Quantifying Confidence

George-Marios Angeletos
MIT and University of Zurich

Fabrice Collard
University of Bern

Harris Dellas
University of Bern

July 31, 2014

preliminary and incomplete

check *http://economics.mit.edu/files/9755* for the latest update

Abstract

We develop a tractable methodology for incorporating higher-order belief dynamics in a large class of DSGE models. This helps us accommodate a a new type of volatility, which we interpret as variation in "confidence" and as symptom of imperfect coordination. We next quantify the role of this volatility in US business cycles under the lenses of both RBC and NK models. We find that it helps match various salient features of the data, from moments and wedges to VAR evidence regarding the sources of the business cycle; it accounts for a significant fraction of the overall volatility in estimated versions of these models that allow for multiple competing structural shocks; and it operationalizes the notion that recessions are periods of weak aggregate demand without the usual reliance on nominal rigidities and policy frictions.

1 Introduction

Recessions are often described as periods of "low confidence" and "weak aggregate demand". But what exactly do these words mean? Are they just buzz words, or do they have a precise and useful meaning within formal models? And how important are changes in fundamentals, such as technology, or nominal rigidities, such as sticky prices, in understanding these notions and in explaining the observed fluctuations?

Macroeconomists have long disagreed on the answers to these questions and on their policy implications. Nonetheless, macroeconomists have reached a certain consensus with regard to the methodological framework they use in order to interpret the data and address these questions. A central feature of this framework, from which we wish to depart in this paper, is that the economy is modeled, in effect, as a complete-information game in which players face no strategic uncertainty and no frictions in coordinating their beliefs and actions.

What explains the ubiquity of this modeling feature is its convenience, not its plausibility. The notion that the millions of firms and households whose interaction determines macroeconomic outcomes can coordinate their beliefs and actions as flawlessly as in representative-agent settings is obviously a gross abstraction. Macroeconomists are accustomed to this abstraction because they appreciate the tractability and the versatility of the models that they can build with its help, as well as because they lack evidence of the quantitative value of departing from it.

By contrast, theorists have sought to operationalize the notion that coordination is imperfect by introducing incomplete information and higher-order uncertainty. The basic insight is that higher-order uncertainty precludes the players from reaching common belief about one another's choices, thus also inhibiting coordination.[1] In so doing, theorists have also formalized a distinct form of volatility—one that can be interpreted as the symptom of frictions in coordination and of shifts in the expectations that agents form regarding one another's choices for given fundamentals (payoffs).

This suggests that the introduction of higher-order uncertainty in macroeconomic models may offer a novel structural interpretation of business-cycle phenomena, along with a certain formalization of the slippery notion of "confidence". The objective of this paper is to elaborate this structural interpretation; to evaluate its quantitative content; and to uncover its implications for the questions mentioned in the top of this page.

Empirical backdrop. Complementing the preceding theoretical motivation, we offer evidence that points towards a different type of volatility in the data than those formalized in macro models. We use a VAR-based strategy that helps us encapsulate the bulk of the observed fluctuations into a single empirical shock, without imposing any structural interpretation of it. The observable properties of this shock are unlike those of the structural shocks that are prominent in the macro literature. An integral part of our contribution in then to show how higher-order uncertainty offers a parsimonious yet potent explanation of this and other salient features of the data.

[1] A complementary approach is to consider weaker solution concepts, such as correlated equilibrium and rationalizability. These two approaches turn out to be equivalent in the sense formalized by Weinstein and Yildiz (2007).

Methodological contribution. Our paper is not the first to consider higher-order beliefs in a macro context. Important precedents include Morris and Shin (2002), who study how higher-order uncertainty matters for the social value of information, and Woodford (2002), who show how higher-order uncertainty can induce inertia in the response of prices to monetary shocks. More recently, Angeletos and La'O (2013) argue that higher-order uncertainty also allows unique-equilibrium models to accommodate "animal spirits", paving the way for what we do in this paper.[2]

Despite these theoretical advances, progress has been limited on the quantitative front. This is because the introduction of informational frictions and higher-order uncertainty in DSGE models raises two technical difficulties—Kalman filtering and infinite state spaces —that hinder the solution and estimation of these models. The broader methodological contribution of our paper rests on bypassing these difficulties: we propose a heterogeneous-prior formulation that permits us to accommodate rich yet tractable higher-order belief dynamics within a large class of DSGE.

In particular, we first characterize the equilibrium as a fixed-point problem between perceived and actual aggregate policy rules. We then obtain the log-linearized solution of the belief-augmented model as a simple transformation of the log-linearized solution of the original model.

To illustrate, consider the baseline RBC model. As well known, the linearized policy rule for the capital stock takes the form $\tilde{k}_{t+1} = \Gamma_k \tilde{k}_t + \Gamma_a \tilde{a}_t$, where $\tilde{k}_t$ is the capital stock and $\tilde{a}_t$ is the technology shock. In general, adding incomplete information causes the state space of the model to explode, because of the "forecasting-the-forecasts-of-others" problem. By contrast, our heterogeneous-prior specification guarantees that the state space expands only by one dimension: the policy rule becomes $\tilde{k}_{t+1} = \Gamma_k \tilde{k}_t + \Gamma_a \tilde{a}_t + \Gamma_\xi \xi_t$, where ξ_t is an exogenous state variable that conveniently summarizes any difference between first- and higher-order beliefs.

Applied contribution. The variable ξ_t introduces variation in the expectations that firms and consumers form about one another's choices, for given preferences, technologies, and other fundamentals. We interpret this as variation in "confidence" (accordingly, we henceforth refer to ξ_t as the "confidence shock") and proceed to evaluate its quantitative importance.

In the first part of the paper, we quantify this kind of confidence-driven fluctuations within a calibrated RBC model. We allow for monopolistic competition in order to link higher-order beliefs to expectations of aggregate demand, but assume flexible prices in order to guarantee that the notion of aggregate demand we accommodate is of a non-monetary nature.

In spite of its parsimony, the model has no difficulty in matching standard business-cycle moments, including: the co-movement of hours, output, investment, and consumption; the high relative volatility of hours; and the nearly zero, if not negative, correlation between hours and labor productivity. This success is not trivial: if we replace the confidence shock with a news shock, a discount-factor shock, an investment-specific technology shock, or certain proxies of financial shocks, the empirical fit of the model is compromised. What explains the success is the distinct way in which confidence shocks impact on expectations, incentives, and thereby on allocations.

[2]See also Bergemman and Morris (2013) and Bergemman, Heumann and Morris (2014) for a systematic investigation of how incomplete information expands the predictions of a class of static games that have a macro flavor.

When a negative confidence shock hits, firms expect aggregate output and hence the demand for their product to be low, which motivates them to hire less labor, to rent less capital, and to decrease utilization. This leads to an actual reduction in output and income, which in turn motivates consumers to cut down on both consumption and saving. All in all, variation in confidence therefore causes employment, output, consumption, and investment to move in the same direction, without a corresponding co-movement in TFP or labor productivity.

Essentially for the same reasons, the confidence shock is a look-alike of the aforementioned empirical shock that encapsulates most of the business cycle in the data. In addition, it tracks empirical measures of confidence such as the University of Michigan Index of Consumer Sentiment. We interpret these facts as additional validations of our approach.

Under the lenses of business-cycle accounting (Chari, Kehoe, McGrattan, 2005), variation in confidence registers as variation in the labor and Euler wedges: a negative confidence shock shows up as a "tax" on both employment and savings. Our approach therefore provides also a simple explanation of the cyclicality and the co-movement of the measured wedges.[3]

All in all, the first part of the paper shows that variation in confidence, modeled as higher-order uncertainty, helps account for a number of salient features of the data even within the context of a plain-vanilla RBC model. The second part of the paper elaborates on the robustness of this permanence as one moves form the RBC to the NK framework and on the quantitative importance of confidence within richer, estimated, models that accommodate multiple business-cycle forces.

More specifically, we estimate, with Bayesian Maximum Likelihood in the frequency domain, both flexible- and sticky-price specifications of an extended model that lets the confidence shock compete with up to ten other types of shocks: permanent and transitory TFP shocks; permanent and transitory investment-specific shocks; news shocks regarding the two types of technology; fiscal and monetary shocks; a markup shock; and a discount-rate shock.[4]

Despite these competing forces, variation in confidence accounts for a significant component of the business cycle: depending on which specification we consider, the estimated confidence shock explains 39% to 55% of the volatility in GDP, and up to 69% of the volatility in hours.

This result is not driven by the choice of Bayesian priors in the estimation: at the priors, the confidence shock explains less than 5% of GDP volatility. Rather, what drives the result is that the confidence shock acts as an "aggregate demand force" in a way that is unmatched by other shocks.

By the same token, the estimated confidence shock continues to be a look-alike of the empirical shock that, according to the aforementioned VAR evidence, captures most of the business-cycle movements in the data. It also continues tracks empirical measures of confidence.

We conclude the paper by revisiting the role of nominal rigidity. In the absence of confidence shocks, nominal rigidity significantly improves the fit of the model vis-a-vis the real quantities in the

[3]By contrast, models that emphasize labor-market frictions generate a pro-cyclical labor wedge, but fail to generate an Euler wedge. The same failure is shared by an important class of models featuring firm heterogeneity and financial frictions; see Buera and Moll (2014).

[4]Note that some of these shocks are often used in the literature as proxies of financial shocks.

data. Furthermore, the model predicts that the observed recessions would be *drastically* different, perhaps even turn to booms, if prices were flexible. We find the second prediction to be unappealing. We show that both predictions are overturned once confidence shocks are introduced in the model: nominal rigidity is no more essential for fitting the real data, the structural interpretation of the observed recessions is not so heavily dependent on the presumed degree of nominal rigidity.

We interpret these findings, in conjunction with our earlier results, as follows. The data favor models that give rise to "aggregate demand fluctuations", by which we mean the kind of empirical regularities documented in Section 2 of our paper and in prior work such as Blanchard and Quah (1989). In the pertinent NK literature, the assumption of significant nominal rigidity, along with additional modeling ingredients such as large investment-adjustment costs, has proved to be *essential* for obtaining fluctuations of this short. Furthermore, these modeling ingredients induce some key structural shocks—such as investment-specific, discount-factor, and news or noise shocks—to work drastically different than in the RBC framework. By contrast, variation in confidence, modeled as higher-order uncertainty, provides a micro-foundation of "aggregate demand fluctuations" that is at once more parsimonious, robust across RBC and NK settings, and quantitatively superior.

Related literature. By enriching the structure of beliefs in macroeconomic models, our paper adds to a literature that seeks to relax the "straight jacket" of the rational-expectations equilibrium concept. This includes work on robustness (Hansen and Sargent, 2007), learning (Evans and Honkapohja, 2001), and ambiguity (Ilut and Schneider, 2013); see also Woodford (2013) and the references therein. What distinguishes our contribution is the focus on higher-order uncertainty, the particular quantitative exercises we conduct, and the particular lessons we deliver.

Our paper adds to the growing literature on informational frictions (Mankiw and Reis, 2001, Sims, 2001, Woodford, 2002, Mackowiac and Wiederholt, 2009), but differentiates from it three ways. First, we propose a modeling trick that bypasses the technical difficulties that have hindered quantitative progress in this literature. Second, whereas some of this literature, especially that on rational inattention, focuses on payoff uncertainty, we isolate the role of higher-order (or strategic) uncertainty. Finally, we quantify the importance of this kind of uncertainty—and in that sense of "confidence"—in explaining the observed business cycles.

Our paper also connects to the literature that uses heterogeneous priors to explain asset-price phenomena such as speculation, bubbles, and leverage (Harrison and Kreps, 1992; Allen, Morris, and Postelwaite, 1993; Scheinkman and Xiong, 2003; Geanakoplos, 2010). One obvious difference is that we instead focus on business-cycle phenomena. Another difference is the particular modeling strategy we develop in order to engineer rich yet tractable higher-order belief dynamics. Whether this strategy can find a fruitful application in finance is an open question for future research.

Above all, our paper builds heavily on insights from game theory. The ideas of using heterogeneous priors as a tool to engineer rich higher-order uncertainty and of using the latter, in turn, as a tool to relax coordination have a long history in game theory. Our contribution rests on developing a particular formulation of these ideas that can be embedded in macroeconomic models without any loss of tractability and, most importantly, on confronting the theory with the data.

2 An Empirical Backdrop

In this section we identify some empirical regularities against which we confront, first, existing formalizations of the business cycle and, later on, our own approach.

We frame our investigation of the data by posing the following question. Suppose the macroeconomist seeks to capture the bulk of the business-cycle variation in the data as the response of a model economy to a *single* shock. What should the observable properties of this shock be?

To answer this question, we adopt the following strategy.[5] We run a VAR on eight key macroeconomic variables over the 1960-2007 period: GDP, consumption, hours, investment, the relative price of investment, inflation, the FFR, and government spending. (Here and throughout the paper, we leave out the post-2007 data, because we wish to abstract from the kind of financial phenomena that appear to have played a distinct role in the recent recession.) We assume there are two permanent components, those that drive the long-run movements in labor productivity and the price of investment. We filter out them from the data, thus obtaining six residual transitory components. We finally identify the wanted shock by taking the linear combination of these six components that maximizes the volatility of hours and investment at frequencies of 6 to 32 quarters.[6]

The rational behind this empirical strategy is as follows. We run a "big" VAR in order to extract information from multiple series. We filter out the aforementioned permanent components because they have familiar counterparts in models, as neutral and investment-specific technology shocks; because prior work (Blanchard and Quah, 1989) and our own results (see below) indicate that most business cycles are driven by transitory forces; and finally because we simply wish to focus on such transitory forces. At this point, one can envision two modeling strategies: consider models in which the business cycle is explained by the combination of eight structural shocks (as many as the residuals in the VAR); or consider models in which a single structural shock drives the bulk of the transitory movements in output, hours, and investment. Our empirical strategy seeks to guide the latter modeling strategy: the shock we identify is meant to be an empirical counterpart for the "main shock" in *any* model that aspires to offer a parsimonious theory of the business cycle.

