\frac{1}{\tilde{\rho}} + \tilde{\rho} + \frac{1-\omega}{\tilde{\rho}}\frac{\tilde{\sigma}_a^2+\tilde{\sigma}_u^2}{\tilde{\sigma}_a^2\tilde{\sigma}_u^2}\tilde{\sigma}_\xi^2\right) - \sqrt{\left(\frac{1}{\tilde{\rho}} + \tilde{\rho} + \frac{1-\omega}{\tilde{\rho}}\frac{\tilde{\sigma}_a^2+\tilde{\sigma}_u^2}{\tilde{\sigma}_a^2\tilde{\sigma}_u^2}\tilde{\sigma}_\xi^2\right)^2 - 4}\,\right] \\ \psi &=& \Psi\left(\tilde{\theta},\omega\right) \equiv \dfrac{\omega\Phi(\tilde{\theta},\omega)}{\tilde{\rho}\left(1-\omega^2\frac{\tilde{\rho}\tilde{\sigma}_a^2+\Phi(\tilde{\theta},\omega)\tilde{\sigma}_u^2}{\tilde{\rho}\tilde{\sigma}_a^2+\tilde{\rho}\tilde{\sigma}_u^2}\right)}\tilde{\sigma}_a \end{array}\right\} \tag{25}$$

where $\tilde{\theta} \equiv (\tilde{\rho}, \tilde{\sigma}_\xi, \tilde{\sigma}_u, \tilde{\sigma}_a)$.

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. By comparing (24) and (25), we can readily prove that the two economies are observationally equivalent in the following sense.

Proposition 1 *(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 (reps., 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 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, 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

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

[33] That result provides an analytic solution for aggregate output in the case without aggregate TFP shocks. By adding an aggregate TFP shock but also assuming that that this shock is common knowledge, we guarantee that the same solution applies to the output gap of the common-prior economy under consideration.

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 (23) 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 belief-driven fluctuations in the common-prior economy are necessarily stationary. This would be true even if we had 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.[34]

Corollary 1 converts the above properties into restrictions on the parameters in 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 stationary. Part (ii), on the other, provides an upper bound on σ_ξ.

Now suppose that we had data on the output gap that allowed us to estimate the AR(1) process described in (23).[35] And also that we possessed information on the value for $\tilde{\sigma}_a$, perhaps from micro-economic observations. The common prior model could not be consistent with the data if the estimated values of ϕ and ψ violated the bounds in Proposition 2. By contrast, a heterogeneous-prior specification would not have this problem. Consequently, a key question is

[34]The bound is also a function of φ, because the presence of Bayesian learning ties the persistence of the fluctuations in higher-order beliefs to their volatility; and it is a function of ω, because stronger complementarity allows the same stochastic process in higher-order beliefs to generate both larger and more persistent fluctuations in output.

[35]Alternatively, think of the researcher estimating φ and ψ jointly with all the other parameters, as a part of a structural-estimation exercise.

whether the quantitative success of our mechanism hinges on the freedom to bypass the bounds that a realistic common-prior specification would have imposed.

Answering this question requires one to take a stand on what qualifies as "realistic". This is beyond the scope of this paper. A first pass, however, is provided by Figure 3.

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.[36] 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 calibrate the present model in a way similar to that in Section 5, namely by fixing $\rho_\xi = 0.75$ and setting 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. As is evident from the figure, the calibrated economy lies comfortably within the set of values that can be rationalized by a common-prior variant under a plausible value for $\tilde{\sigma}_a$.

Figure 3: The bounds on persistence and volatility

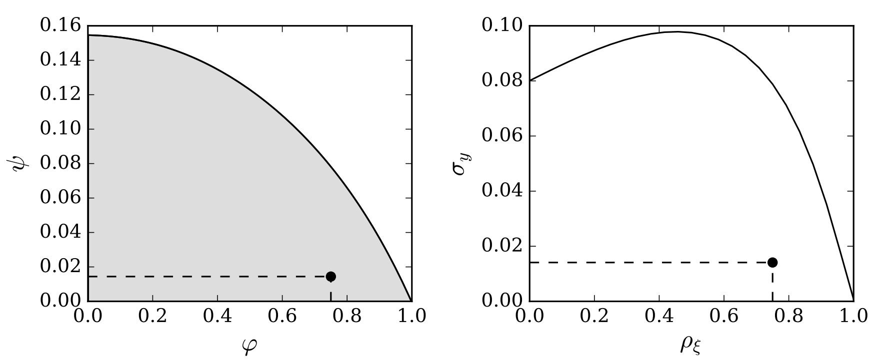

Note that the relevant bound would still be satisfied even if we were to reduce $\tilde{\sigma}_a$ to 4%. Such a value would not appear implausibly large even if we confined first-order uncertainty to concern aggregate fundamentals. For instance, this value is only about twice as large as the standard deviation of the quarterly innovations in the aggregate Solow residual. Furthermore, as 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

[36] 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.

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. As we explain in more detail below, the key to the quantitative performance of our mechanism—and hence also the core of our applied contribution—is that it captures variation in expectations of the short run, as opposed to expectations of the medium to the long run. Whether this variation is the product of rational confusion or irrational "animal spirits" is an important research question with obvious policy ramifications, but it is not a question that impacts on the quantitative success of our model.

6.3 Expectations about the Short Run vs the Long Run

The applied contribution of our paper an be better appreciated by relating to the literature on news and noise shocks. This literature has sought to formalize and evaluate the notion that optimism about the future prospects of the economy can cause a boom at present. This idea, however, does not work well within established versions of either the RBC or the NK framework.

In the RBC model, positive news or noise shocks raise estimates of permanent income. Households respond by raising their demand for goods and leisure. Employment falls. Furthermore, higher consumption spending pushes interest rates up, crowding out investment. All in all, news and noise shocks induce *negative* co-movement between consumption and the other macroeconomic variables. The same problem is present in the NK framework.

To overcome the negative co-movement problem, the literature has utilized a variety of new features. 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 accommodate the shifts in consumer expectations; that is, monetary policy induces pro-cyclical deviations from the underlying flexible-price allocations. The mechanism seems plausible, yet it is not sufficiently strong to offset the negative co-movement of the underlying flexible-price quantities once investment has been added to that model.[37]

The type of optimism and pessimism captured by *all* all these papers is fundamentally different from that in our model: ours is about the short-run economic outlook, theirs is primarily about productivity and income in the medium to long-run.

We illustrate this point in Figure 4. This figure revisits the exercise conducted in Figure 2, but replacing our confidence shock with the appropriate shocks featured in the aforementioned papers. In particular, we adopt the specification found in Barsky and Sims (2012). Accordingly, aggregate

[37] Blanchard, L'Huillier, and Lorenzoni (2013) try to fix this problem by adding investment adjustment costs.

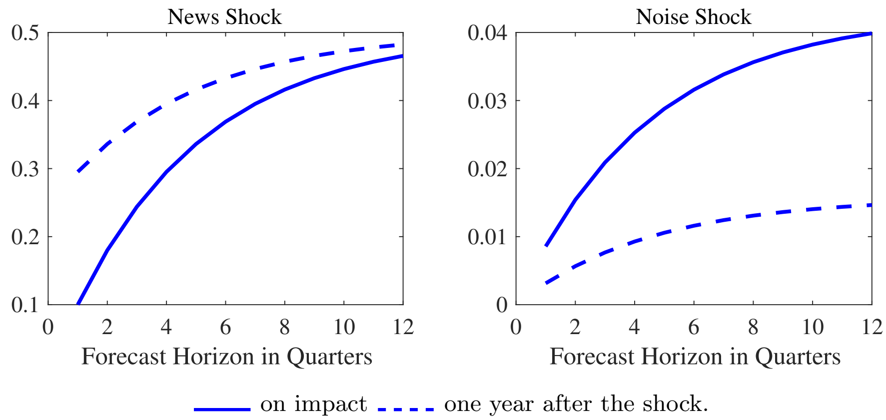

Figure 4: Forecasts of productivity at different horizons, following a news and a noise shock.

productivity is now given by

$$A_t = A_{t-1} + \gamma_{t-1} + \varepsilon_{a,t}$$
$$\gamma_t = \rho_\gamma \gamma_{t-1} + \varepsilon_{\gamma,t}$$

where $\rho \in (0,1)$ and where $\varepsilon_{a,t} \sim \mathcal{N}(0,\sigma_a^2)$ 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, $\varepsilon_{\gamma,t}$ captures news about future productivity that end up materializing, whereas $\varepsilon_{\eta,t}$ captures news that fail to materialize. The former are called "news " and the latter "noise" shocks.

Figure 4 shows how the term structure of expectations of future productivity at different horizons—namely, the locus of $\mathbb{E}_\tau[A_{\tau+k}]$ against k—looks right after (solid line) and also one year after the realization of the shock (dashed line). The expectations of output at different horizons look similar. The two shocks have the same qualitative effect on expectations: they cause agents to become optimistic about aggregate productivity and output in the medium to long run, without a commensurate movement in their expectations of the short run. This is simply because agents can not distinguish between the two shocks when they occur. The nature of optimism is therefore the same for both shocks—and is very different from the type of optimism seen in Figure 2. Finally, the fact that the dashed line lies below the solid line in the right panel of Figure 4 indicates that the effect of noise shock resembles that of our confidence shock in that they are both short lived; yet, in sharp contrast to the confidence shock, the noise shock has a disproportionate impact on long-term relative to short-term expectations regardless of whether we inspect these expectations right after the shock or later.

To recap, while our confidence shock resembles the noise shocks found in, inter alia, Lorenzoni (2009), Barsky and Sims (2012), and Blanchard, L'Huillier, and Lorenzoni (2013) in that both types of shocks are transitory, our mechanism captures a very different type of fluctuations. In those papers, recessions are periods in which the agents expect the economy to do badly for a long time; in our paper, they are periods in which the agents expect the economy to recover after a few quarters. It is precisely this difference, and not the specifics of how we have engineered the variation in expectations, that drives the superior quantitative performance of our mechanism.

The recent work by Ilut and Saijo (2016) further corroborates this point. This paper builds on different micro-foundations than ours, namely on a certain combination of learning and ambiguity aversion as opposed to higher-order beliefs. Yet, it effectively accommodates cyclical variation in the same kind of expectations as our paper, namely in expectations about the short run, and thereby obtains similar quantitative properties as our paper.[38]

6.4 Output Gaps and Aggregate Demand

In the NK model, the notion that recessions are driven by weak aggregate demand is captured by letting nominal rigidity, together with a particular type of monetary policy, generate deviations from the model's underlying flexible-price allocations. These deviations manifest as variation in measured markups and "output gaps", and are accompanied by commensurate movements in inflation.

There are no nominal rigidities in our setting. Nevertheless, because firms make their input choices prior to observing the demand for their products, a drop in confidence manifests itself as an increase in the realized markup; this is essentially the same property as the procylical labor wedge documented in Subsection 5.4. Furthermore, the resulting recession will register as a negative output gap insofar as the latter is measured relative to the underlying RBC benchmark, as illustrated by the example in Subsection 6.2. Consequently, the notion of demand-driven fluctuations formalized in our paper has a similar flavor and similar observable implications as the one formalized in the NK model.

There is, however, an important testable difference: in our setting, fluctuations in the aforementioned kind of output gap may arise without any movements in inflation. This is because our approach does not contain the restriction between the output gap and inflation encapsulated by the New-Keynesian Philips Curve. As a result, our approach avoids the "inflation puzzles" of the NK framework and may help explain, inter alia, why the severe contraction in output and employment during the recent recession was not accompanied by severe deflation.[39]

[38]One can generate similar co-movement patterns also by replacing our confidence shock with a shock that triggers errors in first-order beliefs of TFP in the short run, holding constant first-order beliefs of TFP is the medium to long run. We find this alternative theory less appealing than the one we have developed here, because it ties recessions with the perception of high marginal costs, whereas ours ties them with the perception of low demand. Nonetheless, this alternative helps highlight that the important element is the expectations about the short run, rather that the specific modeling devices through which these expectations arise.

[39]For a recent discussion of the "inflation puzzles" faced by NK models, see Beaudry and Portier (2013). The

We corroborate these points in Appendix C by repeating the horserace of Table 2 within the context of the NK model. The confidence shock continues to outperform the alternatives. Moreover, its real effects mirror those of monetary shocks, yet they come with only modest movements in inflation. A similar picture emerges from the richer, estimated models in Section 7.

One may nevertheless object to the interpretation of our mechanism as a formalization of the notion that the business-cycle is "demand-driven" on the following grounds. In our model, employment and output are fixed in the morning of each period, whereas consumption and investment adjust in the afternoon. In this regard, it appears that supply is determined first and that prices adjust to make sure that demand meets supply—which is the exact opposite of how the aggregate-demand mechanism is formalized within the NK framework. But instead of engaging in an endless debate of the merits and the drawbacks of different formalizations, let us only clarify one point: the ability of our model to generate the relevant co-movements does *not* depend on whether supply or demand is determined first.