Putting aside this rational, what we do is merely to construct a particular filter of the data. By design, this filter helps detect whether most of the variation in hours and investment at the business-cycle frequency—which is what our identified shock picks up—comes together with systematic co-movement in other variables, whether at business-cycle or other frequencies.

The solid lines in Figure 1 report the estimated IRFs to the aforementioned shock. The top two rows in Table 1 report the variance contribution of this shock along with the joint contribution of all the transitory components. Inspecting these results reveals the following properties.

First, and as anticipated, the transitory components explain the bulk (more than 75%) of the business-cycle volatility in output, hours, and investment.

[5] See the Appendix for a detailed description of the empirical strategy and the data, and our companion paper (Angeletos, Collard, and Dellas, 2014) for various sensitivity checks and other complementary evidence.

[6] We obtain similar results if we identify the wanted shock by maximizing the business-cycle volatility of output.

Second, our identified shock captures more than two-thirds of the joint contribution of all the transitory components in output, hours, and investment. It also triggers a visible business cycle, with consumption and the aforementioned variables all moving in tandem.

Third, the shock is associated with modest to insignificant movements in labor productivity and the price of investment; with modest pro-cyclical movements in inflation; and with pro-cyclical movements in both the nominal and the real interest rate.

Finally, we obtain similar results if we retain the permanent components of the data when identifying the wanted shock: see Table 2 and the dashed lines in Figure 1. This indicates that the "main business-cycle shock" is transitory by nature of the data, not by assumption. Importantly, although the shock is now allowed to encapsulate both transitory and permanent movements in labor productivity and the price of investment, it ends up explaining only a relatively small fraction of the volatility of these variables at *all* frequencies: see the second row in Table 2.

All in all, the picture that emerges appears to pose a challenge for existing structural interpretations of the observed macroeconomic fluctuations.

Clearly, our identified shock is unlike the kind of exogenous technology shocks, whether neutral or investment-specific, that play a prominent role in RBC and NK models alike. It is also inconsistent with models that emphasize either "news" or "noise shocks": in these models, the business-cycle movements in output, employment, and investment are correlated with future movements in labor productivity or the price of investment—a prediction contradicted by our evidence.[7]

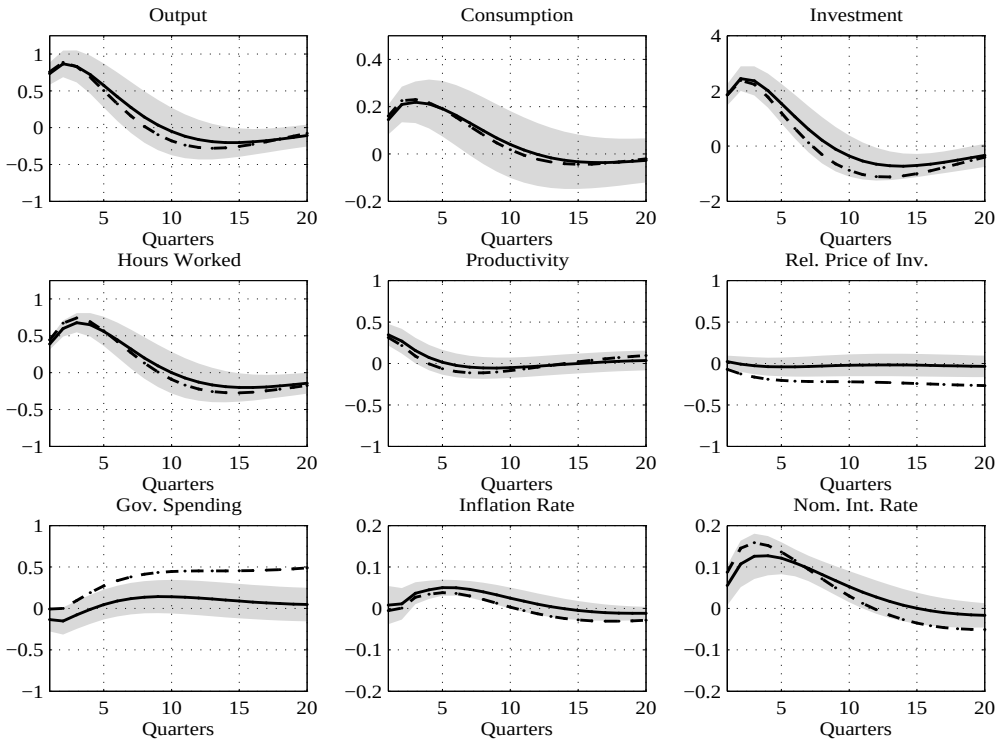
The evidence is also at odds with any model in which the business cycle is tied to *endogenous* movements in aggregate TFP and labor productivity. This is true whether these movements are themselves triggered by uncertainty shocks (Bloom et al, 2012), financial shocks (Kahn and Thomas, 2012, Buera and Mall, 2014), or even sunspots (Benhabib and Farmer, 1992).

Following Blanchard and Quah (1989), one may wish to interpret our identified shock as some kind of "aggregate demand shock". However, existing formalizations of this notion in the NK framework do not fit the picture that well. First, monetary shocks are associated with counter-cyclical movements in interest rates, which is the opposite of what we see in Figure 1. Second, the quantitative importance of either monetary or monetary shocks in estimated models is much smaller than that of our identified shock. Finally, discount-rate preference shocks and other micro-founded shocks that have a "demand flavor" have difficulty in generating the kind of co-movement between consumption and investment seen in Figure 1 (see Section ?? for more on this claim).

Like any other evidence, the evidence presented above does not pinpoint to a unique "right" theory. It could also be consistent with models that feature a rich combination of structural shocks. Notwithstanding these qualifications, in the sequel we show that variation in confidence—modeled as variation in higher-order beliefs—offers a parsimonious explanation of this evidence, along with other salient features of the data.

[7] This finding contradicts Beaudry and Portier (2005), but is consistent with Barsky and Sims (2010) and Barsky, Basu and Lee (2014). See our companion paper (Angeletos, Collard, and Dellas, 2014) for more on this topic.

Figure 1: IRFs to "Main Business-Cycle Shock"



Note: ——— shock that maximizes volatility in hours and investment at business-cycle frequencies, after filtering out the two permanent components; ▬ 68% HPDI; - - - - the aforementioned shock without filtering out the permanent components.

Table 1: Contribution of "Main Shock" at Business Cycle Frequencies (6–32 quarters)

	Y	C	I	h	Y/h	P_i	G	π	R
All Transitory	78.47	64.36	79.64	75.99	79.36	61.25	65.89	68.66	73.51
	[68.91,88.11]	[50.22,78.44]	[70.39,88.91]	[65.75,86.41]	[69.11,89.76]	[42.57,80.28]	[51.45,80.32]	[55.41,81.68]	[59.14,87.80]
Main Transitory	51.69	25.55	55.78	49.36	21.42	6.07	9.89	18.50	29.94
	[42.74,60.69]	[15.11,35.97]	[46.81,64.78]	[40.66,58.08]	[12.24,30.48]	[1.60,10.69]	[3.93,15.91]	[9.76,26.63]	[15.25,44.94]

Note: *All Transitory* denotes the contribution of all transitory shocks to the dynamics of the variables at the business cycle frequencies (6–32 quarters). *Main Shock* denotes the contribution of our identified shock, namely of the linear combination of the transitory components that explain the maximal share of the volatility of investment and hours at business cycle frequencies. 68% HPDI into square brackets.

Table 2: Contribution of Variant "Main Shock" with Permanent Components Included

	Y	C	I	h	Y/h	P_i	G	π	R
6–32 quarters	61.32	32.23	64.59	63.85	24.62	12.43	13.69	20.26	46.58
	[52.93,69.60]	[21.72,42.77]	[56.53,72.57]	[55.97,71.74]	[14.84,34.58]	[3.86,21.08]	[5.27,22.45]	[9.74,30.52]	[26.52,65.87]
32–80 quarters	27.63	13.97	33.89	33.21	12.20	12.83	11.34	18.50	47.03
	[14.80,40.57]	[4.60,23.51]	[20.91,47.14]	[20.78,45.67]	[2.53,22.65]	[2.50,23.75]	[2.83,20.09]	[7.26,29.89]	[31.66,61.88]
80–∞ quarters	10.35	10.23	10.29	10.67	10.35	10.23	11.08	9.74	12.59
	[1.24,20.21]	[1.16,20.02]	[1.12,20.39]	[2.51,19.01]	[1.24,20.21]	[1.17,20.06]	[1.23,21.90]	[2.05,17.49]	[4.03,21.42]

Note: The shock is again identified by maximizing the business-cycle volatility of hours and investment, but now without filtering out the two permanent components of the data. 68% HPDI into square brackets.

3 An RBC Model with Fluctuations in Confidence

In this section we set up our baseline framework: a variant of the RBC model, augmented with a tractable form of higher-order belief dynamics. We first describe the physical environment, which is standard. We then specify the belief structure, which is where the only novelty rests.

Geography, markets, and timing. There is a continuum of islands, indexed by i, and a mainland. On each of the islands, there is a firm that produces a differentiated good, and a household that consists of a consumer, a worker, and a shopper. Workers can work only in the local firm and, similarly, capital is specific to the island. On the mainland, there is a final-good sector where the intermediate goods are combined to produce the final good, which in turn is used for consumption and investment. The final-good producers are competitive, whereas the intermediate-good producers are monopolists. Prices are flexible, ruling out nominal rigidities.

Time is indexed by $t \in \{0, 1, ...\}$ and each period contains two stages. In stage 1, the local labor and capital markets operate on each island, the local worker decides how much to work, and the local intermediate-good firm decides how much labor to hire and how much capital to utilize. Importantly, these choices are made under imperfect knowledge of the concurrent level of economic activity on other islands. In stage 2, the intermediate-good producer and the shopper of each island visit the mainland. At this point, the final-good sector purchases the intermediate goods, producse the final good, and sells it to the shoppers, which in turn take it back to their home island so that it can be consumed and/or invested in the local firm.[8]

Households. Consider household on island i. Its 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} denotes consumption, n_{it} denotes employment (hours worked), and U is the per-period utility function. The latter takes the form

$$U(c, n) = \log c - \frac{n^{1+\nu}}{1+\nu}$$

where $\nu \geq 0$ is the inverse of the Frisch elasticity of labor supply. The household's period-by-period budget constraint is given by

$$c_{it} + i_{it} = w_{it} n_{it} + r_{it} k_{it} + \pi_{it}$$

where 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 given by

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

where δ is the depreciation rate.

[8] We abstract from centralized financial markets; such markets can easily be incorporated in stage 2, when information has been revealed.

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

$$y_{it} = A_t(n_{it})^{1-\theta}(k_{it}u_{it})^{\theta}$$

where A_t is the aggregate TFP level, k_{it} is the local capital stock, and u_{it} is capital utilization. Its profit is given by

$$\pi_{it} = p_{it}y_{it} - w_{it}n_{it} - r_{it}k_{it} - \Psi(u_{it})k_{it}.$$

where p_{it} is the price of the good and $\Psi(u_{it})k_{it} = \psi_0 u_{it}^{\frac{1}{1-\psi}} k_{it}$ is the resource cost of utilization, with $\psi_0, \psi > 0$ and $\psi \in (0,1)$. The firm is a price-taker in the local labor and capital markets, but a monopolistic competitor in the mainland commodity market.

Final-good retailers. The final-good sector is competitive. Its output is given by the following CES aggregator of the intermediate goods:

$$Y_t = \left(\int_0^1 y_{it}^{1-\varrho}\right)^{\frac{1}{1-\varrho}} \text{ with } \rho \in (0,1).$$

Profits (which are zero in equilibrium because of CRS) are given by $Y_t - \int p_{it}y_{it}di$. It follows that the demand faced by the intermediate-good firm of island i takes the following familiar form:

$$p_{it} = Y_t^{\varrho} y_{it}^{-\varrho}. \tag{1}$$

(Note that we have normalized the aggregate price level to 1.)

Technology shocks. TFP follows a random walk with drift: $\log A_t = \gamma + \log A_{t-1} + v_t$, where $\gamma > 0$ pins down the long-run growth of the economy and v_t is the period t innovation. The latter is drawn from a Normal distribution with mean 0 and variance σ_A^2.

Higher-order uncertainty and confidence shocks. In the standard RBC model, the realization of A_t is assumed to be common knowledge, ruling out higher-order uncertainty. Under the standard rational-expectations equilibrium concept, this also rules out strategic uncertainty and ultimately imposes that the only source of volatility in equilibrium beliefs and outcomes is the technology shock. We instead open the door to higher-order uncertainty by relaxing the assumption that A_t is commonly known. We then engineer rich yet tractable dynamics in higher-order beliefs by adopting a particular heterogeneous-prior specification of the information structure.

First, we let each island i observe the following private (local) signal:

$$x_{it} = \log A_t + \varepsilon_{it},$$

where ε_{it} is an island-specific noise term. Next, we assume that each island believes that its own signal is unbiased, whereas the signal of others is biased, and let ξ_t denote the bias that each island perceives to exist in the information of other islands. Formally, for every i, the prior of island i is (i) that the noise ε_{it} in its own signal is drawn from a Normal distribution with mean 0 and variance σ^2 and (ii) that the average noise in the rest of the economy is drawn from a Normal distribution with the same variance and mean equal to ξ_t. The latter is a random variable, independent of A_t,

and commonly observed in the beginning of period t. Finally, we assume that the actual signals are unbiased (so that agents are wrong to think that others are biased) and take the limit as $\sigma \to 0$.[9]

These modeling choices are admittedly exotic, yet they strike a delicate balance between the desire for tractability and the desire for accommodating rich higher-order belief dynamics. By taking the limit as $\sigma \to 0$, we force all agents to behave as if they were perfectly informed and to take the same action in equilibrium. We thus bypass the technical challenges that have hindered quantitative work in models with informational frictions and guarantee a minimal increase in the state space of the model. At the same time, by introducing heterogeneous priors, we make sure that higher-order beliefs do not collapse to first-order beliefs *even* in that limit. Finally, we control higher-order beliefs by specifying an appropriate stochastic process for ξ_t.

Let us expand on the modeling role of ξ_t. When a positive innovation in ξ_t occurs, neither the true fundamentals nor any agent's perception of them changes. Yet, every agent believes that the agents in all other islands are optimistic about the productivity, and the beliefs, of others. Conversely, a negative innovation in ξ_t triggers a negative movement in higher-order beliefs. At this point, it is not obvious how exactly these higher-order belief movements impact on equilibrium outcomes; to figure this out, we need to solve for the equilibrium strategies, which we do in the next section. Nonetheless, we can safely guess that, one way or another, the ξ_t shock will cause equilibrium beliefs and outcomes to fluctuate away from and around the outcomes predicted by the standard, common-knowledge, RBC benchmark. From this perspective, variation in ξ_t operationalizes the notion that coordination is imperfect: the resulting fluctuations can be interpreted as the symptom of "coordination mistakes". As it turns out, these fluctuations are also associated with co-movemement in the equilibrium expectations that firms form about consumer demand and profitability, on the one hand, and consumers form about employment and income, on the other hand. With this in mind, we henceforth interpret the variation in ξ_t as variation in "confidence".

In our framework, these fluctuations are hard-wired in biased priors. However, Angeletos and La'O (2013) obtain similar fluctuations in a common-prior, rational-expectations setting by allowing for sufficiently rich informational frictions. Furthermore, following Bergemman and Morris (2013) and Bergemman, Heumann, and Morris (2013), one can show that this kind of fluctuations is *generic* to the incompleteness of information. Finally, if we were to focus on static settings, which is what these papers did, we could establish an exact observational equivalence between the higher-order belief variation we engineer here and the one obtained in those papers under the common-prior restriction. This is because the empirical content of higher-order uncertainty does not rest on the precise details of the underlying information structure. With these points in mind, we invite the reader to interpret our heterogeneous-prior formulation as a convenient proxy for informational frictions as opposed to taking it too literally.

This interpretation also implies that ξ_t oughts to be transitory. If the agents share the same

[9] In practice, we simply set $\sigma = 0$. In principle, however, we have in mind a sequence of incomplete-information economies in which the noise in first-order beliefs is vanishing, yet the variation in higher-order beliefs remains bounded away from zero.

prior and keep learning (or communicating) over time, their beliefs are likely to converge, at least insofar as there are no exogenous innovations in their information. This means that the belief fluctuations we seek to capture in this paper are inherently short-lived. To capture this property in our setting, we assume that ξ_t follows an $AR(1)$ process:

$$\xi_t = \varrho \xi_t + \zeta_t,$$

with $\varrho \in (0,1)$. One can then think of ϱ as a proxy for the speed in which beliefs converge to their common-knowledge counterparts, and of ζ_t as the exogenous innovations in information.[10]

4 Equilibrium characterization and solution method

To solve the model, we proceed in three steps. First, we introduce a fictitious local planner, whose choices replicate the equilibrium allocations of an island for given aggregate outcomes. Next, we study the fixed point between the policy rules of the local planners of all the islands, which gives us the general equilibrium of the economy. We conclude by obtaining the log-linear approximation of the equilibrium—and by discussing how our method extends to a large class of DSGE models.

A fictitious planner

To simplify the exposition, we momentarily abstract from utilization. Because the local labor and capital markets are competitive, the local equilibrium outcomes of any given island coincide with the solution to the following planning problem:

$$\max \mathbb{E}_i \left[\sum_{t=0}^{\infty} \beta^t \left(\log c_{it} - \frac{n_{it}^{1+\nu}}{1+\nu} \right) \right]$$

subject to

$$c_{it} + i_{it} = Y_t^{\varrho} y_{it}^{1-\varrho} \tag{2}$$

$$k_{i,t+1} = (1-\delta) k_{it} + i_{it} \tag{3}$$

$$y_{it} = A_t (n_{it})^{1-\theta} (k_{it})^{\theta} \tag{4}$$

[10] Apart from guaranteeing that our belief fluctuations are transitory, the AR(1) specification is in line with conventional specifications of shock processes in the DSGE literature, which is useful for comparison purposes. Note, however, that the belief fluctuations we are after may naturally feature hump-shaped dynamics, because of contagion effects; see Angeletos and La'O (2013) for a concrete illustration. Hump-shaped belief dynamics also obtain in epidemiological models of belief formation, such as Carroll (2001) and Burnside, Eichenbaum, and Rebelo (2011). To capture this possibility, we have experimented with an $AR(2)$ specification for ξ_t, and have found that this actually improves the empirical performance of our models. Finally, note that our approach presumes that the confidence shock is correlated across agents. As noted by Angeletos and La'O (2013), such correlation may emerge endogenously from communication: when agents talk to one another, learn from one another's actions, or observe the same economic statistics, their beliefs become correlated.

The above conditions are familiar: condition (2) is the local resource constraint, which is obtained by combining the household's budget with the firm's profit and replacing p_{it} from (1); condition (3) is the law of motion for local capital; and condition (4) is the local production function. It is only worth noting that the resources that are available in any given period for consumption and investment depend, not only on the island's own output, but also on the aggregate output of the economy. This captures the aggregate demand externality that is present in modern DSGE models and identifies the precise channel through which higher-order beliefs matter in our setting: ξ_t affects equilibrium outcomes because, and only because, it triggers extrinsic variation in expectations of "aggregate demand".

Recursive equilibrium

Recall that we are interested in the limit as $\sigma \to 0$. In this limit, every island's posterior belief of A_t during stage 1 converges to a Dirac distribution centered around the island's private signal, which in turn converges in distribution to the true shock A_t. Furthermore, each island recognizes that its *own* signal converges in distribution to A_t, but believes that the signals of *other* islands converge to $A_t + \xi_t$. By the same token, the first-order beliefs of all islands collapse to A_t, but their higher-order beliefs remain displaced away from a A_t by a factor that is proportional to ξ_t. These observations suggest that, relative to the common-knowledge RBC benchmark, the state space of the model has been extended only by the introduction of one additional state variable, namely the ξ_t shock. We thus propose that the equilibrium dynamics of aggregate capital satisfy

$$K_{t+1} = \Gamma(A_t, \xi_t, K_t),$$

for some function Γ, and that the equilibrium levels of aggregate output, consumption, investment, and hours can similarly be expressed as functions of (A_t, ξ_t, K_t).

When characterizing the equilibrium, we nevertheless need to keep track of the difference between the realized level of aggregate output and the belief that the agents hold about it during stage 1 of each period. To this goal, let us denote with $\bar{x}_t$ the average signal in the population and let us express aggregate aggregate output as

$$Y_t = \mathcal{Y}(A_t, \bar{x}_t, \xi_t, K_t),$$

for some function $\mathcal{Y}$. This notation distinguishes the direct effect that TFP has on aggregate output for given capital and labor, from the effect that the average signal of TFP has on the labor decisions of the islands. Of course, the actual realization of $\bar{x}_t$ coincides with A_t, which means that the *actual* level of aggregate output is ultimately given by

$$Y_t = \mathcal{Y}(A_t, A_t, \xi_t, K_t).$$

Yet, when the islands make their employment and production choices during state 1 of period t, they believe that the average signal $\bar{x}_t$ in the rest of the population is biased away from the actual

TFP. In particular, when an island receives a signal x_t, it believes with probability one that $A_t = x_t$ and that $\bar{x}_t = x_t + \xi_t$, which means that its *perceived* level of aggregate output is given by

$$\hat{Y}_t = \mathcal{Y}(x_t, x_t + \xi_t, \xi_t, K_t).$$

With these points in mind, we now proceed to develop a recursive representation of the planning problem that each island faces. Let V_1 and V_2 denote the value functions corresponding to stages 1 and 2 of each period. Consider now stage 1, at which point any given island must choose its employment and production before observing the level of economic activity in other islands and before trading selling its intermediate good in the mainland. For any given values (k, x) of the local capital stock and the local signal, the island solves the following problem:

$$\begin{aligned} V_1(k, x; \xi, K) \;=\; & \max_n \quad V_2(\hat{m}; x, \xi, K) - \tfrac{1}{1+\epsilon} n^{1+\epsilon} \\ & s.t. \quad \hat{m} = \hat{Y}^{\varrho} \hat{y}^{1-\varrho} + (1-\delta)k \\ & \qquad \hat{y} = x n^{\theta} k^{1-\theta} \\ & \qquad \hat{Y} = \mathcal{Y}(x, x+\xi, \xi, K) \end{aligned} \tag{5}$$

In the above, $\hat{y}$ and $\hat{Y}$ stand for the island's perceptions of the local and aggregate levels of output, while $\hat{m}$ stands for the resources (the quantity of the final good) that the island expects to collect in stage 2 from the sale of its intermediate good. Consider next stage 2, at which point any given island makes its consumption and investment choices after having collected its share of the final good (and having also observed the realized values of all the relevant economic variables). For any given value m of resources, the island solves the following problem:

$$\begin{aligned} V_2(m; A, \xi, K) \;=\; & \max_{(c,k')} \quad U(c) + \beta \int V_1(k'; A', \xi', K') df(A', \xi' | A, \xi) \\ & s.t. \quad c + k' = m \\ & \qquad K' = \Gamma(A, \xi, K) \end{aligned} \tag{6}$$

where f stands for c.d.f. of the (exogenous) joint distribution of the technology and confidence shocks.

In solving the above two problems, each island takes the functions $\mathcal{Y}$ and Γ as given: these functions summarize the beliefs the island holds regarding the strategies of other island, or equivalently the statistical model it has in its mind regarding the evolution of aggregate economic outcomes. In equilibrium, these functions must in turn be consistent with the policy rules that obtain from the solution of these two problems. Thus let us denote with $n(k, x; \xi, K)$ the policy rule for employment, which obtains from the stage-1 problem, and with $g(m; A, \xi, K)$ the policy rule for capital, which obtains from the stage-2 problem. When all islands follow these rules, aggregate output is given $Y_t = A_t n(K_t, \bar{x}_t; \xi_t, K_t)$ and aggregate capital evolves according to $K_{t+1} = g(M_t; A_t, \xi_t, K_t)$, where $M_t = Y_t + (1-\delta)K_t$. It follows that, in equilibrium, the functions $\mathcal{Y}$ and Γ must satisfy

$$\mathcal{Y}(A, \bar{x}, \xi, K) = A n(K, \bar{x}; \xi, K), \tag{7}$$

for all $(A, \bar{x}, \xi, K)$, and

$$\Gamma(A, \xi, K) = g\left(\mathcal{Y}(A, A, \xi, K) + (1-\delta)KlA_t, \xi_t, K_t\right), \tag{8}$$

for all (A, ξ, K).

Putting the above points together, we conclude that an equilibrium is a fixed point between the Bellman problems in (5)-(6) and the aggregation conditions (7)-(8).

Definition 1 *An equilibrium is given by a combination of individual and aggregate policy rules,* (n, g) *and* $(\mathcal{Y}, \Gamma)$*, such that: (i) the individual rules* (n, g) *solve the problems in (5) and (6) taking the aggregate rules as given; and (ii) the aggregate rules* $(\mathcal{Y}, \Gamma)$ *are obtained from the individual ones according to (7) and (8).*

Log-linear solution

The basic idea behind the equilibrium characterization in the previous section was to look for a fixed point between the perceived and the actual laws of motion. In principle, one can use this approach to obtain the non-linear solution of the model numerically. In Appendix A, however, we show how the same basic idea delivers a tractable, closed-form solution once implemented on the log-linear approximation of the equilibrium. In particular, we show that the log-linear policy rules of our model take the following form:

$$\begin{aligned} (\tilde{Y}_t, \tilde{C}_t, \tilde{I}_t, \tilde{N}_t) &= \Gamma_1 \tilde{K}_t + \Gamma_2 \tilde{A}_t + \Delta_1 \xi_t \\ \tilde{K}_{t+1} &= \Gamma_3 \tilde{K}_t + \Gamma_4 \tilde{A}_t + \Delta_2 \xi_t \end{aligned}$$

where the tilde ($\sim$) over a variable indicates the log-deviation of that variable from its steady-state value. The same solution applies to the standard RBC model, except that in that case $\Delta_1 = \Delta_2 = 0$. Therefore, if we put aside the underlying economics and focus on the technical aspects, all we have done at the end the day is to introduce a single additional state variable, namely ξ_t. This state variable summarizes the role of higher-order beliefs, and is conveniently modeled here as an exogenous stochastic process.

In Appendix A, we explain how these properties generalize to a large class of DSGE models. We first show how a generic DSGE model can be augmented with the type of higher-order belief dynamics we are proposing in this paper. We then develop a simple algorithm that obtains the solution of the belief-augmented model as a simple transformation of the solution of the original model. This facilitates the simulation, calibration, and estimation of the belief-augmented model, as if it were a standard DSGE model. The results we develop in Appendix A are therefore instrumental to the broader methodological contribution of our paper: researchers can readily use these results to embed higher-order belief dynamics in their preferred model, so as to study how imperfect coordination and higher-order uncertainty matter for the questions of their interest.