We establish this by letting consumption and investment be fixed in the morning of each period and employment and output to be determined in the afternoon. By assuming the reverse timing protocol, this variant seems more in the spirit of the Keynesian view that "demand drives supply". Yet, as is evident in Figure 5, the observable implications differ very little across the two protocols.

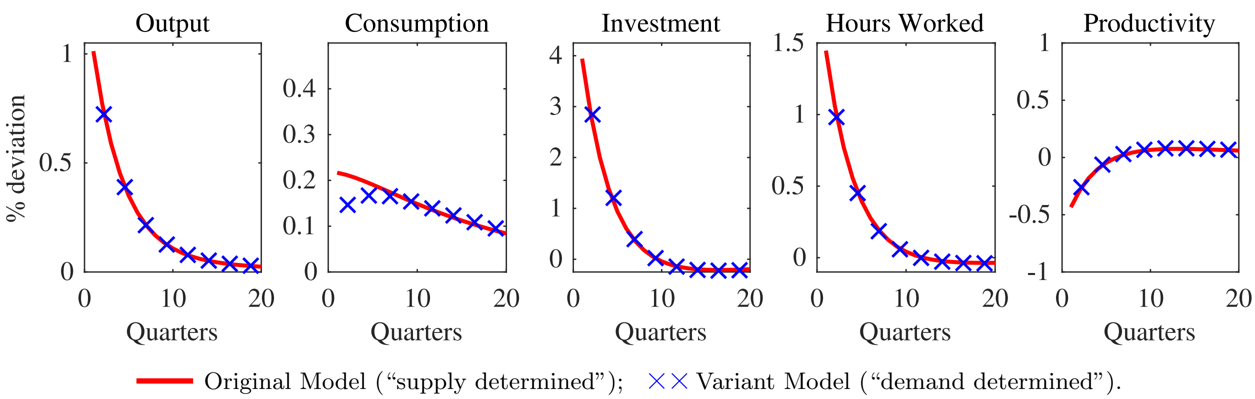

Figure 5: Impulse Responses to a Positive Confidence Shock, under Different Timing Protocols

The solid red lines in Figure 5 report the impulse response functions of the predetermined supply variant. For comparison, the blue crosses give the impulse response functions of the demand predetermined variant. 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 key business-cycle moments of the two models (see Table 8 in Appendix B), as well as to their predictions about the labor wedge.

Let us explain why. In the original model, a positive confidence shock causes employment and output to increase in the morning. This means that total spending *has* to increase in the afternoon.

extant literature has sought to minimize these puzzles by assuming, in effect, that the slope of the NK Philips Curve is extremely flat.

The composition of spending, 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 new model, 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 optimistic about the short run; if, instead, the shock caused the agents to become optimistic about the medium or long run, consumption would increase but investment would fall, and overall spending and employment would fall as well. We conclude that the timing protocol is not important.

7 Extension and Estimation

The analysis in Section 5 showed that the inclusion of our confidence mechanism in the RBC model produces a parsimonious account of multiple salient features of the data. In this section, we extend the analysis to a pair of richer, "medium-scale", DSGE models, which we estimate on US data. This exercise serves two goals: (*i*) to illustrate the robustness of our mechanism to the inclusion of multiple competing mechanisms; and (*ii*) to strengthen the claim that our mechanism can be viewed as a substitute—or, a complement—to the NK formalization of demand-driven fluctuations.

In the rest of this section, we first describe briefly the features of the two models and the rational behind the inclusion of these features. We next describe the estimation procedure. We finally review and interpret the main findings from the estimation of the two models. Many of the details as well as several additional results are delegated to Appendix D.

7.1 Two medium-scale models

The two models in this section share the same backbone with the baseline model, but include a number of additional ingredients that correspond to standard DSGE practices and help illustrate the robust performance of our mechanism.

To accommodate price-setting behavior, we now 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. 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; we refer to this model as the flexible-price model. In the other model, firms can instead adjust prices only infrequently, in the familiar Calvo fashion; we refer to this model as the sticky-price model and we close it by adding a conventional Taylor rule for monetary policy.

To accommodate the possibility that the business cycle is explained by multiple forces, and to let these forces compete with our mechanism, we include the following 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. Although this menu of shocks is not exhaustive, it allows us to embed a variety of mechanisms that seem a priori plausible and have also been found to be quantitatively significant in prior structural estimations.[40]

Finally, we allow for adjustment costs in investment (IAC) and habit persistence in consumption (HP). Although these modeling features lack compelling micro-foundations, they have become standard in the DSGE literature following the appearance of Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2007). This is 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. Our preceding analysis has already established that these modeling features are *not* needed for our mechanism to deliver realistic fluctuations. Here, we incorporate them in order to keep our exercise as close as possible to standard DSGE practice, and to give a better chance to the aforementioned competing shocks to outperform the confidence shock.

7.2 Estimation

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

Rational. 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 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

[40]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 and Buera and Moll (2012) for a thorough analysis of how different types of financial frictions map to different wedges. Fourth, the introduction of multiple transitory shocks, whatever their interpretation, increases the chance that these shocks, rather than our confidence shock, will pick up the transitory fluctuations in the data.

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 6 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.[41] 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.

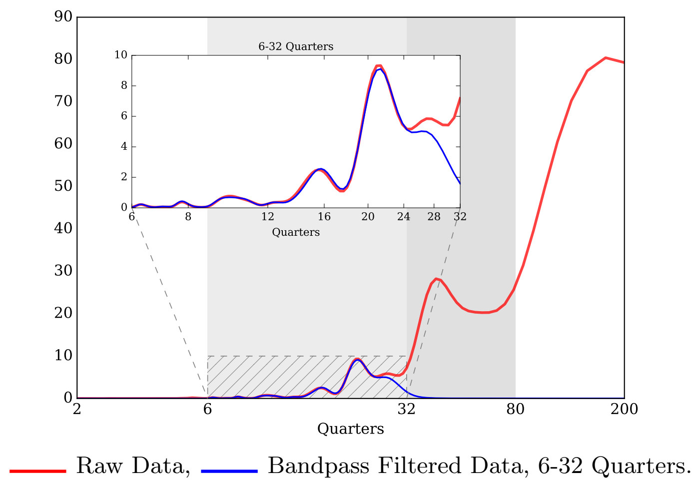

Figure 6: Spectral Density of Hours, 1960Q1-2007Q4.

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 the models 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 "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.[42]

[41]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.

[42]This is, of course, a premise that cannot be tested because we do not know what the "true" data-generating

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. The first four variables are in logs and linearly de-trended; the remaining two are in percentage points.[43] Our sticky-price model is estimated on the basis of all these six variables. By contrast, our 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. For pedagogical reasons, however, we find it useful to augment the estimated flexible-price model with a monetary policy rule that implements zero inflation.[44]

Priors. The priors used in the estimation are reported in Tables 10 and 11 in Appendix D. The priors for the preference, technology, and monetary parameters are in line with the literature. The priors for the shocks, on the other hand, are guided by two principles. First, we do not want a priori to favor the confidence shock: in either model, the confidence shock explains less than 3% of output volatility under the assumed priors. Second, we impose a tight prior only on the persistence of the transitory shocks, in order to help the estimation disentangle them from the permanent shocks.

Posteriors. Posterior distributions were obtained with the MCMC algorithm. The resulting estimates of the parameters are reported in the aforementioned tables. The estimated values of the preference, technology, and monetary parameters are close to previous estimates in the literature, indicating that the only major difference from the state of the art is the importance attributed to the novel structural shock, namely, the confidence shock.

7.3 Results

The picture that emerges from the estimation of two medium-scale models is consistent with the one that was painted in Section 5. We review some of the key findings below.

The confidence shock. Figure 7 reports the estimated IRFs to a positive confidence shock. 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 co-movement patterns found in the baseline model. This

process is. We nevertheless prefer this to the alternative. For instance, consider Smets and Wouters (2007). In an attempt to insulate the estimation of the business-cycle properties of their model from the observed low-frequency movements in inflation and hours, Smets and Wouters (2007) added ad-hoc slow-moving shocks to the price and the wage markups in their model. Nevertheless, as indicated by the findings of Sala (2015), the estimates of the business-cycle properties of that model are not immune to the exclusion of the low frequencies.

[43]Following Justiniano, Primiceri, and Tambalotti (2010), we do not include the price of investment in our data so as to accommodate both supply-side and demand-side interpretations of the investment shocsk. Also, we identify hours with hours in the non-farm business sector in order to make our results directly comparable to the extant DSGE literature. However, we have also repeated the estimation with total hours and have obtained similar results.

[44]This policy rule is equivalent to allowing the intercept in the Taylor rule to track the underlying natural rate. When this rule is appended to the sticky-price model, it replicates the underlying flexible-price allocations. Hence, one can re-interpret our flexible-price model as the sticky-price model augmented with this policy rule and re-estimated.

underscores the robustness of the key positive implications of our mechanism as we move between RBC and NK settings, or as we add the aforementioned modeling ingredients.

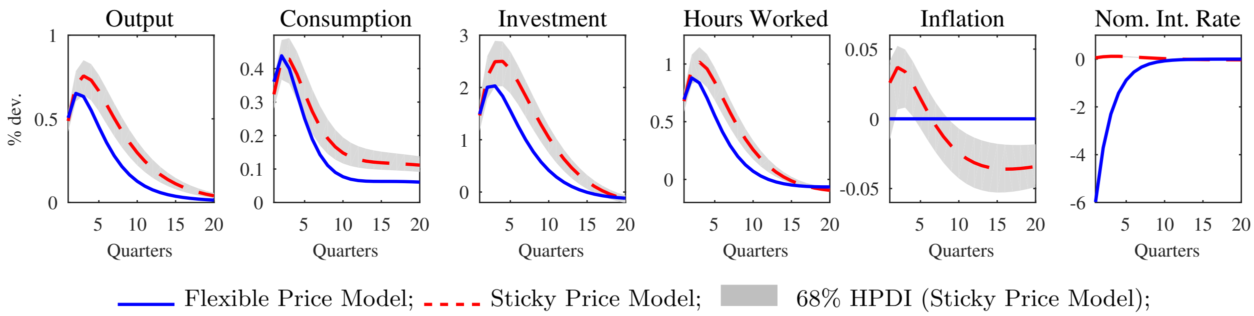

Figure 7: Theoretical IRFs to Confidence Shock

What differs, however, is the behavior of inflation and interest rates. In response to the confidence shock, as well as to other shocks, the flexible-price model predicts implausibly large movements in the real interest rate, due to the inclusion of the particular types of investment-adjustment costs and habit persistence. This compromises the model's performance vis-a-vis inflation and interest rates.[45] By contrast, because nominal rigidity permits the actual real interest to deviate from its natural level, the sticky-price model is able to accommodate simultaneously modest movements in the real and the nominal interest rate, as well as in inflation.

Table 4: Contribution of Confidence Shocks to Volatilities (6–32 Quarters)

	Y	C	I	h	π	R
Flexible Price	44.83	44.48	39.95	63.75	0.00	99.37
Sticky Price	51.69	46.38	48.07	63.91	8.44	32.50

Table 4 presents 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 45% of the business-cycle volatility in output in the flexible-price model, and for 52% in the sticky-price model.

Table 5 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 co-movement of all these variables.

These findings are not driven by the priors in the estimation: as noted earlier, the variance contribution of the confidence shock to output under the assumed priors is less than 3% in both

[45] Since we have augmented our flexible-price model with a monetary policy that stabilizes inflation, the large volatility of the real rate manifests fully in the nominal rate. But even if we had assumed a different monetary policy, the model would not manage to produce modest movements in both inflation and the nominal interest rate.

Table 5: Contribution of Confidence Shocks to Comovements (6–32 Quarters)

	Covariance					
	(Y, h)	(Y, I)	(Y, C)	(h, I)	(h, C)	(I, C)
Flexible Price	63.89	48.76	57.22	59.17	99.94	89.21
Sticky Price	65.98	55.26	63.86	61.85	100.12	95.18

models. What explains these findings is that the data favor a mechanism that triggers strong procylical movements in hours, investment, and consumption without commensurate movements in labor productivity, TFP, inflation, and interest rates. Our mechanism is well positioned to generate this kind of co-movement within workhorse macroeconomic models.

The other shocks. To economize on space, the estimated role of the other shocks is reported in Appendix D. One finding is nevertheless worth mentioning. The co-movement implications and the estimated importance of investment-specific, discount-rate, and news shocks change substantially depending on whether we turn on or shut down the particular forms of adjustment costs in investment and habit persistence in consumption that are popular in the DSGE literature.[46] By contrast, neither the co-movement implications of the confidence shock nor its estimated contribution are unduly sensitive to the inclusion or exclusion of these modeling features.