5 Quantitative evaluation

This section explores the qualitative and quantitative properties of our baseline model. We fix the preference and technology parameters to conventional values; we next calibrate the two shocks to match the volatility of output and hours; we then explore the fit of the model vis-a-vis a number of salient features of the data. This strategy helps reveal in what dimensions, and by how much, variation in "confidence" improves the empirical performance of the prototype RBC model.

Calibration

Table 3 reports the (conventional) values we pick for the preference and technology parameters of model. The discount factor, β, is set to 0.99, meaning an annual discount rate of 4%. The Firsch elasticity of labor supply, $1/\nu$, is set to 2. The capital share in production, α, is set to 0.3. The elasticity of the utilization resource cost, $1/\psi$, is set to 0.2; this is in line with values traditionally considered in the RBC and DSGE literatures (e.g., King and Rebelo, 1999). Finally, the parameter ρ, which governs the degree of strategic complementarity across islands, is set to 0.80.

Table 3: Parameters

Parameter	Role	Value
β	Discount Rate	0.99
ν	Inverse Elasticity of Labor Supply	0.50
α	Capital Share in Production	0.30
ψ	Inverse Elastiticity of Utilization Cost	0.50
ρ	Complementarities	0.80
σ_a	St.Dev. of Technology Shock	0.60
σ_ξ	St.Dev. of Confidence Shock	0.75
ϱ	Persistence of Confidence Shock	0.89

Since the technology shock follows a random walk, three parameters remain to be set: the standard deviations of the two shocks and the persistence of the confidence shock. For the latter, we target a half-life of 6 quarters, which gives $\varrho = 0.89$. This choice is motivated by two considerations: first, that higher-order uncertainty can only be moderately persistent, because learning guarantees convergence of beliefs over time; and second, that the empirical shock documented in Section 2 has as a half-life of a similar magnitude. Finally, as already mentioned, the standard deviations of the two shocks are set so as to match the standard deviations of output and employment. This gives $\sigma_a = 0.60$ for the technology shock and $\sigma_\xi = 0.75$ for the confidence shock.[11]

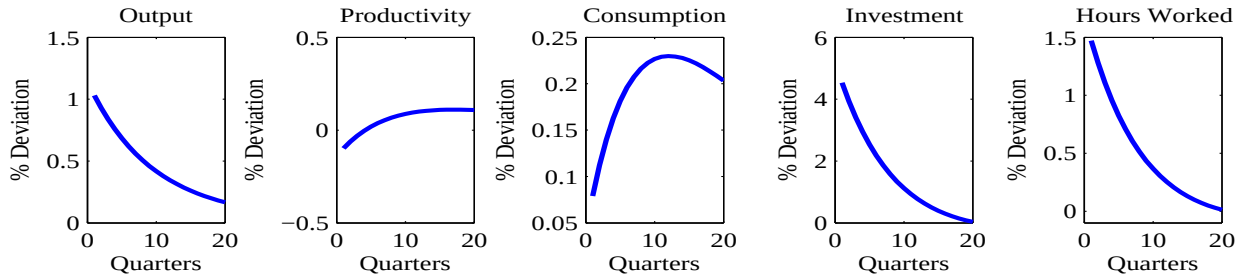
[11] These values for σ_a and σ_ξ mean that the higher-order uncertainty we assume is rather modest compare to the underlying payoff uncertainty. See footnote ?? for a detailed explanation.

Impulse Responses

To understand better the nature of our confidence-driven fluctuations, we start by studying the impulse response functions (IRFs) to the confidence shock. The impulse responses to the technology shock are the same as in the standard RBC model and are thus omitted.

Figure 2 shows that a positive innovation in ξ_t causes a transitory boom in output, consumption, hours, and investment. At the same time, labor productivity stays nearly constant, falling a bit in the beginning and increasing a bit later on. At least qualitatively, the confidence shock therefore looks like the empirical shock that we documented in Section 2 and that has traditionally been interpreted as an "aggregate demand shock". We elaborate on the quantitative content of this finding, and on its position in the literature, in the sequel. But before we do that, we want to explain the economic forces that operate behind these IRFs.

Figure 2: Impulse responses to a positive confidence shock



In general, the effects of higher-order uncertainty on incentives and equilibrium outcomes hinge on the details of the model under consideration. In our model, higher-order beliefs matter *only* through expectations of aggregate capital and aggregate output.[12] In particular, a positive innovation in ξ_t causes, in equilibrium, any given island to update upwards its expectations of the output, consumption, and investment in all other islands, relative to what these expectations would have been under perfect coordination (i.e., in the absence of higher-order variation, or, equivalently, in the standard RBC model). It is in this sense that higher-order beliefs in our model are tied to expectations of "aggregate demand".

Furthermore, because capital adjusts and in addition the shock to higher-order beliefs is itself moderately persistent, aggregate demand is expected to boom, not only in the current period, but also for the next few quarters. As a result, the firms in each island find it optimal to raise their demand for both labor and capital, which in turn tends to raise both wages and capital returns, for few quarters. Other things equal, this motivates the household to supply more labor as well as to save more, because of the familiar static and intertemporal substitution effects. A countervailing wealth effect is also at work, because the households expect (and experience) a

[12] This should be evident from the recursive representation we developed in the previous section: the problem faced by the local planner of an island depends on the strategies of other islands only through the policy rules for aggregate capital and aggregate output.

boom in their income. However, because the boom is transitory, the wealth effect is weak, while the intertemporal substitution effect is strong. This guarantees that households find it optimal to increase their labor supply, as well as to split the increase in their income between consumption and saving. Finally, because TFP is fixed and production exhibits diminishing returns in labor, the increase in employment contributes towards lower labor productivity in the beginning, while the accumulation of capital helps increase it later on; and because utilization is variable, the labor-productivity effects are modest.

Together, the above arguments explain why confidence shocks cause the equilibrium levels of employment, investment, output, and consumption to move in tandem, along with no variation in the underlying TFP or other payoff-relevant fundamentals and with only small variation in measured labor productivity. It follows that the confidence shock in the model mimics the empirical shock we documented in Section 2.

The finding that confidence shocks trigger co-movement in employment, investment, output, and consumption stands in sharp contrast to many other structural shocks considered in the literature. Consider, in particular, the recent work on news and noise shocks. Because these shocks affect expectations of future permanent movements in TFP, they entail strong wealth effects and delayed substitution effects. As a result, these shocks fail to generate the right co-movement with the standard RBC model: positive news or noise shocks raise consumption but reduce hours, output, and investment. Moving to the NK framework can correct this failure, but only insofar as one combines sticky prices with significant adjustment costs to investment and a sufficiently accommodative monetary policy, by which we mean a policy that expands the economy relative to the underlying-flexible price allocation. A similar point applies to investment-specific shocks, discount-factor shocks, and certain financial shocks: these shocks fail to generate the right co-movement in baseline calibrations of either the RBC or the NK framework. (For a detailed explanation of these claims, we refer the reader to the Appendix.) We conclude that the ability of confidence shocks to generate the "right" macroeconomic co-movement is neither trivial nor easily matched by other structural shocks considered in the literature.

Business-Cycle Moments

Now that we have studied the basic mechanism through which higher-order uncertainty operates in the model, we start the evaluation of its quantitative performance by looking at its ability to match standard business-cycle moments. Table 4 reports various volatilities and correlations implied by the model, and compares them to those found in the data. All moments are computed on Bandpass-filtered series.[13] For comparison purposes, we also study two variant specifications of the model. In the one, we shut down the confidence shock, in which case our model reduces to the standard RBC

[13]Throughout the paper, we use the Bandpass filter on frequencies 6-32 quarters. This is considered optimal in the empirical business-cycle literature, because it removes, not only low-frequency movements, but also high-frequency noise. We nevertheless obtain nearly identical results if we use the HP filter.

model. In the other, we shut down the technology shock, which helps isolate the role of confidence shocks. Together, these two variants help identify the distinct ways through each of the two shocks contributes towards the overall empirical performance of the model.

As evident from the first two columns of Table 4, the model does a good job in matching the relevant volatilities and correlations. Furthermore, this success owes to a delicate balance between the contribution of the two shocks: as shown in the last two columns of the table, the specifications that shut down either shock fail in different dimensions, but the combination of the shocks gets multiple things right at the same time.

Table 4: Band–passed Filtered Moments

	Data	Baseline	Technology	Confidence
σ_y	1.46	1.46	1.46	1.46
σ_h/σ_y	1.17	1.17	0.32	1.43
σ_c/σ_y	0.59	0.43	0.70	0.15
σ_i/σ_y	3.95	3.77	2.11	4.41
$\sigma_{y/h}/\sigma_y$	0.68	0.52	0.86	0.12
σ_{sr}/σ_y	0.86	0.63	0.95	0.34
$\rho_{c,y}$	0.86	0.68	1.00	0.54
$\rho_{i,y}$	0.94	0.95	1.00	0.99
$\rho_{h,y}$	0.88	0.89	0.99	0.99
$\rho_{y/h,y}$	0.52	0.46	1.00	-0.76
$\rho_{y/h,h}$	0.07	0.00	0.98	-0.82

In particular, consider the specification of the model that shuts down the confidence shock. In this case, once the standard deviation of the TFP shock is calibrated to match the volatility of output, the generates a very low volatility in hours; conversely, if the model were calibrated to match the volatility of hours, then the model would generate counterfactually levels of volatility in output, as well as all other variables. This reflects the well-known failure of the baseline RBC model to match the relative volatility of hours unless it assumes an implausibly high value for the Frisch elasticity of labor supply. The model also tends to generate too much volatility in consumption and too little volatility in investment. Most crucially, the model predicts a nearly perfect positive correlation between labor productivity and either output or hours, whereas the data feature only a modest correlation between labor productivity and output and a zero correlation between labor productivity and hours. This fact is closely connected to the fact we documented in Section 2, namely that the bulk of the business cycle movements in hours, as well as in investment, appear to be disconnected from movements in technology.

Now consider the specification of the model that shuts down the technology shock, thus isolating

the role of confidence-driven fluctuations. Once calibrated to match the volatility of output, the model predicts a counterfactually high volatility in hours. This is because because diminishing returns impose that output fluctuates less than employment for given capital, and capital in turn moves only very slowly. Furthermore, the model predicts a rather low volatility in consumption and a rather high volatility in investment. This is because the confidence shock is quite transitory. Finally, the model predicts a strongly negative correlation between labor productivity and either output or hours. As already discussed, this is because diminishing returns causes labor productivity to fall in response to an expansionary confidence shock.

Note then that, in all the moments that the technology shock fails to match in one way, the confidence shock happens to contribute in exactly the opposite way. This helps explains why the model that combines the two shocks has a better chance of matching all the moments at the same time. That the match turns out to be nearly perfect—especially with regard to the correlations between output, hours, investment and labor productivity—was a welcome surprise in our eyes.

This performance is not merely the result of adding a second structural shock to the baseline RBC model: other shocks that have been considered in the literature could not have delivered the same performance. We illustrate this point in Table 5 by running a horse-race between our model and five competing models that replace the confidence shock with one of the following, more conventional, structural shocks: (i) an investment-specific shock, which could be interpreted either as a technological shock or as a financial shock affecting the firms' shadow cost of investment; (ii) a discount-rate shock, which could be interpreted either as preference shock or as a financial shock affecting the households' shadow cost of consumption; (iii) a shock to government spending; (iv) a news shock, that is, a shock to expectations of future permanent TPF movements; and (v) a markup shock, that is, a shock to the monopoly power of the firms.

As evident in Table 5, the only competing model that replicates the empirical performance of our model is the one featuring the markup shock. The main reason behind the failure of the other four competing models is that the corresponding structural shocks are unable to generate positive co-movement between output, employment, consumption, and investment.[14] By contrast, the last model matches the empirical performance of our model because the markup shock and the confidence shock have similar effects on real allocations: a temporary reduction in monopoly power has a similar stimulating effect on input and output choices as optimistic beliefs about demand. The two shocks, however, tend to have different implications for inflation: a reduction in monopoly power induces firms to reduce prices, whereas optimistic beliefs about demand motivates them to raise prices.[15] The evidence presented in Section 2 therefore appears to favor confidence shocks over markup shocks—and so does an estimation we conduct later on.

[14]These properties are evident from the corresponding IRFs, which we report in the Appendix. One can also see there that the real implications of markup and confidence shocks are nearly, but not exactly, the same.

[15]A caveat with this argument is that it is based on partial-equilibrium reasoning: in general equilibrium, the response of inflation to *any* shock hinges on the response of monetary policy. The argument therefore carries over insofar as monetary policy is accommodative of (does not overturn) the underlying price pressures.

Table 5: A Horserace

	Data	Conf	Invt	News	Disc	Mkup
standard deviations						
σ_y	1.46	1.46	1.46	1.46	1.46	1.46
σ_h/σ_y	1.17	1.17	1.17	0.87	1.17	1.17
σ_c/σ_y	0.59	0.43	0.93	0.27	0.92	0.43
σ_i/σ_y	3.95	3.77	6.73	5.09	6.74	3.78
$\sigma_{y/h}/\sigma_y$	0.68	0.52	0.59	0.39	0.58	0.51
σ_{sr}/σ_y	0.86	0.63	0.61	0.66	0.60	0.63
correlations						
$\rho_{c,y}$	0.86	0.68	-0.37	-0.33	-0.38	0.68
$\rho_{i,y}$	0.94	0.95	0.88	0.98	0.89	0.95
$\rho_{h,y}$	0.88	0.89	0.86	0.99	0.87	0.89
$\rho_{y/h,y}$	0.52	0.46	0.28	0.88	0.27	0.46
$\rho_{y/h,h}$	0.07	0.00	-0.23	0.79	-0.23	0.01
$\rho_{c,h}$	0.87	0.28	-0.78	-0.48	-0.79	0.28
$\rho_{c,i}$	0.77	0.43	-0.76	-0.50	-0.76	0.43

Mistakes and Wedges

Higher-order uncertainty lets the equilibrium beliefs of the incomplete-information specification of a model to fluctuate away from, and around, the equilibrium beliefs of its complete-information counterpart. This is similar to introducing a random "wedge" to any condition of the complete-information model that features an expectation operator. In our model, expectations of aggregate output enter the optimality conditions for hours and savings. It follows that variation in confidence registers as variation in the labor and capital wedges.

Thus let τ_t^n and τ_t^k denote the labor and capital wedges in our model, defined by the following:

$$\begin{aligned} V'(N_t) &= \tau_t^n U'(C_t)(1-\alpha)\frac{Y_t}{N_t} \\ U'(C_t) &= \beta\tau_t^k \hat{E}_t\left[U'(C_{t+1})\left(1-\delta+\alpha\frac{Y_{t+1}}{K_{t+1}}\right)\right] \end{aligned}$$

where $\hat{E}_t$ denotes the objective expectations operator. In the appendix we show that, up to log-linearization, these wedges are increasing functions of the ξ_t shock alone: the incentive effects of a negative confidence shock are akin to these of a joint tax on labor and capital. Our model therefore predicts that the two wedges are pro-cyclical, as well as positively correlated with each other.

To appreciate the bite of these predictions, consider the alternative scenario in which the agents have completely irrational beliefs. In this case, their expectations of the returns to labor and capital

need not be consistent with each other, nor with their expectations of other economic outcomes, and therefore the wedges would not have to be positively correlated either with each other or with output. The same is true if we allow for arbitrary "trembles" in consumption, saving, and employment choices, unless these trembles happen to be tied together in a specific fashion—which is effectively what our model imposes. Finally, consider Benhabib and Farmer (1992), which seeks to accommodate "animal spirits" in the RBC context by introducing increasing returns and multiple equilibria. In that model, both wedges are zero: animal spirits drive business cycles only by driving the Solow residual. These points underscore that our model's predictions about the wedges are neither trivial nor shared by alternative formalizations of extrinsic uncertainty.

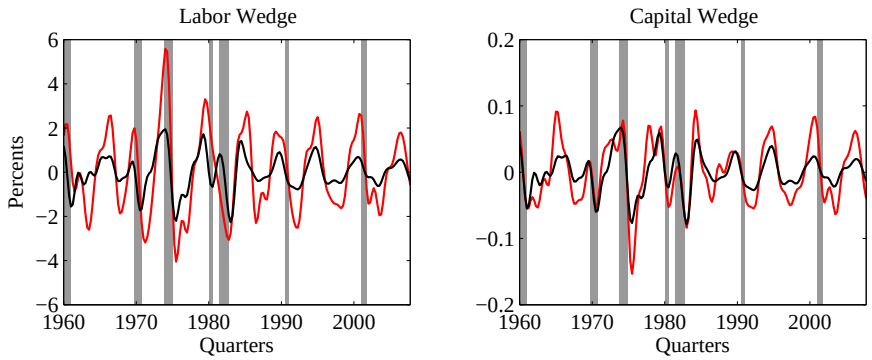
Table 6: Wedges

$Corr(y_t, \tau^h_{t+j})$	j= -2	-1	0	+1	+2
data	0.30	0.51	0.69	0.80	0.81
model	0.49	0.72	0.80	0.70	0.46
$Corr(y_t, \tau^k_{t+j})$	j= -2	-1	0	+1	+2
data	0.25	0.34	0.40	0.42	0.39
model	0.49	0.72	0.80	0.70	0.46
$Corr(\tau^h_t, \tau^k_{t+j})$	j= -2	-1	0	+1	+2
data	0.63	0.69	0.67	0.56	0.39
model	0.59	0.89	1.00	0.89	0.59

Importantly, these predictions are in line with the data. This is illustrated in Table 6, which compares the cyclical properties of the wedges in our model to those measured in US data under standard business-cycle accounting (Chari, Kehoe and McGrattan, 2005). Evidently, our model captures quite well, not only the procycicality of the two wedges with output at various leads and lags, but also their co-movement. The only caveat is that our model over-estimates these correlations, and under-estimates the standard deviations of the wedges. This need not be surprising since we have ruled out other shocks or frictions that could add independent variation in the wedges.

Figure 3 further illustrates the quantitative fit of our model by considering the following exercise. We first find the series of the two shocks that minimize the distance of the model's predicted series for output, employment, and investment from the corresponding ones in the data; we next feed the confidence shock *alone* in the model to derive the model's predicted time series of two wedges; we finally plot these series against those obtained from the data. As is evident from this figure, the confidence-driven wedges in the model track quite well the actual wedges in the data.

Figure 3: Wedges

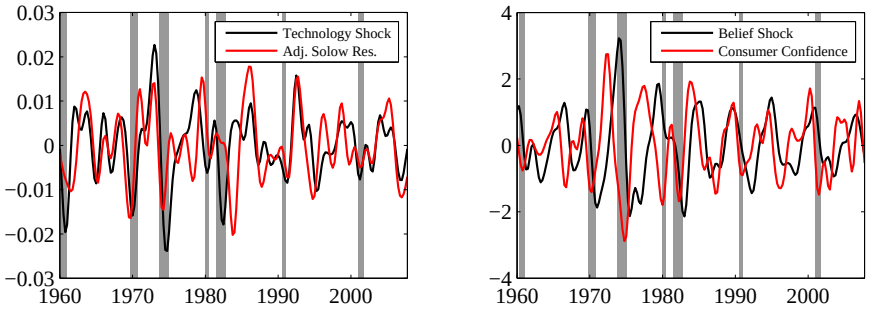


Note: Data, Model. Shaded areas correspond to US recessions, as identified by the NBER.

Empirical counterparts to structural shocks in the model

We have seen that our model matches quite well a variety of empirical regularities regarding the joint dynamics of output, employment, investment and consumption. However, it is unclear whether the technology and confidence shocks in the model are closely connected to, respectively, production possibilities and "market psychology" in the real world. This is an important consideration that extends well beyond our exercise: the inability to directly measure the structural shocks featured in modern DSGE models limits their empirical appeal and often gives them a black-box flavor.

Figure 4: Shocks vs Empirical Counterparts



Note: Data, Model. The Consumer Sentiment Index (data) and the confidence shock (model) were centered and standardized for illustrative purposes.

We seek to mitigate this concern by investigating the connection between the shocks in our model and some plausible real-world counterparts. In particular, we first recover the series of the two structural shocks that minimize the distance between the model's predicted series of output, hours, and investment, and the corresponding series in the data. We then contrast the recovered series of the shocks with the series of Fernald's utilization-adjusted TFP measure (a proxy for technology) and the University of Michigan Consumer Sentiment Index (a proxy for market sentiment).

Figure 4 reports the findings of this exercise (the left panel for the technology shock, the right

panel for the confidence shock). Even though we did not use any information on Fernald's TFP measure and the Sentiment Index either when we calibrated the model or when were recovered the structural shocks, the latter turn to be highly correlated with their empirical counterparts. The same property holds if we replace the Michigan Sentiment Index with either of the Conference Board's Indices of Consumer or Producer Confidence (all these indices are highly correlated at business-cycle frequencies), as well as if we replace the confidence shock with a plausible theoretical analogue of these indices.[16] Notwithstanding the inherent difficulty in interpreting these indices and in mapping them to the model, we see this finding as an additional validation of our interpretation of the two structural shocks and of the usefulness of the model.

6 Extension and Estimation

The preceding analysis has shown that the introduction of higher-order uncertainty within a simple, calibrated version of the RBC model offers a parsimonious yet potent explanation of a number of key business-cycle facts. In this section, we expand on the quantitative role of "confidence" within richer, estimated, models of both the RBC and the NK variety.

The exercises we conduct are centered around the following questions: (i) Does variation in confidence account for a significant fraction of the observed business-cycle volatility once other plausible structural shocks are allowed to drive the business cycle? (ii) Does the estimated confidence shock track the empirical shock we identified in the data? (iii) Does the introduction of confidence shocks improve the empirical performance of both RBC and NK models? (iv) Are nominal rigidities essential for capturing the observed business cycles and the apparent deficiency in demand during recessions?

To address these questions, we extend the analysis to two estimated variants of our baseline model that allow for multiple structural shocks in addition to our confidence shocks. The two variants differ with regard to whether prices are flexible, as in our baseline model, or sticky, as in the NK paradigm. In the flexible-price/RBC model, we allow for two permanent technology shocks, one on neutral TFP and another on investment-specific technology; for two transitory technology shocks; for two news shocks on the two types of technological change; and for a shock to government spending. In the sticky-price/NK model, we also add a monetary shock.

The choice of these shocks is motivated by multiple considerations. First, we rule out ad-hoc shocks to the labor and Euler wedges and, instead, allow for shocks that, in our eyes, have a plausible structural interpretation. Second, some researchers have argued that investment-specific technology shocks are more important than neutral, TFP shocks in driving business cycles. Third, news shocks are obvious competitors to our confidence shocks insofar as expectation-driven fluctuations

[16] In particular, we use the model to create time series of the typical agent's expectation of the level of output in the next quarter, the next year, and the next five years. We take an average of them and interpret this as the model-based analogue of the Michigan Sentiment Index. Using this measure instead of the confidence shock gives a somewhat better fit than the one seen in Figure 4.

are concerned, and have also been on center stage in recent business-cycle research. Fourth, when prices are sticky, monetary, fiscal shocks, and news shocks can operate as aggregate demand shocks, thus directly competing with the mechanics of our confidence shocks. Finally, the introduction of transitory technology and policy shocks maximizes the chance that shocks other than our confidence shocks pick up the transitory fluctuations in the data once the model is estimated. (With the last point in mind, in the end of this section we consider a further extension that adds a markup shock as well certain proxies of financial shocks.)

In addition to these shocks, we allow for investment adjustment costs (IAC) as in Christiano, Eichenbaum and Evans (2005). This modeling feature has played a dual role in the literature: as a source of inertia and persistence in the response of macroeconomic quantities to the underlying shocks; and as a mechanism that helps transform various shocks, such as news, investment-specific, and discount-rate shocks, into something that looks like an aggregate demand shock.

We are suspicious about both roles. First, like Sims (2001), we suspect that much of the apparent inertia in the response of macroeconomic quantities to various shocks originates from informational frictions rather than adjustment costs. Second, we do not believe that the notion of "aggregate demand shocks" has to depend critically on this modeling feature. (We revisit this point in the next section.) Last but not least, we have not seen any compelling micro-level evidence indicating the presence of this kind of adjustment costs at the levels required by the existing DSGE literature. Despite these concerns, we opt to incorporate this modeling feature in order to keep our exercise close to DSGE practice and to give a better chance to the aforementioned shocks to pick up the observed fluctuations. Shutting down this feature only reinforces the results we present below.

What comes out of all this? Confidence shocks persevere. (i) They remain the main driver of business cycles in our estimated RBC and NK models, despite the presence of the aforementioned competing shocks. (ii) They mirror the empirical shock we identified in Section 2. (iii) They improve the empirical fit of both our RBC and NK models, but more so in the RBC case. (iv) They displace the essential role that nominal rigidities play in helping the NK model match the observed fluctuations and the apparent deficiency in aggregate demand during recessions.

Estimation

The two models under consideration are estimated by a Bayesian maximum likelihood method in the frequency domain, as in Christiano and Vigfusson (2002) and Sala (2013). The details are spelled out in the appendix. The benefit of this approach, relative to estimation in the time domain, is that it allows us to guide the estimation of the model from its performance at business-cycle frequencies (i.e., between 6 and 32 quarters). We find this a priori desirable, because we are interested in the ability of RBC and NK models to capture business-cycle phenomena—not phenomena working at very low frequencies, such as demographics and trend growth.[17]

[17]By the same token, we avoid having to introduce ad hoc modeling features, such as slow-moving shocks in the disutility of labor or in monetary policy, which DSGE exercises such as Smets and Wouters (2007) typically require

The dataset comprises information about GDP, consumption, investment, hours worked, the relative price of investment, the inflation rate, and the federal fund rate for the period 1960Q1 to 2007Q4. The first 5 variables are evaluated in logarithm; they are also de-trended, as our model abstracts from long-run growth. Note that we exclude the Great Recession from our dataset. We do so because the RBC and NK models we consider in this paper abstract from financial frictions, and are therefore less suited for studying this particular episode.

Our NK model is estimated on the basis of all the aforementioned seven variables. By contrast, our RBC model is estimated on the basis of real quantities only (GDP, consumption, investment, hours worked, and the relative price of investment), because this kind of model is not meant to capture the nominal side of the economy. For comparison purposes, however, we find it useful to obtain that model's predictions for interest rates and inflation. We thus augment our estimated RBC model with a Taylor rule for monetary policy and, in what follows, report the model's predictions for both real and nominal variables.[18]

The priors we assume for the preference and technology parameters are in line with our earlier calibration exercise, as well as with the pertinent literature. The priors for the nominal side of the model (the Calvo probability of resetting prices and the specification of the Taylor rule) are also drawn from the literature. These priors are reported in Table 13 in the Appendix. Finally, the priors for the shocks are reported in Table 14 and are guided by two principles. First, we do not a priori favor either the belief shock or any other shock. And second, we impose a tight prior only on the persistence of the transitory shocks, in order to assist the estimation disentangle them from the permanent shocks. Accordingly, the standard deviations of the innovations to all shocks are assumed to follow uninformative inverse Gamma distributions centered at 1, while the autocorrelation parameters of the AR(1) shocks to technology, government spending, and beliefs are assumed to be Beta distributed around 0.5 with standard deviation 0.1.