Business-cycle moments. Both models do a good job in matching key business-cycle moments; see Table 13 in Appendix D. This is not surprising given that Section 5 documented a good fit under a more rigid theoretical structure. It nevertheless underscores the robustness of the empirical performance of our mechanism.

Shocks vs empirical proxies. To further corroborate our theory, we now consider the following question: do the technology and confidence shocks in our theory have any resemblance to empirical measures of, respectively, TFP and "market psychology" in the real world?

We address this question in Figure 8 by comparing the estimated series of the technology and the confidence shocks in our models with, respectively, the series of Fernald (2014) utilization-adjusted TFP measure and the University of Michigan Consumer Sentiment Index. Even though we did not use any information on these two empirical measures in the estimation of the models, the theoretical shocks turn to be highly correlated with their empirical counterparts. The same property holds if we replace the Michigan Sentiment Index with the Conference Board's Indices of Consumer or Producer Confidence, because these indices are highly correlated at business-cycle frequencies. Notwithstanding the inherent difficulty of interpreting such indices and of mapping them to the theory, we view this finding as providing additional validation to our mechanism and its interpretation as variation in "confidence".

[46] This is true, not only in our models, but also in the relevant literature (e.g., Justiniano, Primiceri, and Tambalotti, 2010, Blanchard, L'Huillier, and Lorenzoni, 2013).

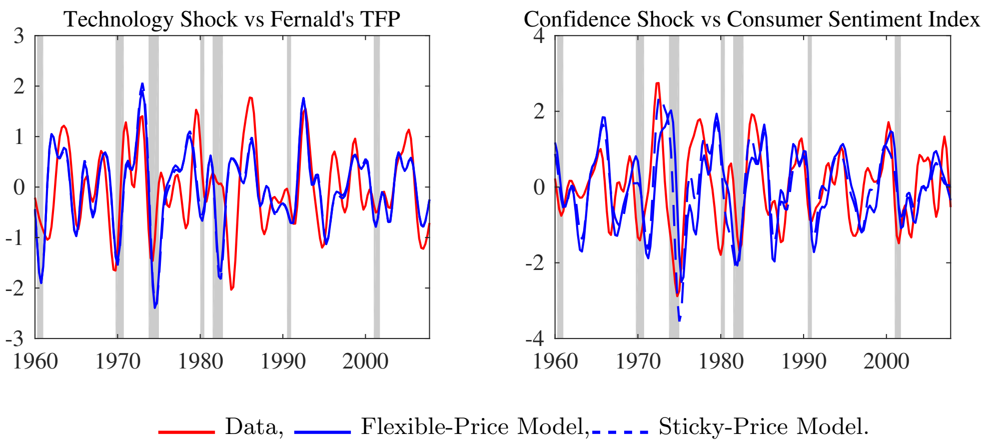

Figure 8: Shocks vs Empirical Proxies

On aggregate demand. As noted before, our approach can be thought as an alternative to the NK formalization of the notion that business cycles are driven by aggregate demand. Of course, the two are not mutually inconsistent: both nominal rigidity and frictions in coordination/beliefs may be important for understanding business cycles. Notwithstanding this point, we now evaluate the relative potency of the two mechanisms, when each one operates in isolation, with regard to the following question. Which of the two mechanisms improves by a larger margin the ability of the RBC framework to capture the joint dynamics of employment, output, investment, and consumption over the business cycle?[47]

We answer this question by computing the posterior odds that the data are generated by each one of the following three possible models: the variant of our flexible-price model that shuts down the confidence shock; the variant of our sticky-price model that also shuts down the confidence shock; and finally the flexible-price model that contains the confidence shock. The first model serves as a benchmark, capturing the RBC core of the other two models. The second model adds sticky prices to the first model, isolating the NK mechanism. The third model replaces the nominal rigidity with a confidence shock, isolating our own mechanism. Since we are interested in comparing these models on the basis of their ability to fit the co-movement of the real quantities, we estimate them using real quantities only. We finally compute the posterior odds that the data are generated by the sticky-price model rather than by either of the flexible-price models, starting with an uniform prior over the three models. These odds provide a metric of how well a given model captures the data relative to another model.[48]

Consider first the pair-wise comparison between the sticky-price model and the flexible-price model that abstracts from the confidence shock. In this case, the sticky-price model wins: the posterior odds that the data are generated by that model are greater than 99%. Consider next how this comparison is affected when the flexible-price model is augmented with the confidence shock.

[47]This question abstracts from monetary phenomena, which our flexible-price models are not equipped to address.
[48]See Table 12 in Appendix D for a tabulation of the posterior odds of the three models.

The odds are now completely reversed: the probability that the data are generated by the sticky-price model are 0%. By this metric, nominal rigidity is essential for the ability of the theory to match the real data when our mechanism is absent, but not if it is present: our mechanism appears to to be *more* potent than the NK mechanism when their relative performance is evaluated in terms of likelihood, as described above.

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. More importantly, 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" in large part to a coordination failure and a lack of confidence as opposed to merely nominal rigidity.

8 Conclusion

By relying on a particular solution concept together with complete information, 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 proposed a certain relaxation of this structure and explored its quantitative implications. In particular, we developed a method for augmenting macroeconomic models with a tractable form of higher-order belief dynamics that can proxy for incomplete information and imperfect coordination and that ultimately helps capture waves of optimism and pessimism about the short-term outlook of the economy. We interpreted these waves as variation in "confidence" and documented their quantitative importance within RBC and NK 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.
- It appears to outperform alternative structural mechanisms that are popular in the literature.
- It offers a potent formalization of the notion of demand-driven fluctuations that can serve as either an alternative or a complement to the NK formalization of this notion and that can also bypass the inflation/deflation puzzles of the NK literature.
- 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. This limitation, however, is not specific to what we do in this paper: any formal model must ultimately attribute the business cycle to some exogenous trigger, whether this is a technology shock, a discount-rate shock, a financial shock, or even a sunspot. Therefore, although we remain agnostic about the deeper foundations of the belief waves we have accommodated in this paper, we hope to have provided a useful gauge of their potential quantitative importance and a stepping stone for future research.

A few papers are already pushing the frontier. Huo and Takayama (2015b) obtain quantitative findings that are broadly consistent with ours while maintaining the common-prior assumption. Myohl and Stucki (2016) study the interaction of our mechanism with financial frictions. Ilut and Saijo (2016) consider a model that endogenizes the level of confidence and lets it covary with other structural shocks; this has the interesting implication that a drop in confidence may be triggered by an adverse financial shock, while a boost in confidence may be accomplished by a fiscal stimulus. Angeletos, Collard, and Dellas (2016) provide VAR-based evidence that the business cycle in the US data appears to be driven 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.

Open directions for future research include the investigation of policy implications; the application of our approach to the recent recession; the use of our methodology towards a quantitative evaluation of the potential inertia in the response of expectations to shocks in fundamentals such as technology and policy; and the empirical investigation of the distinct types of expectations we have highlighted in this paper.

APPENDICES

A. Data

In this appendix we describe the data we use in this paper to (i) obtain the various business-cycle moments and (ii) to estimate the models considered in Section 7.

The data is from the Saint–Louis Federal Reserve Economic Database. 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 6: 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

Table 6 summarizes information about the data. 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, H, are measured by hours of all persons in the non-farm business sector. Labor productivity, Y/H, 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, the Total Factor Productivity corresponds to the TFP adjusted for utilization as reported in Fernald (2014).

B. Baseline Model: Anatomy and Wedges

In the first part of this Appendix, we elaborate on the distinct quantitative roles that the technology and the confidence shock play in our RBC prototype. In the second part, we report the business-cycle moments of the variant that changes the timing protocol. In the last part, we spell out the computation of the model-based wedges and the estimation of their empirical counterparts.

An Anatomy. Table 7 reports the business-cycle moments predicted by the versions of our model that shut down either of the two shocks (see the last two columns) and compares them to those in the data (first column) and in our full model (second column).

By isolating the role of the technology shock, the third column in this table revisits, in effect, the baseline RBC model. The most noticeable, and well known, failures of this model are its inability to generate a sufficiently high volatility in hours; its prediction of a counterfactually strong correlation between hours and labor productivity. An additional failure is that the model generates a counterfactually low volatility in investment.

Consider now the fourth column, which isolates the confidence shock.[49] The key failures are now the counterfactually high volatility in hours, the perfectly negative correlation between labor productivity and either hours or output, and the counterfactually low volatility in consumption. The first two properties follow directly from the fact that technology is fixed and exhibits diminishing returns in labor, while the last property is driven by the transitory nature of the confidence shock.

To sum up, neither the standard RBC mechanism nor our mechanism deliver a good fit when working in isolation. But once they work together, the fit is great.

Timing Protocol. Table 8 reports the business-cycle moments predicted by the variant of our baseline model that changes the time protocol, letting consumption and investment to be set in the morning and employment and output to adjust in the afternoon. Comparing this table to Table 2 in the main text reveals that the timing protocol has only a negligible effect on the empirical fit of our theory.

[49] This version is similar to the model studied in Section 7 of Angeletos and La'O (2013).

Table 7: The Role of Each Shock

	Data	Our Model	A Only	ξ Only
Standard deviations				
stddev(y)	1.41	1.52	1.10	1.03
stddev(h)	1.55	1.52	0.32	1.48
stddev(c)	0.76	0.69	0.64	0.26
stddev(i)	5.12	5.00	2.88	4.06
stddev(y/h)	0.76	0.91	0.79	0.46
Correlations				
corr(c, y)	0.85	0.89	0.99	0.85
corr(i, y)	0.94	0.97	0.99	0.99
corr(h, y)	0.87	0.82	0.98	1.00
corr(c, h)	0.83	0.49	0.93	0.80
corr(i, h)	0.82	0.92	1.00	1.00
corr(c, i)	0.73	0.76	0.96	0.78
Correlations with productivity				
corr($y, y/h$)	0.07	0.29	1.00	-0.96
corr($h, y/h$)	-0.43	-0.30	0.96	-0.98
corr(y, sr)	0.82	0.73	1.00	-0.97
corr(h, sr)	0.47	0.22	0.98	-0.94

Note: The first column reports moments in the US data, bandpass-filtered, over the 1960-2007 period. The second column reports the moments in our model. The rest of the table reports the moments in the two variants that shut down either the TFP or the confidence shock..

Table 8: Different Timing Protocols

	Data	determined in the morning: supply (h, y)	determined in the morning: demand (c, i)
Standard deviations			
stddev(y)	1.41	1.52	1.52
stddev(h)	1.55	1.52	1.49
stddev(c)	0.76	0.69	0.68
stddev(i)	5.12	5.00	5.23
stddev(y/h)	0.76	0.91	0.92
Correlations			
corr(c, y)	0.85	0.89	0.83
corr(i, y)	0.94	0.97	0.96
corr(h, y)	0.87	0.82	0.81
corr(c, h)	0.83	0.49	0.37
corr(i, h)	0.82	0.92	0.93
corr(c, i)	0.73	0.76	0.65
Correlations with Productivity			
corr$(y, y/h)$	0.07	0.29	0.32
corr$(h, y/h)$	-0.43	-0.30	-0.28
corr(y, sr)	0.82	0.73	0.74
corr(h, sr)	0.47	0.22	0.22

Note: The first column reports moments in the US data, bandpass-filtered, over the 1960-2007 period. The second and third column report the moments in our model under the two alternative timing protocols.

Wedges. We now show how the various wedges in the model are defined and computed. All the wedges are denoted by τ_t, with an appropriate superscript, and are measured in terms of log-deviations, like all other variables. We adopt the convention that $\tau > 0$ means an implicit tax and, conversely, $\tau < 0$ means an implicit subsidy.

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

$$c_{it} = \Gamma_K^c K_t + \Gamma_z^c z_{it} + \Gamma_{\overline{z}}^c \overline{z}_t + \Gamma_a^c A_t + \Gamma_\xi^c \xi_t$$

Therefore, the stage-1 expectation of own consumption is given by

$$\mathbb{E}_{it}[c_{it}] = \Gamma_K^c K_t + (\Gamma_z^c + \Gamma_{\overline{z}}^c + \Gamma_a^c) A_t + (\Gamma_\xi^c + \Gamma_{\overline{z}}^c)\xi_t$$

Realized consumption, on the other hand, is given by

$$C_t = \Gamma_K^c K_t + (\Gamma_z^c + \Gamma_{\overline{z}}^c + \Gamma_a^c) A_t + \Gamma_\xi^c \xi_t$$

Combining, we infer that the household-side component of the labor wedge is given by

$$\tau_t^{nh} \equiv \mathbb{E}_{it}[c_{i,t}] - C_t = \Gamma_{\overline{z}}^c \xi_t$$

For our calibration, we have $\tau_t^{nh} = 0.0152\xi_t$.

The labor wedge on the household 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 the firm-side component of the labor wedge is given by

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

For our calibration, 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:

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

In the equilibibrium of our model,

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

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 calibration, 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 $\tau_t^k = \tau_t^{kh} + \tau_t^{kf} = 0.1366\xi_t$. The former is negatively correlated with the confidence shock, the latter is positively correlated.

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^n, a capital wedge, τ_t^k, and a resource wedge τ_t^g. Accordingly, the system to be estimated is the following:

$$\nu N_t + C_t = Y_t - N_t - \tau_t^n \tag{26}$$
$$\mathbb{E}_t[C_{t+1}] - C_t = (1-\beta(1-\delta))(\mathbb{E}_t[Y_{t+1} - K_{t+1}] - \tau_t^k) \tag{27}$$
$$Y_t + (1-\delta)K_t = C_t + K_{t+1} + \tau_t^g \tag{28}$$
$$Y_t = \tau_t^e + \alpha K_t + (1-\alpha)N_t \tag{29}$$

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 3.

C. The Horse-race Repeated Within the NK Framework

This appendix revisits the "horserace " introduced in Section 5. Recall that the competing models were variants of our baseline model that replace the confidence shocks with one of the following: a transitory technology shock; a transitory investment-specific shock; a transitory discount-factor shocks; or a news shocks about future TFP. Also recall that the analysis was limited to the RBC context. Here, we complement that analysis by repeating the same horserace in the context of the NK framework; we also monetary shock to the set of shocks contained in the horserace .

The NK extension of our analysis is standard and is described in the beginning of Section 7. For the calibrated versions considered in this appendix, the markup rate is set to 15%; the Calvo probability of resetting prices is set so that the average length of an unchanged price is 4 quarters; and monetary policy is assumed to follow a Taylor rule with coefficient on inflation of 1.5 and on output of 0.05, and a degree of interest rate smoothing of 0.8. For illustration purposes, we also consider the alternative case in which the monetary authority completely stabilizes the nominal interest rate. The preference and technology parameters remain the same as before. Likewise the standard deviations of the shocks are obtained by minimizing the weighted distance between the volatility of output, consumption, investment and hours between the data and the model, assuming the same persistence for each alternative shock as that of the confidence shock—with the exception of the monetary shock which has persistence 0.15, a value that lies within the range of values usually found in the literature (see for example Smets and Wouters, 2007).

Table 9 reports the same kind of information as Table 2, that is, the business-cycle moments of the different models in the horserace, but now for the NK versions of these models. Five key lessons emerge. First, the superior empirical performance of our baseline model survives when we introduce sticky prices. Second, with the exception of the monetary shock, none of the aforementioned competing structural shocks is able to generate positive co-movement patterns in the real data whether one considers the RBC or the NK version of the models. Third, the monetary shock can match these patterns quite well, but only at the expense of requiring implausibly large monetary shocks (an order of magnitude larger than monetary shocks identified in SVARs) and of overstating the procyclicality of inflation and interest rates. Fourth, the similarity between the real effects of the confidence shock and those of the monetary shock provide further justification for reinterpreting the confidence shock as a formalization of the popular notion of "demand-driven fluctuations". Last but not least, the NK version our model predicts a negligible correlation between inflation and output, underscoring the ability of our mechanism to capture the aforementioned notion while also bypassing the inflation puzzles of the NK framework.

These lessons are not unduly sensitive to the parametrizations chosen. They survive when we move to estimated versions of richer RBC and NK models that allow multiple structural shocks to coexist and that also introduce two propagation mechanism that have played a crucial role in the DSGE literature, namely investment-adjustment costs and habit. See the discussion in Section 7 and the estimated IRFs in Figures 9 and 10 in Appendix D.

Table 9: horserace within the NK Framework

	Data	Our Model		Alternative Two-Shock Models				
		RBC	NK	TFP	Invt	Disc	News	Mon.
Standard deviations								
stddev(y)	1.41	1.52	1.52	1.66	1.35	1.25	1.54	1.52
stddev(h)	1.55	1.52	1.46	0.94	1.33	1.05	1.07	1.46
stddev(c)	0.76	0.69	0.66	0.61	0.75	0.94	0.64	0.66
stddev(i)	5.12	5.00	5.44	6.23	6.38	6.46	6.57	5.42
stddev(y/h)	0.76	0.91	0.88	0.83	0.81	0.78	0.74	0.90
Correlations								
corr(c, y)	0.85	0.89	0.77	0.75	0.20	0.10	0.41	0.78
corr(i, y)	0.94	0.97	0.96	0.97	0.91	0.84	0.95	0.96
corr(h, y)	0.87	0.82	0.82	0.95	0.82	0.78	0.90	0.82
corr(c, h)	0.83	0.49	0.30	0.50	-0.40	-0.53	-0.01	0.30
corr(i, h)	0.82	0.92	0.94	0.99	0.98	0.99	0.99	0.94
corr(c, i)	0.73	0.76	0.57	0.57	-0.23	-0.44	0.10	0.58
Correlations with productivity								
corr($y, y/h$)	0.07	0.29	0.35	0.93	0.32	0.54	0.77	0.35
corr($h, y/h$)	-0.43	-0.30	-0.24	0.76	-0.27	-0.10	0.43	-0.24
corr(y, sr)	0.82	0.73	0.76	0.98	0.75	0.83	0.92	0.76
corr(h, sr)	0.47	0.22	0.27	0.89	0.25	0.32	0.67	0.27
Nominal variables								
stddev(π)	0.23	–	0.09	0.06	0.18	0.06	0.09	0.12
stddev(R)	0.35	–	0.10	0.10	0.31	0.09	0.13	0.10
corr(y, π)	0.21	–	0.07	0.77	0.84	0.76	0.47	0.94
corr(y, R)	0.33	–	0.69	0.83	0.82	0.80	0.57	0.66

Note: The first column reports moments in the US data, bandpass-filtered, over the 1960-2007 period. The second and third columns report the moments in, respectively, the RBC and the NK version of our model. The rest of the table reports the moments in the NK versions of the five competing models discussed in the text.

D. Estimated Models

In this appendix we discuss the estimation of the multi-shock RBC and NK models considered in Section 7. In particular, we first fill in the formal details of the two models; we next explain the estimation method, the priors assumed for the parameters, and the posteriors obtained from the estimation; we finally review a number of findings that were omitted, or only briefly discussed, in the main text.

The details of the two models. As mentioned in the main text, the two models we study in Section 7 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. 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 given by the following CES aggregate:

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

where $\eta > 0$ is a parameter that pins down the monopoly power. The technology is the same as before, so that the output of firm j in island i is

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

but now TFP is given by the sum of two permanent components, one corresponding to a standard unanticipated innovation and another corresponding to a news shock, plus a temporary shock. More specifically,

$$\log A_t = a_t^{\tau} + a_t^{p},$$

where a_t^{τ} is the transitory component, modeled as an AR(1), and a_t^{p} is the sum of the aforementioned two permanent components, namely,

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

where ε_t^p is the unanticipated innovation and ζ_t^n captures all the TFP changes that agents anticipated in earlier periods. The latter is given by a diffusion-like process of the form

$$\zeta_t^n = \frac{1}{8} \sum_{j=1}^{8} \varepsilon_{t-j}^n$$

where ε^n_{t-j} is the component of the period-t innovation in TFP that becomes known in period $t-j$.[50] In line with our baseline model, the confidence shock is now modeled as a shock to higher-order beliefs of a^p_t.

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

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

where ζ^c_t 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.[51]

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

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

where Z_t denotes the relative price of investment, G_t denotes government spending (the cost of which is assumed to equally spread across the islands), and $\exp(Z_t)\Psi(u_{it})$ denotes the resource cost of utilization per unit of capital. The latter is scaled by $\exp(Z_t)$ in order to transformed the units of capital to units of the final good, and thereby also guaranteed a balanced-growth path. Z_t is modeled as a random walk: $Z_t = Z_{t-1} + \varepsilon^z_t$. 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. Government spending is given by $G_t = \bar{G}\exp(\tilde{G}_t)$, where $\bar{G}$ is a constant and

$$\tilde{G}_t = \zeta^g_t + \tfrac{1}{1-\alpha}a_t - \tfrac{\alpha}{1-\alpha}Z_t.$$

In the above, ζ^g_t denotes a transitory shock, modeled as an AR(1), and the other terms are present in order to guarantee a balanced-growth path. The utilization-cost function satisfies $u\Psi''(u)/\Psi'(u) = \frac{\psi}{1-\psi}$, with $\psi \in (0,1)$.

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:

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

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^i_t is a temporary shock, modeled as an AR(1) and shifting the demand for investment, as in Justiniano, Primiceri, and Tambalotti (2010).

[50] We have experimented with alternative forms of diffusion, as well as with specifications such as $\zeta^n_t = \varepsilon^n_{t-4}$, and we have found very similar results.

[51] 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.

The above description completes the specification of the flexible-price model of Section 7. 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 (40) 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:

$$\tilde{\lambda}_{it} = \zeta_t^c - \tfrac{1}{1-b}\tilde{c}_{it} + \tfrac{b}{1-b}\tilde{C}_{t-1} \tag{30}$$

$$\zeta_t^c + \nu\tilde{n}_{it} = \mathbb{E}_{it}\left[\tilde{\lambda}_{it} + \tilde{s}_{it} + \tilde{Y}_t - \tilde{n}_{it}\right] \tag{31}$$

$$Z_t + \tfrac{1}{1-\psi}\tilde{u}_{it} = \tilde{s}_{it} + \tilde{Y}_t - \tilde{k}_{it} \tag{32}$$

$$\tilde{y}_{it} = a_t + \alpha(\tilde{u}_{it} + \tilde{k}_{it}) + (1-\alpha)\tilde{n}_{it} \tag{33}$$

$$\tilde{Y}_t = s_c\tilde{c}_{it} + (1-s_c-s_g)(Z_t + \tilde{\imath}_{it}) + s_g\tilde{G}_t + \alpha\tilde{u}_{it} \tag{34}$$

$$\tilde{k}_{it+1} = \delta\tilde{\imath}_{it} + (1-\delta)\tilde{k}_{it} \tag{35}$$

$$\tilde{q}_{it} = (1+\beta)\varphi\tilde{\imath}_{it} - \varphi\tilde{\imath}_{t-1} - \beta\varphi\mathbb{E}'_{it}\tilde{\imath}_{it+1} + Z_t - \zeta_t^i \tag{36}$$

$$\tilde{R}_t = \tilde{\lambda}_{it} - \mathbb{E}'_{it}\left[\tilde{\lambda}_{it+1} - \tilde{\pi}_{it+1}\right] \tag{37}$$

$$\tilde{\lambda}_{it} + \tilde{q}_{it} = \mathbb{E}'_{it}\left[\tilde{\lambda}_{it+1} + (1-\beta(1-\delta))\left(\tilde{s}_{it+1} + \tilde{Y}_{t+1} - \tilde{u}_{it+1} - \tilde{k}_{it+1}\right) + \beta(1-\delta)\tilde{q}_{it+1}\right] \tag{38}$$

$$\tilde{X}_t = s_c\tilde{C}_t + (1-s_c-s_g)(Z_t + \tilde{I}_t) + s_g\tilde{G}_t \tag{39}$$

$$\tilde{R}_t = \kappa_R\tilde{R}_{t-1} + (1-\kappa_R)\left(\kappa_\pi\tilde{\pi}_t + \kappa_y(\tilde{X}_t - \tilde{X}_t^F)\right) + \zeta_t^m \tag{40}$$

$$\tilde{\pi}_{it} = (1-\chi)(1-\beta\chi)\mathbb{E}'_{it}\tilde{s}_{it} + (1-\chi)\mathbb{E}'_{it}\tilde{\Pi}_t + \beta\chi\mathbb{E}'_{it}\tilde{\pi}_{it+1} \tag{41}$$

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_t denotes the measured aggregate GDP, X_t^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 familiar. Condition (30) gives the marginal utility of consumption. Conditions (31) and (32) characterizes the equilibrium employment and utilization levels. Condition (33) and (34) give the local output and the local resource constraint. Conditions (35) and (36) give the local law of motion of capital and the equilibrium investment decision. Conditions (37) and (38) are the two Euler equations: the first corresponds to optimal bond holdings and gives the relation between consumption growth and the interest rates, while the second corresponds to optimal capital accumulation and gives the evolution of Tobin's Q. Condition (39) gives the measured aggregate GDP. Conditions (40) gives the Taylor rule for monetary policy. Finally, condition (41) 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 (37) in terms of the real interest rate; and we drop the nominal side of this system, namely conditions (40) and (41).

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 $t = 1 \dots T$) 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:

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

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 Hannan (1970) and Harvey (1991), the likelihood function is asymptotically given by

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

where $\omega_j = 2\pi j/T$, $j = 1 \dots T$ 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.

Parameters: priors. The first three columns in Tables 10 and 11 report the priors used in the estimation of the parameters of the two models. The logic behind our choice of priors for the shock processes was discussed in the main text; we now briefly discuss the rest of the priors.