Posterior distributions are obtained relying on the MCMC algorithm. We generated 2 chains of 100,000 observations each. The posterior estimates of the parameters are reported in the last two columns of Tables 13 and 14.

As seen in Table 13, the posterior estimates of the Frisch elasticity of labor supply, the capital share, and the degree of complementarity are similar across the two versions of the model. By contrast, the two models require different utilization and investment-adjustment margins. In particular, the elasticity of utilization is lower in the sticky price version of the model than in the flexible price version.[19] Perhaps more interestingly, the flexible-price model validates the level of IAC estimated by Eberly, Rebelo and Vincent (2012) on firm-level data, and the sticky-price ver-

in order accommodate the medium-run dynamics in hours and inflation.

[18]Note that the familiar neoclassical dichotomy applies: once we fix the real side of the model, the specification of the Taylor rule matters only for nominal variables, not for real quantities. Accordingly, we first estimate the real side of the model on real quantities alone, and then estimate the Taylor rule on interest rates and inflation.

[19]Note that the log-linearized version of the utilization decision is given by $\widehat{r}_t^k = \frac{\sigma_z}{1-\sigma_z}\widehat{z}_t$, where $\widehat{r}_t^k$ denotes the rental rate of capital and $\widehat{z}_t$ the utilization rate (both in log-deviations from steady state). It follows that the elasticity of utilization with respect to the rental rate of capital is given by $(1-\sigma_z)/\sigma_z$.

sion turns out to require even lower adjustment costs (0.11 vs 0.28). This contrasts existing DSGE exercises, which typically require much larger IAC costs. We revisit this point in the sequel.

Finally, two points are worth noting regarding the posteriors of the shock parameters, which are reported in Table 14. First, the estimated persistence of the transitory shocks is quite low. This is driven partly by of our choice of priors, but it also in line with the evidence we provided in Section 2, according to which most of the business-cycle movements in the data appear to be driven by transitory forces. Second, the estimated size of the higher-order uncertainty in the model (the confidence shock) is sizable, but not implausibly large relatively to the estimated size of the payoff uncertainty (the shocks to TFP and other fundamentals).[20]

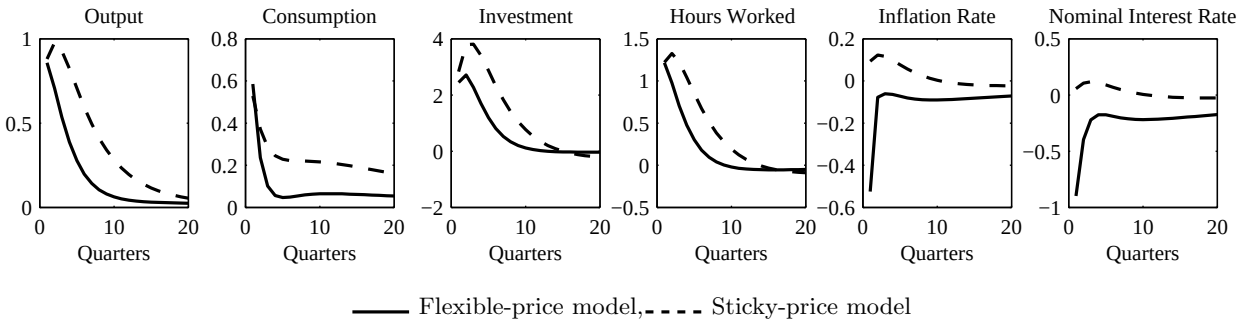
IRFs

Figure 5 depicts the IRFs of our estimated flexible-price and sticky-price models to a positive confidence shock. As far as real variables are concerned, the IRFs are similar between the two models (and are also in line with those in our baseline model). Where the two models differ is in the responses of nominal variables: the flexible-price model features a strongly countercyclical response in inflation and the nominal interest rate, whereas the sticky-price model predicts a modest pro-cyclical response. Therefore, if we interpret the confidence shock as a formal counterpart of the empirical shock we documented in Section 2, both models do a good job in capturing the real dimension of the data, but only the sticky-price model does a good job in the nominal dimension.

The IRFs for the other, more conventional, shocks are reported in the Appendix. They are all in line with those found in the literature. What we nevertheless find noteworthy is that the IRFs (and also the variance contributions) of some of these shocks are distinctly different between the flexible- and the sticky-price model. This underscores a point we anticipated earlier, namely that some of these conventional shocks operate very differently under RBC and NK models. Furthermore, this difference increases with the presumed level of IAC. Given that we do not have a strong prior about what the "right" model is, we view it as an advantage of our approach that the dynamics effects of confidence shocks on real quantities turn out to be robust across the two classes of models.

[20] Because of persistence of the shocks and the forward-looking aspects of our model, it is not appropriate to measure the relative magnitude of different shocks by simply comparing the standard deviations of the corresponding innovations. Instead, some kind of present-value metric seems desirable. Thus let F_t^T denote the typical agent's first-order belief of the present value of TFP from period t to period $t+T$, discounted by β, and evaluated conditional on the information available at the end of period $t-1$; let S_t^T denote the corresponding second-order belief; and let $B_t^T \equiv S_t^T - F_t^T$ denote the difference between the two, which obtains only because of the ξ_t shock. We can think of the variance of F_t^T as a measure of the TFP uncertainty faced by the agent, as of period $t-1$ and over the next T periods; of the variance of B_t^T as a measure of the corresponding higher-order uncertainty; and of the ratio of the latter to the former as a metric of their relative importance. For the flexible-price model, this ratio turns out to be 0.13 when $T = 8$ years and 0.04 when $T = \infty$. For the sticky-price model, the corresponding ratios are are 0.42 and 0.13. Finally, these ratios would be even lower if we were to take into account the other shocks in the model (investment-specific, fiscal, monetary, etc). This explains the metric by which the estimated higher-order uncertainty is modest relative to the estimated payoff uncertainty.

Figure 5: Theoretical IRFs to a confidence shock



Variance contribution

We now turn attention to the contribution of the various shocks to macroeconomic volatility over business-cycle frequencies (6-32 quarters). The relevant variance decompositions are reported in Table 7 for our estimated flexible-price model, and in Table 8 for its sticky-price variant.

The most noteworthy piece of information in these tables is that the confidence shock emerges as the single most important source of fluctuations in output, employment, and investment. In the flexible-price model, the confidence shock explains 39% of the volatility in output, 34% of the volatility in investment, and 68% of the volatility in hours. In the sticky-price model, the corresponding numbers are 55%, 47% and 68%. To our surprise, the inclusion of nominal rigidities therefore appears to increase the estimated significance of confidence shocks.

This finding is not driven by our choice of parameters/priors. In the flexible-price model, the variance contribution of the confidence shock evaluated at our priors is 5% for output, 4% for investment, and 17% for hours; and in the sticky-price model, these numbers are 2%, 1%, and 3%. It is the data that favor the confidence shock as the main driver of the business cycle.

Due to its transitory nature, the confidence shock has a more modest effect on consumption: it explains between 19% and 25% of the volatility of consumption, depending on which model we look at. This finding may contradict the view that business cycles are driven by shifts in consumer spending: under that view, one would probably expect the underlying shock to drive a larger fraction of the variance in consumption than that in output, due to countervailing movements in investment. Indeed, this is precisely what happens with two known formalizations of this view: with discount-factor shocks in the calibrated NK model we review in the Appendix, and with the noise shock in Blanchard et al (2012). Yet, the evidence presented in Section 2 favors the picture painted by our model: the identified business-cycle shock had a larger contribution to the volatility of output, employment and investment than to the volatility of consumption.

Another noteworthy finding is that the two news shocks explain a sizable component of the volatility in investment, but a relatively smaller component of output and hours. Furthermore, as can be seen in Tables 15 and 16 in the Appendix, news shocks drive a relatively small fraction of the co-variance between these variables: in the sticky-price model, for example, the combination

Table 7: Variance Decomposition (6–32 Quarters) in Flexible-Price Model

	Y	I	h	Y/h	P_i	C
Confidence shock	39.20	34.23	68.45	1.41	0.00	18.53
TFP shock: permanent	17.75	1.59	1.05	50.47	0.00	37.96
TFP shock: transitory	22.41	18.43	6.09	35.17	0.00	12.91
Investment shock: permanent	0.97	0.43	1.69	3.89	80.47	0.37
Investment shock: transitory	0.55	1.98	0.95	0.02	11.10	1.75
News shock: TFP	13.29	24.08	11.89	8.88	0.00	13.55
News shock: investment	5.17	16.80	8.97	0.16	8.43	13.36
Fiscal shock	0.65	2.46	0.90	0.01	0.00	1.57

Table 8: Variance Decomposition (6–32 Quarters) in Sticky-Price Model

	Y	I	h	Y/h	P_i	π	R	C
Confidence shock	54.59	46.90	68.15	17.58	0.00	23.06	25.55	24.57
TFP shock: permanent	11.16	1.71	0.60	44.78	0.00	1.21	1.71	26.02
TFP shock: transitory	4.85	4.01	0.68	27.13	0.00	2.89	5.73	2.70
Investment shock: permanent	1.14	1.32	1.36	1.42	74.42	2.77	3.91	0.97
Investment shock: transitory	1.54	3.14	1.89	0.13	8.89	2.56	4.26	1.69
News shock: TFP	6.79	10.00	3.27	8.21	0.00	3.75	5.61	8.82
News shock: Investment	10.58	25.67	12.46	0.63	16.69	32.62	49.54	29.04
Fiscal shock	1.19	0.80	1.34	0.03	0.00	0.55	1.04	0.61
Monetary shock	8.17	6.46	10.26	0.08	0.00	30.59	2.65	5.58

of the two news shocks explains 17% of the co-variance between hours and output and 26% of the co-variance between investment and either output or hours, whereas the corresponding numbers for the confidence shock are between 54% and 66%. Last but not least, both types of news shocks contribute towards a *negative* co-movement between consumption and the rest of the macroeconomic quantities, leaving the confidence shock to capture more than 100% of the co-variance between consumption and either hours or investment. Therefore, although our findings are consistent with the notion that a significant fraction of investment fluctuations is driven by shifts in expectations of long-run profitability, transitory fluctuations in confidence emerge as the dominant force behind the business-cycle co-movement in all the key macroeconomic variables.[21]

Business-cycle moments

We now proceed to evaluate the empirical performance of our estimated RBC and NK models in regard to standard business-cycle moments, namely, the volatilities and correlations of the Band-Pass filtered series of the seven macro variables in our data set. There are three questions of interest here. First, how well the two models reproduce the moments in the data. Second, how their fit depends on the presence of either confidence shocks or nominal rigidities. Finally, how the two models compare to the state of the art in DSGE literature.

To address these questions, Table 9 reports the following information: a standard variety of business-cycle moments in the data (first column); the corresponding moments predicted by our two models (second and third column); those predicted by variants of these models that remove the confidence shock (fourth and fifth column); and finally the moments predicted by the model of Smets and Wouters (2007) (last column).[22]

The following key lessons emerge from inspecting Table 9.

First, our two models both do a very good job in reproducing the key business-cycle moments in the data. The flexible-price model does marginally better with regard the the correlation of labor productivity with either output or hours, but it misses by a significant margin the corresponding correlations for inflation. By contrast, the sticky-price model does a very good job on both the real

[21] In understanding these findings, it is important to underscore the role played by IAC. Suppose we re-estimate our NK after imposing a much higher value for IAC than our current estimate, namely $\phi \geq 4$ instead of $\phi = 0.2$. Under this scenario, news (or noise) shocks trigger positive co-movement between consumption, output, investment and hours. This scenario corresponds, in effect, to the one studied in Blanchard et al (2013), where the IAC is in the neighborhood of $\phi = 4$, or Lorenzoni (2009), where the IAC is infinite since that paper assumes away investment. But recall that such high values for ϕ are inconsistent with the firm-level estimates of Eberly, Rebelo and Vincent (2012), which hover in the neighborhood of 0.3. Furthermore, if we allow for a diffuse prior centered at $\phi = 4$, then our estimation brings ϕ down to values close to 0.3, recovering the results we have documented. This highlights that, within the NK framework, certain parameterizations guarantee that news and noise shocks trigger realistic business cycles, yet these parameterizations do not appear to be favored by the data once our confidence shocks are present.

[22] In our view, Smets and Wouters (2007) has too many ad hoc ingredients whose structural interpretation is stretched. Yet, it severs as a useful metric for evaluating the relative performance of our models, both because it is a benchmark in the DSGE literature and because it was designed to capture the business-cycle data as well as possible.

Table 9: Band-pass Filtered Moments

		With Confidence Shock		Without Confidence Shock		
	Data	Flex-Price	Sticky-Price	Flex-Price	Sticky-Price	SW (2007)
Standard Deviation						
y	1.46	1.46	1.76	2.39	2.22	1.43
c	0.86	0.91	0.84	1.70	1.16	1.03
i	5.77	5.96	7.83	9.31	11.22	4.61
h	1.71	1.59	2.20	3.60	2.91	0.91
q	0.71	0.70	0.80	2.47	1.05	0.43
y/h	0.99	0.92	0.81	1.64	1.00	0.86
π	0.23	0.28	0.37	0.31	0.33	0.30
R	0.35	0.41	0.34	0.46	0.51	0.33
Correlation with output						
c	0.86	0.48	0.41	-0.11	-0.40	0.66
i	0.94	0.88	0.92	0.86	0.95	0.73
h	0.88	0.89	0.92	0.94	0.94	0.81
q	-0.31	-0.13	-0.17	-0.23	-0.49	-0.27
y/h	0.52	0.59	0.11	0.43	0.74	0.79
π	0.21	-0.35	0.68	-0.02	0.75	0.15
R	0.37	-0.57	0.55	-0.12	0.79	0.07
Correlation with hours						
c	0.87	0.23	0.26	-0.33	-0.57	0.57
i	0.85	0.89	0.91	0.91	0.96	0.66
q	-0.30	-0.18	-0.20	-0.31	-0.56	-0.34
y/h	0.07	0.18	-0.25	0.23	0.49	0.29
π	0.39	-0.33	0.79	0.00	0.85	0.25
R	0.60	-0.55	0.64	-0.10	0.88	0.25

and the nominal dimensions at the same time. In this respect, the presence of nominal rigidities is not essential for matching the real fluctuations in the data, but it is crucial for matching the cyclical behavior of inflation and interest rates.

Second, the only noticeable failure is that both models underestimate the pro-cyclicality of consumption. What causes this failure is the presence of news shocks: as explained earlier on, these shocks contribute towards counter-cyclical consumption. Removing these shocks from our two models, and re-estimating them, fixes this dimension without any noticeable sacrifice in the rest of their empirical fit. In fact, the fit of the models remains roughly as good if we remove either the two news shocks, or the transitory shocks to both types of technology and to government spending, or all these five shocks at the same time (see Table ?? in the Appendix). Finally, the monetary shock is crucial only in matching the nominal quantities. In short, a bare-bone version of our RBC model that features only the two permanent technology shocks and the confidence shock delivers roughly the same empirical performance as the richer models we are considering.

Third, crucial for this empirical performance is the presence of the confidence shock: if we remove this shock and re-estimate our models, both models end up delivering too high volatilities for multiple real variables and, most importantly, a strongly counter-cyclical consumption. Among the types of structural shocks we have considered, the confidence shock therefore emerges as the the only one that can deliver the right co-movement among all variables at the same time.[23]

Fourth, although the introduction of sticky prices helps improve the overall fit of the model (compare the third to the second column), the presence of confidence shocks appears to be more crucial for the fit of the model than the assumption of sticky prices. In particular, the sticky-price model without confidence shocks (fifth column) is clearly outperformed by the flexible-price model that contains confidence shocks (second column). We will revisit the relative importance of confidence shocks and nominal rigidities in the last part of the paper.

Finally, our NK model fares quite well in comparison to Smets and Wouters (2007) vis-a-vis both the real and the nominal aspects of the data. In particular, our model overestimates the cyclicality of inflation, but does a better job in matching the co-movement of nominal interest rates with real economic activity (check the correlations between the interest rate and either output or hours reported in the third and last columns of Table 9). Although we are primarily interested in understanding the observed fluctuations in real economic activity as opposed to the accompanying monetary phenomena, we find this last point to be remarkable, given that Smets and Wouters (2007) incorporate a more elaborate Taylor rule and a richer specification of shocks specifically designed to capture the nominal side of the data.

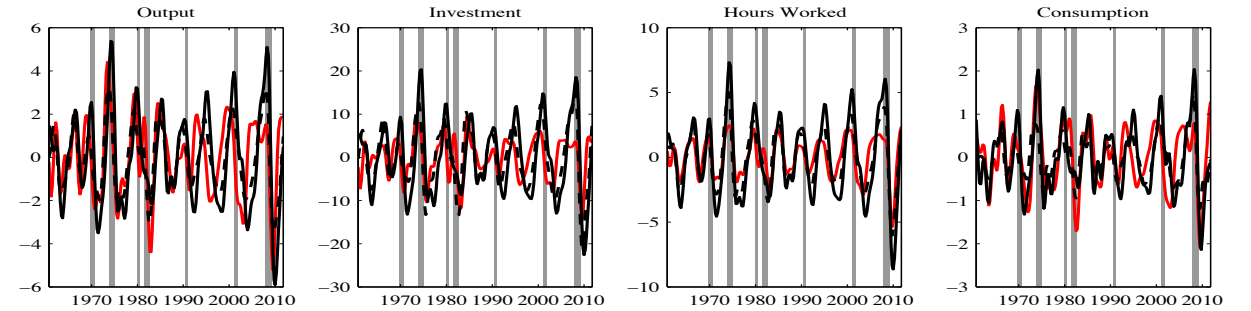
[23]Of course, this last statement applies to the models we have consider here, not other models. In Smets and Wouters (2007), the right co-movement is delivered primarily by the combination of the shocks to the Euler wedge, the labor wedge, and the monopoly markups.

Going back to the VAR evidence

We have already discussed that the IRFs of the confidence shock in the model mirror the IRFs to the main transitory shock we documented in Section 2. Figure 6 shows that this resemblance carries over to the estimated paths of the two types of shocks.

The solid black lines in the figure report the counterfactual paths for output, hours, investment and consumption that obtain from our flexible-price model, when we feed it with the estimated path of the confidence shock only. The dashed black lines give the same paths for the the sticky-price model. Finally, the solid red lines give the counterfactual paths that obtain from the VAR we run in Section 2, when we feed it with the identified "main shock" only. All series are band-pass filtered, in order to reveal their co-movement (or lack thereof) at business-cycle frequencies.

Figure 6: Main transitory shock (data) vs confidence shock (model)



main transitory component in the US data (as identified in Section 2), ——— confidence-driven fluctuations in the flexible-price model, - - - - confidence-driven fluctuations in the sticky-price model

The match is quite good: the time series corresponding to the confidence shock in either the flexible-price or the sticky-price specification of the model co-move almost perfectly with the times series corresponding to the main transitory shock in the data. The only noticeable discrepancy is in the magnitudes of the two series: the paths predicted by the theory are more volatile than those generated by the VAR, suggesting that the confidence shock is perhaps *too* potent within the theory. Importantly, this success is specific to the confidence shock: if we consider any of the other structural shocks in the model, there is no resemblance to the main transitory shock in the data, whether in terms of IRFs or in terms of paths.

The are two caveats, however, in the conclusions we are drawing from the above findings. First, we could be subject to the critique of Chari, Kehoe and McGrattan (2011). And second, the empirical shock we have identified could be a convoluted combination of multiple structural shocks, even in the case that our model is the "true" data generating process.

To address these issues and shed further light on the empirical performance of our models, we consider the following exercise: we generate artificial data from our two models and run on them exactly the same VAR as the one we run on actual US data in Section 2. The idea is that we stop giving any structural interpretation to the evidence of Section 2 and, instead, use the VAR merely

as a filter: we ask whether applying this filter on our models gives the same picture as applying this filter on US data.

Figure 7: IRFs to "Main Shock" in the Data vs in the Models

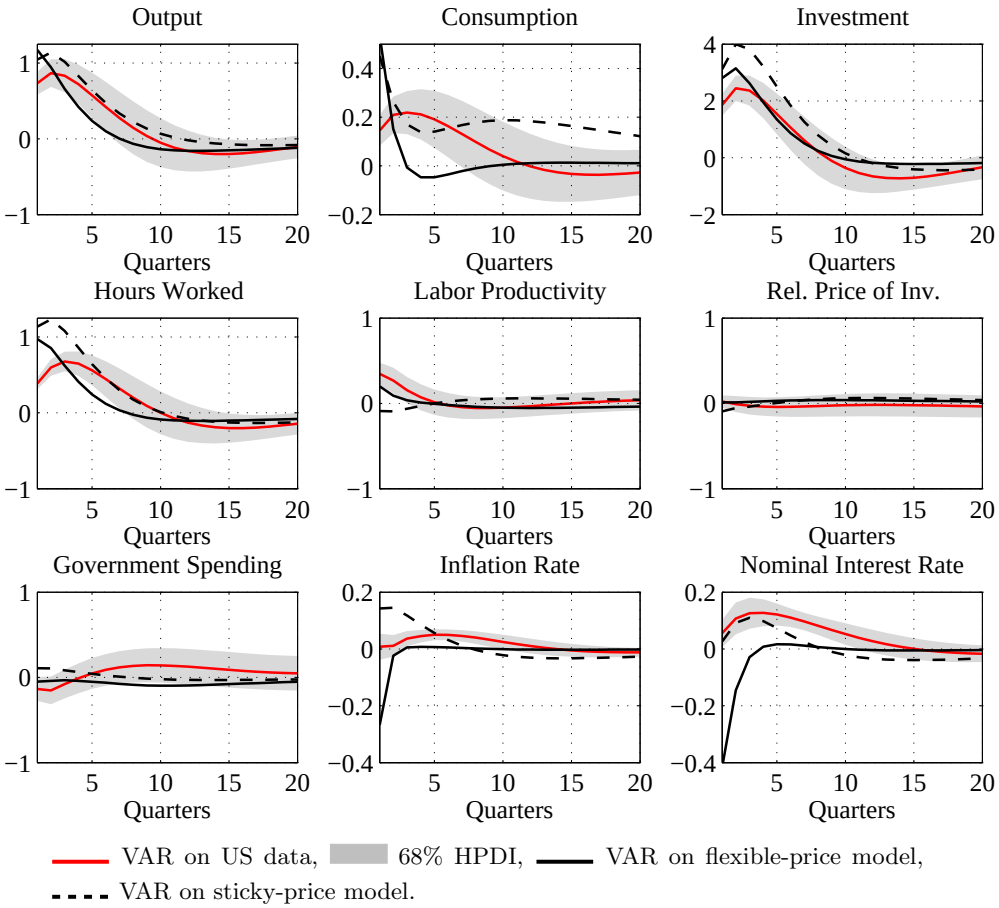


Table 10: Contribution of Transitory Shocks (6–32 Quarters)

	Y	C	I	h	Y/h	P_i	G	π	R
Identified Main Shock									
Data	51.69	25.55	55.78	49.36	21.42	6.07	9.89	18.50	29.94
	[42.74,60.69]	[15.11,35.97]	[46.81,64.78]	[40.66,58.08]	[12.24,30.48]	[1.60,10.69]	[3.93,15.91]	[9.76,26.63]	[15.25,44.94]
Flex. Price	41.46	21.73	46.99	51.36	11.13	8.02	8.36	13.42	13.55
Sticky Price	58.02	18.50	59.30	65.14	9.81	10.94	7.04	37.94	33.09
All Transitory Shocks									
Data	78.47	64.36	79.64	75.99	79.36	61.25	65.89	68.66	73.51
	[68.91,88.11]	[50.22,78.44]	[70.39,88.91]	[65.75,86.41]	[69.11,89.76]	[42.57,80.28]	[51.45,80.32]	[55.41,81.68]	[59.14,87.80]
Flex. Price	56.28	68.72	55.43	61.07	69.30	40.46	71.42	69.21	68.36
Sticky Price	72.92	64.06	70.17	78.27	67.78	46.77	78.55	77.45	72.73

Note: *All Transitory* denotes the contribution of all transitory shocks to the dynamics of the variables at the business cycle frequencies (6–32 quarters). *Main Transitory* denotes the main transitory component of the data, as obtained by the combination of VAR residuals that explain the maximal share of the volatility of investment and hours at business cycle frequencies. *Main Total* denotes the main shock that obtains if we refrain from filtering out the permanent shocks. 68% HPDI into square brackets.

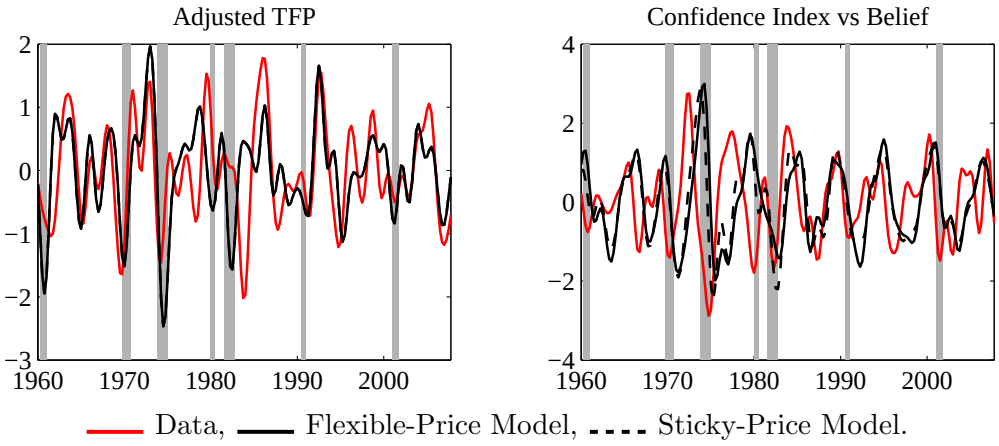
Figure 7 answers this question in terms of the IRFs to the identified main shock. Both models pass the test with regard to the real quantities: the IRFs that obtain once we run the VAR on the model mirror quite well those that obtained when ran the VAR on US data. One exemption to this statement is consumption: the theory predicts an overshooting in consumption that is not seen in the data. Finally, where the sticky-price model outperforms the flexible-price one is with regard to the behavior of inflation and interest rates.[24]

Table 10 compares the variance decompositions that obtain from running the VAR on our two models to those that obtain from running the VAR on US data. These decompositions do not have a structural interpretation, because there is no one-to-one mapping between the shocks in the model and the shocks in the VAR. They nevertheless provide evidence of the ability of our models to capture the permanent and transitory components of the data: whether we look at the decomposition between permanent and transitory components, or at the variance contribution of the identified main transitory shock, the match is, once again, excellent.

Shocks vs empirical proxies

Figure 8 revisits an exercise we conducted in our baseline, calibrated model: it compares the estimated technology and confidence series from our models to their empirical counterparts, the Fernald's utilization-adjusted TFP measure and the University of Michigan Consumer Sentiment Index. Once again, the match is good.