Table 10: Estimated Parameters, Part I

	Priors			Posteriors	
	Shape	Mean	Std. Dev.	FP	SP
Real Aspects					
ν	Gamma	0.500	0.250	0.538	0.576
				[0.121, 1.006]	[0.184, 1.030]
α	Beta	0.300	0.050	0.271	0.281
				[0.242, 0.301]	[0.254, 0.309]
ψ	Beta	0.300	0.150	0.371	0.442
				[0.140, 0.630]	[0.238, 0.663]
φ	Gamma	2.000	1.000	2.036	1.954
				[0.847, 3.345]	[0.763, 3.240]
b	Beta	0.500	0.250	0.711	0.647
				[0.619, 0.797]	[0.537, 0.756]
Nominal Aspects					
χ	Beta	0.500	0.250	–	0.724
					[0.660, 0.784]
κ_R	Beta	0.750	0.100	–	0.372
					[0.242, 0.500]
κ_π	Normal	1.500	0.250	–	1.123
					[0.994, 1.265]
κ_y	Normal	0.125	0.050	–	0.134
					[0.065, 0.203]

Note: 95% HPDI into brackets.

Table 11: Estimated Parameters, Part II

	Priors			Posteriors	
	Shape	Mean	Std. Dev.	FP	SP
Shocks: persistence					
ρ_a	Beta	0.500	0.100	0.477	0.481
				[0.285, 0.667]	[0.289, 0.686]
ρ_i	Beta	0.500	0.100	0.494	0.481
				[0.304, 0.678]	[0.313, 0.648]
ρ_c	Beta	0.500	0.100	0.487	0.492
				[0.305, 0.671]	[0.295, 0.691]
ρ_g	Beta	0.500	0.100	0.623	0.598
				[0.442, 0.790]	[0.419, 0.769]
ρ_m	Beta	0.500	0.100	–	0.432
					[0.289, 0.576]
ρ_ξ	Beta	0.500	0.100	0.610	0.765
				[0.435, 0.777]	[0.727, 0.803]
Shocks: volatilities					
$\sigma_{a,p}$	Inv. Gamma	1.000	4.000	0.544	0.516
				[0.377, 0.707]	[0.362, 0.666]
$\sigma_{a,t}$	Inv. Gamma	1.000	4.000	0.395	0.401
				[0.228, 0.582]	[0.238, 0.585]
$\sigma_{a,n}$	Inv. Gamma	1.000	4.000	0.477	0.490
				[0.251, 0.759]	[0.313, 0.686]
$\sigma_{i,p}$	Inv. Gamma	1.000	4.000	1.402	1.582
				[0.303, 2.629]	[0.486, 2.619]
$\sigma_{i,t}$	Inv. Gamma	1.000	4.000	1.914	2.737
				[0.247, 4.698]	[0.917, 4.751]
σ_c	Inv. Gamma	1.000	4.000	0.883	0.479
				[0.269, 1.564]	[0.237, 0.767]
σ_g	Inv. Gamma	1.000	4.000	1.853	1.799
				[1.394, 2.341]	[1.341, 2.283]
σ_m	Inv. Gamma	1.000	4.000	–	0.261
					[0.206, 0.322]
σ_ξ	Inv. Gamma	1.000	4.000	2.886	0.453
				[1.147, 5.621]	[0.241, 0.729]

Note: 95% HPDI into brackets.

The inverse labor supply elasticity, ν, is Gamma distributed around 0.5 with standard deviation 0.25. The capital share, α, is Beta distributed around 0.3 with standard deviation 0.05. The utilization elasticity parameter, ψ, is Beta distributed around 0.3 with standard deviation 0.15. The habit persistence parameter, b, is Beta distributed around 0.5 with standard deviation 0.25. The parameter governing the size of investment adjustment costs, φ, is Gamma distributed around 2 with a standard deviation 1; the relatively high standard deviation reflects our own uncertainty about this modeling feature, but also allows the estimation to accommodate the higher point estimates required by the pertinent DSGE literature, in case that would improve the empirical performance of the model. The Calvo probability of not resetting prices, χ, is Beta distributed around 0.5 with standard deviation 0.25. The persistence parameter in the Taylor rule, κ_R, is Beta distributed around 0.75 with standard deviation 0.1; the reaction coefficient on inflation, κ_π, is Normally distributed around 1.5 with standard deviation 0.25; and the reaction coefficient on the output gap, κ_y, is also Normally distributed with mean 0.125 and standard deviation 0.05. Finally, the following three parameters are fixed: the discount factor β is 0.99; the depreciation rate δ is 0.025; and the CES parameter η is such that the monopoly markup is 15%, which is eliminated in steady state by a subsidy.

Parameters: posteriors. Posterior distributions were obtained with the MCMC algorithm. We generated 2 chains of 100,000 observations each. The posteriors for all the parameters of our two models are reported in the last two columns of Tables 10 and 11. The posteriors for the preference, technology, and monetary parameters are broadly consistent with other estimates in the literature. Below, we find it useful only to comment on the estimated size of σ_ξ, the standard deviation of the innovation in the confidence shock.

IRFs. Figures 9–10 report the IRFs of our estimated models with respect to all the structural shocks. In the main text, we mentioned that the inclusion of investment adjustment costs (IAC) and habit (HP) plays a crucial role in the existing DSGE literature, but has only a modest effect on the performance of our own mechanism. To illustrate this point, Figures 9–10 report the IRFs of the various shocks both in the versions of our models that include these propagation mechanisms and in those that shut them down.

Business-cycle moments. Table 13 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) and our own baseline model (last column). Inspection of this table leads to the following conclusions.

First, both of the estimated models do a good job on the real side of the economy. Perhaps this is not surprising given that Section 5 documented a good fit under a more rigid theoretical structure. But this only underscores the good empirical performance of our mechanism.

Second, in comparison to Smets and Wouters (2007), our sticky-price model does a good job in matching, not only the real, but also the nominal side of the data. 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, first, that the inclusion of our mechanism in NK models does not interfere with their ability to capture the nominal side of the data and, second, that our mechanism itself is robust to the introduction of realistic nominal rigidities.

Variance/Covariance Decompositions. Tables 14 and 15 report 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.) 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 DSGE 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 co-movement patterns.

Posterior odds. We conclude this appendix by reporting the posterior odds of the models used at the end of Section 7, for the purpose of comparing our mechanism to the NK mechanism. Table 12 reports the posterior odds of four alternative 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 a confidence shock or not. Since we have dropped 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 a fair chance to each model to match the real data.

Table 12: Posterior Odds of Model A vs Model B

Model A →	flex prices	sticky prices	
Model B ↓	without	with	without
flex prices, with confidence	0.00	0.81	0.00
flex prices, without confidence	–	1.00	1.00
sticky prices, with confidence	0.00	–	0.00

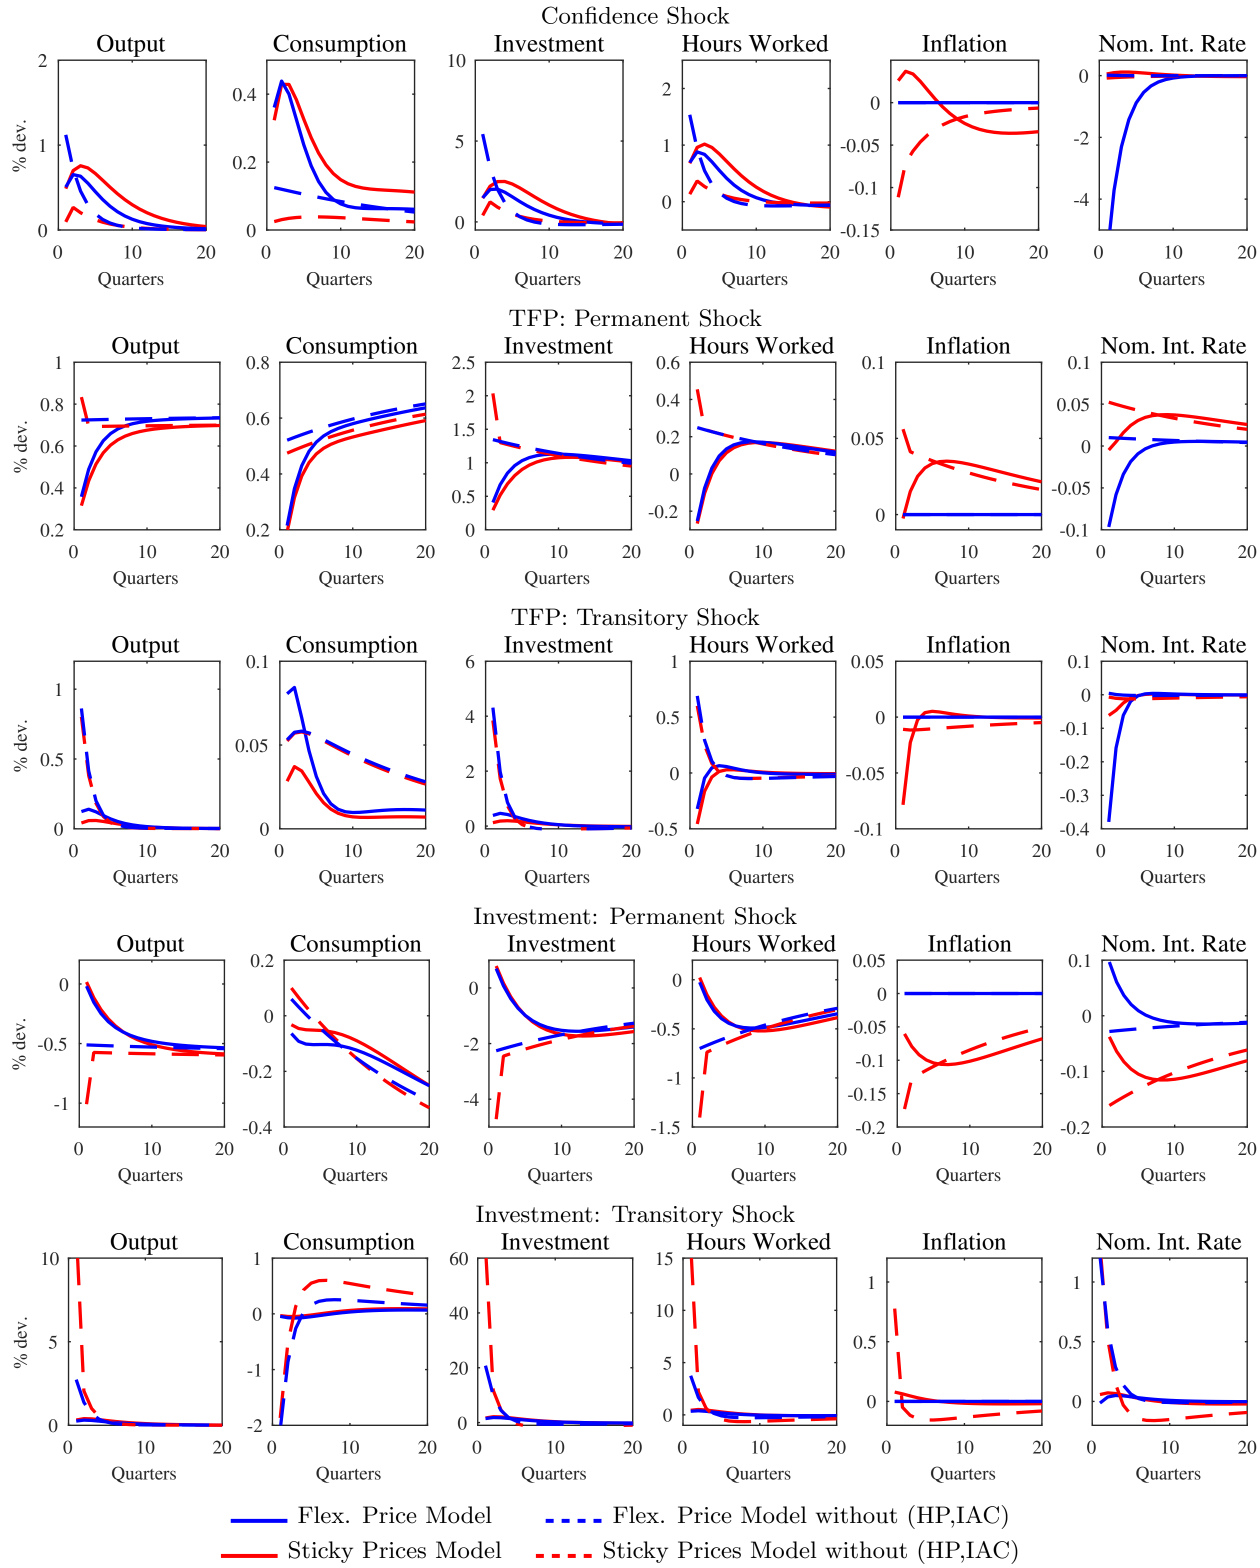

Figure 9: Theoretical IRFs, Part I

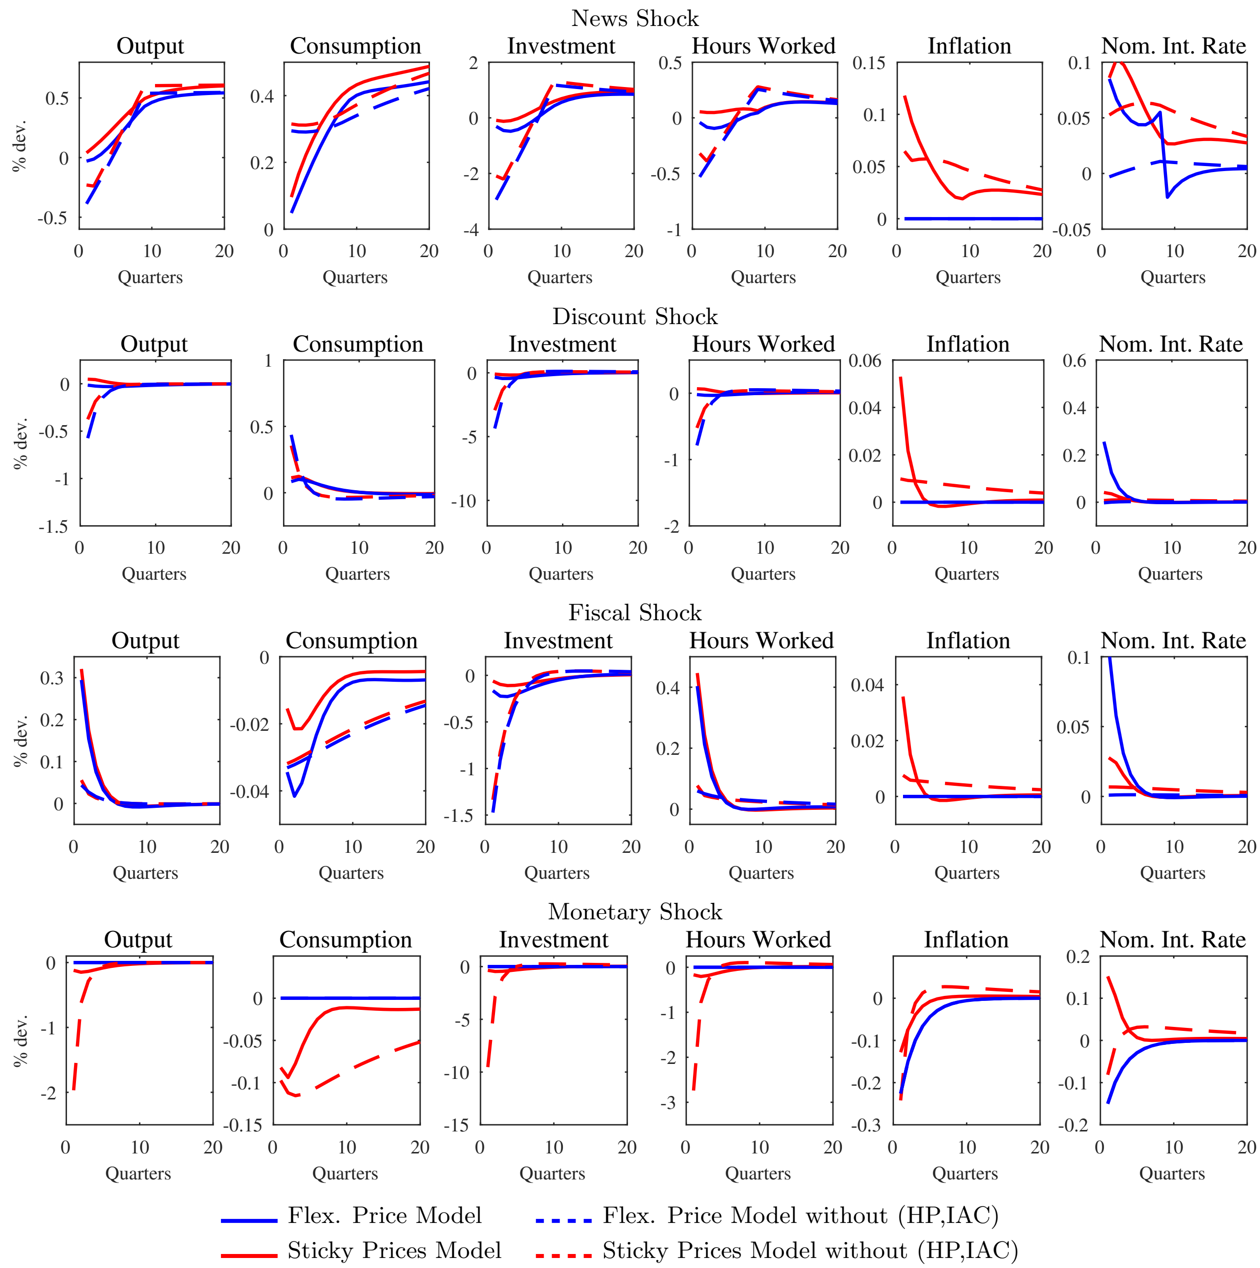

Figure 10: Theoretical IRFs, Part II

Table 13: Band-pass Filtered Moments

	Data	FP	SP	SW	Baseline	Data	FP	SP	SW	Baseline
	Standard Deviations					*Correlations with Output*				
y	1.41	1.32	1.43	1.42	1.52					
i	5.12	4.55	5.15	4.86	5.00	0.94	0.86	0.89	0.74	0.97
h	1.56	1.54	1.83	0.97	1.52	0.87	0.83	0.87	0.81	0.82
c	0.76	0.83	0.81	1.11	0.69	0.85	0.76	0.73	0.67	0.89
y/h	0.76	0.85	0.92	0.84	0.91	0.07	0.03	-0.18	0.74	0.29
π	0.23	0.21	0.25	0.34	–	0.21	0.01	0.46	0.13	–
R	0.35	5.35	0.32	0.35	–	0.33	-0.54	0.58	0.06	–
	Correlations with Investment					*Correlations with Hours*				
h	0.82	0.84	0.87	0.67	0.92					
c	0.73	0.44	0.44	0.30	0.76	0.83	0.52	0.52	0.59	0.98
y/h	0.07	-0.20	-0.36	0.47	0.08	-0.43	-0.52	-0.64	0.22	-0.30
π	0.09	0.01	0.41	0.18	–	0.44	0.00	0.59	0.23	–
R	0.23	-0.46	0.57	0.23	–	0.61	-0.63	0.69	0.21	–

Note: FP and SP: our estimated flexible- and sticky-price models, SW: the model in Smets and Wouters (2007), Baseline: the calibrated RBC model studied in Section 5.

Table 14: Contribution of Shocks to Volatilities (6–32 Quarters)

	Y	C	I	h	π	R
Permanent TFP Shock						
Flexible Price	24.15	35.05	5.16	4.14	0.00	0.03
Sticky Price	16.42	30.65	2.96	3.11	2.65	1.58
In the absence of belief shock						
Flexible Price	37.00	20.75	1.19	17.79	0.00	0.12
Sticky Price	23.90	55.22	3.39	4.82	0.78	1.20
Transitory TFP Shock						
Flexible Price	1.82	1.49	1.82	2.03	0.00	0.28
Sticky Price	0.32	0.32	0.28	3.14	4.69	3.64
In the absence of belief shock						
Flexible Price	2.32	0.27	0.97	8.00	0.00	5.48
Sticky Price	0.59	1.04	0.15	1.90	5.67	8.36
Permanent Investment Shock						
Flexible Price	7.88	1.27	12.62	10.63	0.00	0.03
Sticky Price	6.92	0.29	11.83	7.78	22.06	12.48
In the absence of belief shock						
Flexible Price	14.04	9.61	8.32	17.70	0.00	0.16
Sticky Price	23.72	5.39	27.52	33.72	4.30	6.71
Transitory Investment Shock						
Flexible Price	9.37	3.16	30.95	12.67	0.00	0.02
Sticky Price	13.47	2.42	32.98	15.58	15.42	11.48
In the absence of belief shock						
Flexible Price	1.21	0.01	0.98	1.55	0.00	0.00
Sticky Price	34.67	6.33	63.86	47.32	9.16	14.20

	Y	C	I	h	π	R
News Shock						
Flexible Price	8.58	11.33	7.11	2.29	0.00	0.04
Sticky Price	5.87	13.99	2.06	0.16	22.44	16.40
In the absence of belief shock						
Flexible Price	11.44	5.74	5.17	9.31	0.00	1.25
Sticky Price	11.35	24.14	4.44	1.74	31.85	49.45
Discount Shock						
Flexible Price	0.10	2.82	1.89	0.12	0.00	0.14
Sticky Price	0.18	3.88	0.22	0.19	2.56	1.99
In the absence of belief shock						
Flexible Price	28.17	63.61	83.31	37.96	0.00	90.98
Sticky Price	0.10	5.15	0.26	0.40	9.06	13.40
Fiscal Shock						
Flexible Price	3.28	0.39	0.50	4.37	0.00	0.03
Sticky Price	3.33	0.12	0.09	3.92	1.19	0.93
In the absence of belief shock						
Flexible Price	5.84	0.02	0.06	7.69	0.00	0.09
Sticky Price	5.33	0.53	0.09	8.78	0.94	1.43
Monetary Policy Shock						
Flexible Price	0.00	0.00	0.00	0.00	100.00	0.07
Sticky Price	1.79	1.94	1.52	2.21	20.55	19.00
In the absence of belief shock						
Flexible Price	0.00	0.00	0.00	0.00	100.00	1.91
Sticky Price	0.34	2.21	0.30	1.32	38.24	5.26

Table 15: Contribution of Shocks to Comovements (6–32 Quarters)

	Covariance					
	(Y,h)	(Y,I)	(Y,C)	(h,I)	(h,C)	(I,C)
Permanent TFP Shock						
Flexible Price	4.55	12.76	37.65	2.86	10.83	29.55
Sticky Price	2.56	7.48	30.20	1.89	6.88	20.14
In the absence of belief shock						
Flexible Price	-2.20	9.78	1781.85	0.63	10.29	7.20
Sticky Price	8.37	9.69	112.94	3.94	602.20	198.44
Transitory TFP Shock						
Flexible Price	-0.51	2.09	2.09	-0.29	-1.28	3.43
Sticky Price	-0.44	0.33	0.42	-0.30	-1.05	0.60
In the absence of belief shock						
Flexible Price	-0.91	2.22	50.78	-0.29	-1.16	0.71
Sticky Price	-1.02	0.32	2.42	-0.13	-79.79	5.56
Permanent Investment Shock						
Flexible Price	10.80	10.09	1.83	10.95	3.78	-0.29
Sticky Price	8.31	9.25	-0.44	9.25	-0.52	-2.03
In the absence of belief shock						
Flexible Price	26.37	8.12	612.09	7.62	29.60	-1.62
Sticky Price	19.45	27.53	-17.90	29.11	498.11	-25.21
Transitory Investment Shock						
Flexible Price	12.78	19.70	-6.74	23.02	-11.19	-21.09
Sticky Price	16.40	23.50	-7.04	25.42	-10.88	-18.09
In the absence of belief shock						
Flexible Price	2.27	1.61	-7.48	1.74	-0.32	-0.15
Sticky Price	56.57	55.79	-37.49	63.82	-1355.77	-320.72

	Covariance					
	(Y,h)	(Y,I)	(Y,C)	(h,I)	(h,C)	(I,C)
News Shock						
Flexible Price	3.83	7.00	9.66	4.75	0.93	3.26
Sticky Price	0.51	2.73	10.12	0.27	1.06	2.16
In the absence of belief shock						
Flexible Price	10.06	7.74	315.44	9.81	-5.73	-1.58
Sticky Price	5.27	6.73	40.94	3.20	154.88	42.54
Discount Shock						
Flexible Price	0.12	0.46	-0.57	0.57	-1.07	-5.00
Sticky Price	0.22	-0.11	1.02	-0.12	1.47	-1.65
In the absence of belief shock						
Flexible Price	53.15	70.93	-2639.62	80.92	-132.06	-104.60
Sticky Price	0.28	-0.03	1.73	-0.12	102.85	-14.99
Fiscal Shock						
Flexible Price	4.53	-0.85	-1.14	-1.03	-1.93	0.93
Sticky Price	4.16	-0.28	-0.60	-0.31	-0.90	0.21
In the absence of belief shock						
Flexible Price	11.27	-0.40	-13.06	-0.45	-0.61	0.04
Sticky Price	10.18	-0.41	-4.71	-0.51	-158.24	2.98
Monetary Policy Shock						
Flexible Price	0.00	0.00	0.00	0.00	0.00	0.00
Sticky Price	2.29	1.83	2.44	2.05	3.82	3.48
In the absence of belief shock						
Flexible Price	0.00	0.00	0.00	0.00	0.00	0.00
Sticky Price	0.90	0.38	2.07	0.68	135.76	11.40

E. 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 we consider the general case, it is useful to review the linearized version of our baseline model. This helps fix some key ideas.

As noted in Section 4, the log-linearized equilibrium conditions of the model are given by (7)-(11). The interpretation of these conditions is familiar. The only peculiarity is the presence of two distinct expectation operators $\mathbb{E}_{it}$ and $\mathbb{E}'_{it}$, which denote the local expectations, respectively, stage 1 and stage 2 of period t. 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. Along with the timing convention we have adopted, this explains why 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. The method we develop in the sequel 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 the sequel, 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:

$$\begin{aligned}
M_{yy}y_{it} &= M_{yx}x_{it}^b + M_{yX}X_t^b + M_{yY}\mathbb{E}_{it}Y_t + M_{yf}\mathbb{E}_{it}x_{it}^f + M_{yF}\mathbb{E}_{it}X_t^f + M_{ys}z_{it} \\
M_{xx0}x_{it+1}^b &= M_{xx1}x_{it}^b + M_{xX1}X_t^b + M_{xy1}y_{it} + M_{xY1}Y_t + M_{xf1}x_{it}^f + M_{xF1}X_t^f + M_{xs1}s_t \\
M_{ff0}\mathbb{E}'_{it}x_{it+1}^f &= M_{fF0}\mathbb{E}'_{it}X_{t+1}^f + M_{ff1}x_{it}^f + M_{fF1}X_t^f + M_{fx0}x_{it+1}^b + M_{fx1}x_{it}^b + M_{fX1}X_t^b \\
&\quad + M_{fy0}\mathbb{E}'_{it}y_{it+1} + M_{fY0}\mathbb{E}'_{it}Y_{t+1} + M_{fy1}y_{it} + M_{fY1}Y_t + M_{fs0}\mathbb{E}'_{it}s_{t+1} + M_{fs1}s_t \\
s_t &= Rs_{t-1} + \varepsilon_t \\
\xi_t &= Q\xi_{t-1} + \nu_t
\end{aligned}$$

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 variable 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 4.) 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:

$$\begin{aligned} Y_t &= M_X X_t^b + M_{EY} Y_t + M_F X_t^f + M_s s_t \\ X_{t+1}^b &= N_X X_t^b + N_Y Y_t + N_F X_t^f + N_s s_t \\ (P_{f0} - P_{F0})\mathbb{E}_t X_{t+1}^f &= P_{F1} X_t^f + P_{Y0}\mathbb{E}_t Y_{t+1} + P_X X_t^b + P_{Y1} Y_t + P_s s_t \\ s_t &= R s_{t-1} + \varepsilon_t \\ \xi_t &= Q \xi_{t-1} + \nu_t \end{aligned}$$

(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:

$$y_{it} = M_x(x^b_{it} - X^b_t) + M_X X^b_t + M_{EY}\mathbb{E}_{it}Y_t + M_f\mathbb{E}_{it}(x^f_{it} - X^f_t) + M_F\mathbb{E}_{it}X^f_t + M_s z_{it} \quad (42)$$

$$x^b_{it+1} = N_x(x^b_{it} - X^b_t) + N_X X^b_t + N_y(y_{it} - Y_t) + N_Y Y_t + N_f(x^f_{it} - X^f_t) + N_F X^f_t + N_s s_t \quad (43)$$

$$P_{f0}\mathbb{E}'_{it}x^f_{it+1} = P_{f1}(x^f_{it} - X^f_t) + P_{F0}\mathbb{E}'_{it}X^f_{t+1} + P_{F1}X^f_t + P_x(x^b_{it} - X^b_t) + P_X X^b_t +$$
$$+ P_{y0}(\mathbb{E}'_{it}y_{it+1} - \mathbb{E}'_{it}Y_{t+1}) + P_{Y0}\mathbb{E}'_{it}Y_{t+1} + P_{y1}(y_{it} - Y_t) + P_{Y1}Y_t + P_s s_t \quad (44)$$

where

$$M_x = M^{-1}_{yy}M_{yx},\ M_X = M^{-1}_{yy}(M_{yx} + M_{yX}),\ M_{EY} = M^{-1}_{yy}M_{yY},$$
$$M_f = M^{-1}_{yy}M_{yf},\ M_F = M^{-1}_{yy}(M_{yf} + M_{yF}),\ M_s = M^{-1}_{yy}M_{ys}$$
$$N_x = M^{-1}_{xx0}M_{xx1},\ N_X = M^{-1}_{xx0}(M_{xx1} + M_{xX1}),\ N_y = M^{-1}_{xx0}M_{xy1},\ N_Y = M^{-1}_{xx0}(M_{xy1} + M_{xY1}),$$
$$N_f = M^{-1}_{xx0}M_{xf1},\ N_F = M^{-1}_{xx0}(M_{xf1} + M_{xF1}),\ N_s = M^{-1}_{xx0}M_{xs1}$$
$$P_{f0} = M_{ff0},\ P_{f1} = M_{ff1} + M_{fx0}N_f,\ P_{F0} = M_{fF0},\ P_{F1} = M_{fF1} + M_{ff1} + M_{fx0}N_F$$
$$P_x = M_{fx1} + M_{fx0}N_x,\ P_X = M_{fX1} + M_{fx1} + M_{fx0}N_X,$$
$$P_{y0} = M_{fy0},\ P_{Y0} = M_{fY0} + M_{fy0},\ P_{y1} = M_{fy1} + M_{fx0}N_y,\ P_Y = M_{fy1} + M_{fY1} + M_{fx0}N_Y,$$
$$P_s = M_{fs0}R + M_{fs1} + M_{fx0}N_s$$

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

$$y_{it} = \Lambda_x(x^b_{it} - X^b_t) + \Lambda_X X^b_t + \Lambda_z z_{it} + \Lambda_\xi \xi_t \quad (45)$$
$$x^f_{it} = \Gamma_x(x^b_{it} - X^b_t) + \Gamma_X X^b_t + \Gamma_z z_{it} + \Gamma_{\bar{z}}\bar{z}_t + \Gamma_s s_t + \Gamma_\xi \xi_t \quad (46)$$

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 (43) and (44), *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 (46); 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 (45) that solves the fixed-point relation between y_{it} and Y_t described in (42) 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 (43) and (44) under the assumption that that the stage-1 outcomes are determined according to an arbitrary rule in (46). 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 (46) 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 (46), 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

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

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 (42), 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,

$$\begin{aligned} y_{it} &= (M_x + M_f\Gamma_x)(x_{it}^b - X_t^b) + (M_X + M_F\Gamma_X)X_t^b + M_{EY}\mathbb{E}_{it}Y_t \\ &\quad + (M_s + M_F(\Gamma_z + \Gamma_{\bar{z}} + \Gamma_s))z_{it} + (M_F\Gamma_\xi + M_F\Gamma_{\bar{z}}\Delta + (M_F - M_f)\Gamma_z\Delta)\,\xi_t \end{aligned} \tag{47}$$

Note that the above represents us 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 (45). 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 (47), we obtain the following best response for island i:

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

For this to be consistent with our guess in (45), we must have

$$\Lambda_x = M_x + M_f\Gamma_x \tag{48}$$

$$\Lambda_X = (I - M_{EY})^{-1}(M_X + M_F\Gamma_X) \tag{49}$$

$$\Lambda_z = (I - M_{EY})^{-1}\left[M_s + M_F(\Gamma_z + \Gamma_{\bar{z}} + \Gamma_s)\right] \tag{50}$$

$$\Lambda_\xi = (I - M_{EY})^{-1}\left\{M_F(\Gamma_\xi + \Gamma_{\bar{z}}\Delta) + (M_F - M_f)\Gamma_z\Delta + M_{EY}\Lambda_z\Delta\right\} \tag{51}$$

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:

$$\begin{aligned} Y_{t+1} &= \Lambda_X X_{t+1}^b + \Lambda_z \bar{z}_{t+1} + \Lambda_\xi \xi_{t+1} \\ \mathbb{E}_{it+1}Y_{t+1} &= \Lambda_X X_{t+1}^b + \Lambda_z z_{it+1} + (\Lambda_\xi + \Lambda_z\Delta)\xi_{t+1} \\ \mathbb{E}'_{it}Y_{t+1} &= \Lambda_X X_{t+1}^b + \Lambda_z R s_t + (\Lambda_\xi + \Lambda_z\Delta)Q\xi_t \end{aligned}$$

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

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

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:

$$\begin{aligned} x^f_{it+1} &= \Gamma_x(x^b_{it+1} - X^b_{t+1}) + \Gamma_X X^b_{t+1} + \Gamma_z z_{it+1} + \Gamma_{\bar{z}} \bar{z}_{t+1} + \Gamma_s s_{t+1} + \Gamma_\xi \xi_{t+1} \\ \mathbb{E}_{it+1} x^f_{it+1} &= \Gamma_x(x^b_{it+1} - X^b_{t+1}) + \Gamma_X X^b_{t+1} + (\Gamma_z + \Gamma_{\bar{z}} + \Gamma_s) z_{it+1} + (\Gamma_\xi + \Gamma_{\bar{z}}\Delta)\xi_{t+1} \\ \mathbb{E}'_{it} x^f_{it+1} &= \Gamma_x(x^b_{it+1} - X^b_{t+1}) + \Gamma_X X^b_{t+1} + (\Gamma_z + \Gamma_{\bar{z}} + \Gamma_s) R s_t + (\Gamma_\xi + \Gamma_{\bar{z}}\Delta) Q \xi_t \end{aligned}$$

For the aggregate next-period forward variables we have

$$\begin{aligned} \mathbb{E}_{it+1} X^f_{t+1} &= \Gamma_X X^b_{t+1} + (\Gamma_z + \Gamma_{\bar{z}} + \Gamma_s) z_{it+1} + (\Gamma_\xi + (\Gamma_z + \Gamma_{\bar{z}})\Delta)\xi_{t+1} \\ \mathbb{E}'_{it} X^f_{t+1} &= \Gamma_X X^b_{t+1} + (\Gamma_z + \Gamma_{\bar{z}} + \Gamma_s) R s_t + (\Gamma_\xi + (\Gamma_z + \Gamma_{\bar{z}})\Delta) Q \xi_t \end{aligned}$$

and therefore

$$\mathbb{E}'_{it}(x^f_{it+1} - X^f_{t+1}) = \Gamma_x(x^b_{it+1} - X^b_{t+1}) - \Gamma_z \Delta Q \xi_t$$

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

$$\begin{aligned} y_{it} - Y_t &= \Lambda_x(x^b_{it} - X^b_t) + \Lambda_z(z_{it} - \bar{z}_t) \\ x^f_{it} - X^f_t &= \Gamma_x(x^b_{it} - X^b_t) + \Gamma_z(z_{it} - \bar{z}_t) \\ Y_t &= \Lambda_X X^b_t + \Lambda_z \bar{z}_t + \Lambda_\xi \xi_t \\ X^f_t &= \Gamma_X X^b_t + (\Gamma_z + \Gamma_{\bar{z}})\bar{z}_t + \Gamma_s s_t + \Gamma_\xi \xi_t \end{aligned}$$

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

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

Equivalently,

$$x^b_{it+1} = \Omega_x(x^b_{it} - X^b_t) + \Omega_X X^b_t + \Omega_z z_{it} + \Omega_{\bar{z}} \bar{z} + \Omega_s s_t + \Omega_\xi \xi_t$$

and hence

$$\begin{aligned} X^b_{t+1} &= \Omega_X X^b_t + (\Omega_z + \Omega_{\bar{z}})\bar{z}_t + \Omega_s s_t + \Omega_\xi \xi_t \\ x^b_{it+1} - X^b_{t+1} &= \Omega_x(x^b_{it} - X^b_t) + \Omega_z(z_{it} - \bar{z}_t) \end{aligned}$$

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

$$\begin{aligned}\mathbb{E}'_{it}x^f_{it+1} &= \Gamma_x(x^b_{it+1} - X^b_{t+1}) + \Gamma_X X^b_{t+1} + (\Gamma_z + \Gamma_{\bar{z}} + \Gamma_s)Rs_t + (\Gamma_\xi + \Gamma_{\bar{z}}\Delta)Q\xi_t \\ &= \Gamma_x\left\{\Omega_x(x^b_{it} - X^b_t) + \Omega_z(z_{it} - \bar{z}_t)\right\} + \Gamma_X\left\{\Omega_X X^b_t + (\Omega_z + \Omega_{\bar{z}})\bar{z}_t + \Omega_s s_t + \Omega_\xi\xi_t\right\} \\ &\quad + (\Gamma_z + \Gamma_{\bar{z}} + \Gamma_s)Rs_t + (\Gamma_\xi + \Gamma_{\bar{z}}\Delta)Q\xi_t\end{aligned}$$

or equivalently

$$\mathbb{E}'_{it}x^f_{it+1} = \Phi_x(x^b_{it} - X^b_t) + \Phi_X X^b_t + \Phi_z z_{it} + \Phi_{\bar{z}}\bar{z}_t + \Phi_s s_t + \Phi_\xi\xi_t \tag{52}$$

where

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

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

$$\mathbb{E}'_{it}X^f_{t+1} = \Phi_X X^b_t + \Phi_z z_{it} + \Phi_{\bar{z}}\bar{z}_t + \Phi_s s_t + (\Phi_\xi + \Gamma_z\Delta Q)\xi_t \tag{53}$$

With the above steps, we have calculated all the objects that enter the Euler condition (44). 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 (44) below:

$$\begin{aligned}P_{f0}\mathbb{E}'_{it}x^f_{it+1} &= P_{f1}(x^f_{it} - X^f_t) + P_{F0}\mathbb{E}'_{it}X^f_{t+1} + P_{F1}X^f_t + P_x(x^b_{it} - X^b_t) + P_X X^b_t + \\ &\quad + P_{y0}(\mathbb{E}'_{it}y_{it+1} - \mathbb{E}'_{it}Y_{t+1}) + P_{Y0}\mathbb{E}'_{it}Y_{t+1} + P_{y1}(y_{it} - Y_t) + P_{Y1}Y_t + P_s s_t\end{aligned}$$

Use now (52) 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

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

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

$$
\begin{aligned}
P_{f0}\Phi_x &= P_x + P_{f1}\Gamma_x + P_{y0}\Lambda_x\Omega_x + P_{y1}\Lambda_x && (54)\\
(P_{f0} - P_{F0})\,\Phi_X &= P_{F1}\Gamma_X + P_X + P_{Y0}\Lambda_X\Omega_X + P_{Y1}\Lambda_X && (55)\\
P_{f0}\Phi_z &= P_{f1}\Gamma_z + P_{y0}\Lambda_x\Omega_z + P_{y1}\Lambda_z && (56)\\
(P_{f0} - P_{F0})\,\Phi_{\bar{z}} &= P_{F0}\Phi_z + (P_{F1} - P_{f1})\Gamma_z + P_{F1}\Gamma_{\bar{z}} + P_{Y0}\Lambda_X(\Omega_z + \Omega_{\bar{z}}) \\
&\quad - P_{y0}\Lambda_x\Omega_z + (P_{Y1} - P_{y1})\Lambda_z && (57)\\
(P_{f0} - P_{F0})\Phi_s &= P_{F1}\Gamma_s + P_{Y0}(\Lambda_X\Omega_s + \Lambda_z R) + P_s && (58)\\
(P_{f0} - P_{F0})\,\Phi_\xi &= P_{F0}\Gamma_z\Delta Q + P_{F1}\Gamma_\xi + P_{Y0}\left\{\Lambda_X\Omega_\xi + \Lambda_\xi Q\right\} + (P_{Y0} - P_{y0})\Lambda_z\Delta Q + P_{Y1}\Lambda_\xi && (59)
\end{aligned}
$$

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 (48)-(51) and system (54)-(59). 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

$$
\begin{aligned}
Y_t &= \Lambda_X^* X_t^b + \Lambda_s^* s_t \\
X_t^f &= \Gamma_X^* X_t^b + \Gamma_s^* s_t
\end{aligned}
$$

It is straightforward to check that the solution to our model satisfies the following:

$$
\begin{aligned}
\Lambda_X &= \Lambda_X^* & \Lambda_z &= \Lambda_s^* \\
\Gamma_X &= \Gamma_X^* & \Gamma_z + \Gamma_{\bar{z}} + \Gamma_s &= \Gamma_s^*
\end{aligned}
$$

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.

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 (58):

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

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

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

and therefore $\Gamma_s = A_S^{-1} B_S$. Next, Γ_z can be obtained from (56):

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

Plugging the definition of Φ_z and Ω_z in the above, 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 (48) and (54) 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 (51) and (59), which we repeat below:

$$\begin{aligned} \Lambda_\xi &= (I - M_{EY})^{-1}\{M_F(\Gamma_\xi + \Gamma_{\bar{z}}\Delta) + (M_F - M_f)\Gamma_z\Delta + M_{EY}\Lambda_z\Delta\} \\ (P_{f0} - P_{F0})\,\Phi_\xi &= P_{F0}\Gamma_z\Delta Q + P_{F1}\Gamma_\xi + P_{Y0}\{\Lambda_X\Omega_\xi + \Lambda_\xi Q\} + (P_{Y0} - P_{y0})\Lambda_z\Delta Q + P_{Y1}\Lambda_\xi \end{aligned}$$

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

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

to substitute away also Ω_ξ and Φ_ξ. We then obtain a single equation in Γ_ξ, which takes the following form:

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

where

$$\begin{aligned} A \equiv& (P_{F0} - P_{f0}) + P_{Y0}(I - M_{EY})^{-1}M_F \\ B \equiv& ((P_{F0} - P_{f0})\Gamma_X N_Y + P_{Y0}\Lambda_X N_Y + P_{Y1})(I - M_{EY})^{-1}M_F + (P_{F0} - P_{f0})\Gamma_X N_F + P_{F1} + P_{Y0}\Lambda_X N_F \\ C \equiv& \left(P_{F0}\Gamma_z\Delta Q + (P_{Y0} - P_{y0})\Lambda_z + (P_{F0} - P_{f0})\Gamma_{\bar{z}} + P_{Y0}(I - M_{EY})^{-1}[M_F\Gamma_{\bar{z}} + (M_F - M_f)\Gamma_z + M_{EY}\Lambda_z]\right)\Delta Q \\ &+ ((P_{F0} - P_{f0})\Gamma_X N_Y + P_{Y0}\Lambda_X N_Y + P_{Y1})(I - M_{EY})^{-1}[M_F\Gamma_{\bar{z}} + (M_F - M_f)\Gamma_z + M_{EY}\Lambda_z]\,\Delta \end{aligned}$$

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 (51). This completes the solution.

References

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

Allen, F., S. Morris, and A. Postlewaite (1993): "Finite Bubbles with Short Sale Constraints and Asymmetric Information," *Journal of Economic Theory*, 61(2), 206–229.

Alvarez, F. E., F. Lippi, and L. Paciello (2011): "Optimal Price Setting With Observation and Menu Costs," *The Quarterly Journal of Economics*, 126(4), 1909–1960.

Amador, M., and P.-O. Weill (2010): "Learning from Prices: Public Communication and Welfare," *Journal of Political Economy*, 118(5), 866 – 907.

Angeletos, G., F. Collard, and H. Dellas (2016): "An Anatomy of the Business Cycle," miméo, MIT.

Angeletos, G., and J. La'O (2013): "Sentiments," *Econometrica*, 81(2), 739–779.

Angeletos, G., and C. Lian (2016a): "Incomplete Information in Macroeconomics: Accommodating Frictions in Coordination," NBER Working Paper 22297.

——— (2016b): "Forward Guidance without Common Knowledge," miméo, MIT.

Angeletos, G., and A. Pavan (2007): "Efficient Use of Information and Social Value of Information," *Econometrica*, 75(4), 1103–1142.

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

Beaudry, P., D. Galizia, and F. Portier (2015): "Reviving the Limit Cycle View of Macroeconomic Fluctuations," miméo, University of British Columbia.

Beaudry, P., and F. 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*, pp. 69–130.

Benhabib, J., and R. E. A. Farmer (1994): "Indeterminacy and Increasing Returns," *Journal of Economic Theory*, 63(1), 19–41.

Benhabib, J., P. Wang, and Y. Wen (2015): "Sentiments and Aggregate Demand Fluctuations," *Econometrica*, 83, 549–585.

Bergemann, D., and S. Morris (2013): "Robust Predictions in Games With Incomplete Information," *Econometrica*, 81(4), 1251–1308.

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

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

Bloom, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77(3), 623–685.

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. J. Terry (2012): "Really Uncertain Business Cycles," NBER Working Paper 18245.

Buera, F. J., and B. Moll (2012): "Aggregate Implications of a Credit Crunch," NBER Working Paper 17775.

Cass, D., and K. Shell (1983): "Do Sunspots Matter?," *Journal of Political Economy*, 91(2), 193–227.

Chahrour, R., and G. Gaballo (2015): "On the Nature and Stability of Sentiments," mimeo, Boston College.

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

Christiano, L., C. Ilut, R. Motto, and M. Rostagno (2008): "Monetary Policy and Stock Market Boom-Bust Cycles," Working Paper Series 0955, European Central Bank.

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

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

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

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

David, J. M., H. A. Hopenhayn, and V. Venkateswaran (2014): "Information, Misallocation and Aggregate Productivity," NBER Working Paper 20340.

Diamond, P. A. (1982): "Aggregate Demand Management in Search Equilibrium," *Journal of Political Economy*, 90(5), 881–894.

Eusepi, S., and B. Preston (2011): "Expectations, Learning, and Business Cycle Fluctuations," *American Economic Review*, 101(6), 2844–72.

Evans, G., and S. Honkapohja (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, NJ.

Farmer, R. E. A. (2012): "Confidence, Crashes and Animal Spirits," *Economic Journal*, 122(559), 155–172.

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

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

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

Guesnerie, R. (1992): "An Exploration of the Eductive Justifications of the Rational-Expectations Hypothesis," *American Economic Review*, 82(5), 1254–78.

Hannan, E. J. (1970): *Multiple Time Series*, Wiley Series in Probability and Statistics. Wiley, John & Sons.

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

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

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

Huo, Z., and N. Takayama (2015a): "Rational Expectations Models with Higher Order Beliefs," miméo, Yale University.

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

Ilut, C., and H. Saijo (2016): "Learning, Confidence and Business Cycle," miméo, Duke University.

Ilut, C. L., and M. Schneider (2014): "Ambiguous Business Cycles," *American Economic Review*, 104(8), 2368–99.

Izmalkov, S., and M. Yildiz (2010): "Investor Sentiments," *American Economic Journal: Microeconomics*, 2(1), 21–38.

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

Justiniano, A., G. E. Primiceri, and A. Tambalotti (2010): “Investment Shocks and Business Cycles,” *Journal of Monetary Economics*, 57(2), 132–145.

Levchenko, A. A., and N. Pandalai-Nayar (2015): “TFP, News, and “Sentiments”: The International Transmission of Business Cycles,” NBER Working Paper 21010.

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

Lucas, R. J. (1972): “Expectations and the neutrality of money,” *Journal of Economic Theory*, 4(2), 103–124.

Mackowiak, B., and M. Wiederholt (2009): “Optimal Sticky Prices under Rational Inattention,” *American Economic Review*, 99(3), 769–803.

Maćkowiak, B., and M. Wiederholt (2015): “Business Cycle Dynamics under Rational Inattention,” *Review of Economic Studies*, 82(4), 1502–1532.

Mankiw, G. N., and R. Reis (2002): “Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve,” *The Quarterly Journal of Economics*, 117(4), 1295–1328.

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

Morris, S., and H. S. Shin (1998): “Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks,” *American Economic Review*, 88(3), 587–597.

——— (2001): “Rethinking Multiple Equilibria in Macroeconomic Modeling,” in *NBER Macroeconomics Annual 2000, Volume 15*, pp. 139–182. National Bureau of Economic Research, Inc.

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

Myohl, C., and Y. Stucki (2016): “Confidence and Financial Frictions,” miméo, University of Bern.

Nimark, K. (2011): “Dynamic Higher Order Expectations,” Working Paper 1118, Universitat Pompeu Fabra.

Phelps, E. (1971): “Inflation Expectations and Economic Theory,” in *Inflation and the Canadian Experience*, pp. 31–47. Industrial Relation Centre, Queen’s University, Kingston, Ontario.

Reis, R. (2006): “Inattentive Consumers,” *Journal of Monetary Economics*, 53(8), 1761–1800.

Rondina, G., and T. B. Walker (2014): "Dispersed Information and Confounding Dynamics," miméo, University of Colorado, Boulder.

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

Schaal, E., and M. Taschereau-Dumouchel (2015): "Coordinating Business Cycles," 2015 meeting papers, Society for Economic Dynamics.

Scheinkman, J. A., and W. Xiong (2003): "Overconfidence and Speculative Bubbles," *Journal of Political Economy*, 111(6), 1183–1219.

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

Stock, J. H., and M. W. Watson (1999): "Business Cycle Fluctuations in US Macroeconomic Time Series," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, vol. 1 of *Handbook of Macroeconomics*, chap. 1, pp. 3–64. Elsevier.

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

Woodford, M. (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.

——— (2013): "Macroeconomic Analysis Without the Rational Expectations Hypothesis," *Annual Review of Economics*, 5(1), 303–346.