Figure 8: Band-passed Shocks (6–32 Quarters)



[24]Interestingly, both features are driven by the assumption of investment adjustment costs (IAC): if we set $\phi = 0$, the aforementioned overshooting in consumption disappears and the counter-factual nominal predictions of the flexible-price model are overturned. Having said this, the estimation appears to favor modest IAC because this generate hump-shaped dynamics and improves the fit of the NK vis-a-vis the nominal variables.

7 On Nominal Rigidities and Deficient Aggregate Demand

We conclude the paper by touching on one of the quintessential questions in macroeconomics: how central are nominal rigidities in understanding the observed fluctuations and the apparent deficiency in aggregate demand during recessions?

As noted in the introduction, there is hardly any consensus among macroeconomists on the answer to this question, or even on what "deficient aggregate demand" might mean. Nevertheless, most macroeconomists would probably agree on the following two premises: (i) in the RBC framework, there is no operational distinction between aggregate demand and aggregate supply; (ii) in the NK framework, such a distinction becomes possible by considering the deviations of the equilibrium sticky-price allocations from the underlying flexible-price ones. Accordingly, the notion of "deficient aggregate demand" is captured in the NK framework by a negative output gap between the sticky- and the flexible-price allocations.

Our approach, instead, permits us to operationalize the notion of "deficient aggregate demand" by considering the gap between the equilibrium allocations that obtain in the presence of a negative confidence shock and those that would have obtained in its absence. And, as already discussed, this gap in turn can be interpreted as the symptom of imperfect coordination.

In our view, it is not clear which approach is more empirically plausible, nor does it have to be that the one excludes the other. There are nevertheless certain properties that we like with our approach and that are not shared by the standard NK approach.

First, our approach works in both the RBC and the NK framework. We view this as an advantage, because we do not believe that the empirical regularities that have been associated with the notion of aggregate demand depend as tightly on nominal rigidities as presumed by the NK framework.

Second, our approach gives a structural interpretation of the observed recessions that is not vastly different as one varies the degree of nominal rigidity. By contrast, this is not the case with standard NK exercises: if one takes an estimated NK DSGE model such as Smets and Wouters (2007) and assume away the nominal rigidities, the model will predict vastly different macroeconomic outcomes.

We illustrate this point in Figures 9 and 10. These figures consider the following exercise. First, take either our estimated NK model or its estimated variant that removes the confidence shock. In each case, fix the estimated paths of the shocks, and consider the counterfactual paths of output and hours that obtain when one assumes away the nominal rigidity. The counterfactuals for our model are given by the blue lines in the figures, while the counterfactuals of the model without the confidence shock are given by the dotted black lines. (The red lines give the actual data, which are perfectly match by both models when the nominal rigidity is present.) fix the estimated shocks, and consider the counter assume away the nominal rigidity, ; this gives the dashed blue lines in the figures. rigidity or (ii) the confidence shock, and obtain the predicted time series for output during the recession episodes in our sample.

Figure 9: Counterfactual Output, Removing Nominal Rigidity

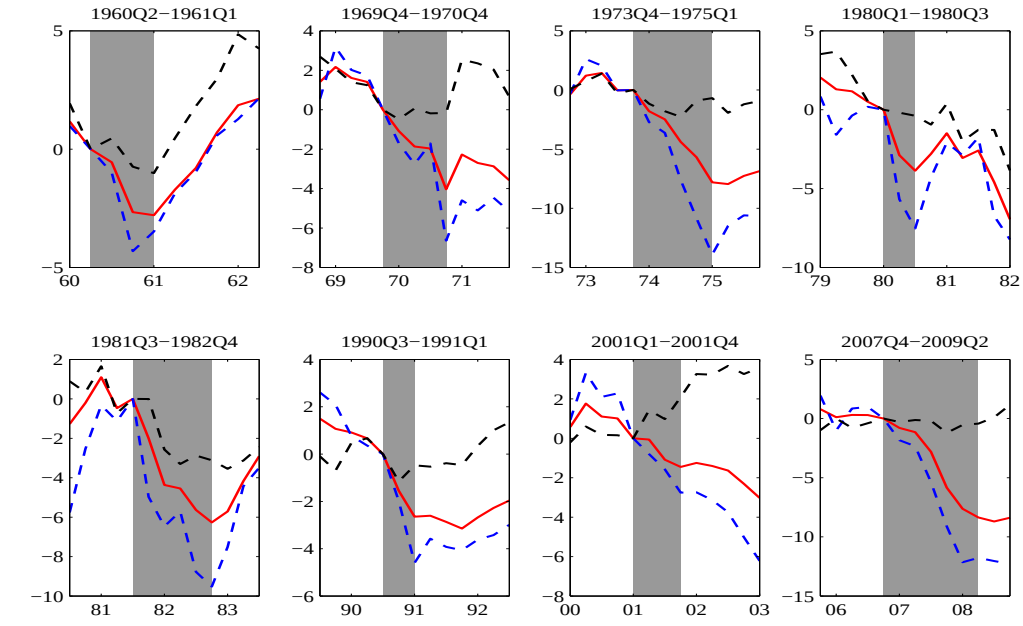
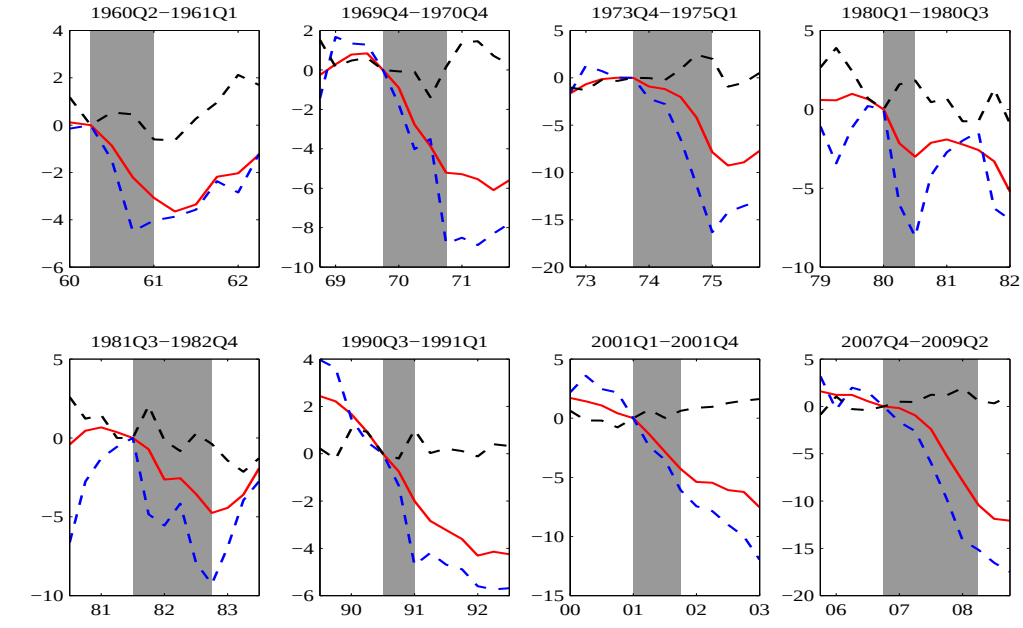


Figure 10: Counterfactual Hours, Removing Nominal Rigidity



The lesson that emerges is quite simple. Under the lenses of the model that contains the confidence shock, the recessions would look very similar if prices had been flexible. In fact, the presence of nominal rigidity appears to mitigate the recession; this is because the nominal rigidity empowers the monetary policy to counteract the recession. By contrast, under the lenses of the model that rules out the confidence shock, the recessions would look very different, and in some cases would have been replaced by booms, if nominal rigidities were absent. This is because nominal rigidity is instrumental, in the model, to transforming certain structural shocks from expansionary to recessionary. Of course, this may not be a problem if one sees nominal rigidities as the source of all evil. In our eyes, however, it looks unappealing to have a theory of recessions that hinges so crucially on the bite of nominal rigidities.

To reinforce this point, we conduct the following formal test. We consider four models, corresponding to different assumptions on whether prices are sticky or flexible and on whether confidence shocks are present or not. Starting from a uniform prior across these four models, and focusing on real quantities, we ask what the posterior odds are that the data are generated by model A vs model B (where model A varies across columns and model B across rows).[25] As evident in the second row of this table, the odds that the data are generated by the flexible-price model that shuts down the confidence shock are zero relative to either of the sticky-price models. In the absence of confidence fluctuations, sticky prices are therefore essential for the ability of the theory to match the observed fluctuations in real economic activity. At the same time, as evident in the top-right element of this table, the odds that the data are generated by the with-confidence flexible-price model rather than the without-confidence sticky-price model are 86%. In this regard, allowing for confidence fluctuations appears to be more essential than allowing for nominal rigidities, insofar as we are concerned only with fluctuations in real economic activity. Sticky prices become essential only if we also wish the theory to match the inflation and the nominal interest rate data.

Table 11: 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	0.61	0.14
flex prices, without confidence	–	1.00	1.00
sticky prices, with confidence	0.00	–	0.10

Note: This table reports the posterior odds of four alternative models, starting from a uniform prior and estimating the models on the real data. The models differ on whether they assume flexible or sticky prices, and on whether they contain a confidence shock or not.

[25] To give all models the best chance to fit this data and make the comparison fair, for this table we estimate all four models on the real data only. See Appendix for the details and also for the complementary exercise that brings the nominal data as well.

We interpret the above findings as follows. The RBC framework alone does not offer a sufficiently good structural interoperation of the observed fluctuations. Instead, the data are asking for something that looks like fluctuations in "aggregate demand fluctuations," by which we mean the empirical regularities documented Section 2. Within the standard paradigm, this favors sticky prices. But once one accommodates the notions of coordination frictions, higher-order uncertainty, and confidence fluctuations in the way we propose in this paper, sticky prices are not more essential: our approach achieves what the data wish whether with or without nominal rigidity.

8 Conclusion

Standard macroeconomic models assume away strategic uncertainty and frictions in coordination. In this paper we propose a modeling trick for incorporating the positive content of such frictions without sacrificing any of the tractability that macroeconomists are accustomed to. We then use this to develop a novel, and quantitatively potent, structural interpretation of the business cycle.

More specifically, we show that the introduction of a certain type of higher-order belief dynamics, which we interpret as fluctuations in "confidence", achieves the following:

- It offers a parsimonious explanation of multiple salient features of the data, including: standard business-cycle moments; the cyclicality and co-movement of measured wedges; VAR evidence about the co-movement of key macroeconomic variables; and the cyclical behavior of empirical measures of confidence.

- It appears to outperform multiple competing structural interpretations of the business-cycle data, such as those based on investment-specific or consumption-specific shocks, news or noise shocks, and certain kinds of financial or uncertainty shocks.

- It operationalizes the notion of "aggregate demand fluctuations" without the assistance of nominal rigidity or any other deviation from the baseline RBC framework, and it replicates the empirical regularities that have historically been associated with that notion.

- It offers a structural interpretation of the observed recessions that is not vastly different as one varies the presumed degree of nominal rigidity or other model ingredient such as the adjustment costs to investment.

The first two points speak directly to the quantitative value and the empirical performance of our approach. How much one appreciates the last two points depends on how strong priors one has about the importance of nominal rigidities, as well as how much useful one finds it to think about the notion of aggregate demand. Putting aside this notion altogether, however, two lessons remain. First, there are important empirical regularities—like the zero correlation between hours and labor productivity and the co-movement properties documented in Section 2 of our paper—that the RBC framework has hard time capturing and that our approach manages to capture in an arguably more

robust way than the NK framework. Second, given our uncertainty about what the "right" model is, we find it reassuring that our approach offers a structural interpretation of the business cycle that is robust as one moves across RBC and NK models.

In closing, we would like to iterate the broader methodological contribution of our paper: we developed a flexible method for incorporating rich higher-order belief dynamics in a large class of DSGE models. This method entails both benefits and costs. The benefit is tractability. The cost is the risk of limited discipline: by abstracting from the details of how agents collect, digest, and exchange information, and by assuming heterogeneous priors, we risk allowing too much freedom in our modeling of higher-order beliefs dynamics. While we don't think that this is a serious caveat for the particular exercises we conducted in this paper, one must be mindful of this risk in general: be aware of the theorist bearing free parameters! We thus view our approach as a complement, rather than a substitute, of alternative approaches that build the dynamics of beliefs from specific information structures, such as the growing work on rational inattention.

Appendices

[warning: the appendices in this draft are still rough and incomplete; please check *http://economics.mit.edu/files/9755* for continuous updates]

A1: Data and Identification

In this appendix we describe the data and the identification we use in Section 2.

Data. [to be added]

Identification. We have a VAR/VECM that admits the following representation:

$$A(L)Y_t = u_t \Longleftrightarrow Y_t = B(L)u_t$$

where Y_t is a N–dimensional random vector, $A(L)$ is a matrix polynomial, $E(u_t u_t') = \Sigma$ and $C(0) = I$. We assume that there exists a linear mapping between reduced form innovations u_t and some structural shocks ε_t such that

$$u_t = S\varepsilon_t$$

where we impose $\mathbb{E}(\varepsilon\varepsilon') = I$, such that the impact matrix S satisfies $SS' = \Sigma$.

The transformation matrix S can be rewritten as $S = \widetilde{S}Q$ where Q is an orthonormal matrix, $QQ' = I$ and $\widetilde{S}$ is an arbitrary orthogonalization matrix. In the sequel, we will assume that the $\widetilde{S}$ matrix is used to identify p permanent shocks, that come first in the identification matrix $\widetilde{S}$, and $N - p$ transitory shocks. This step will be referred to as the first stage in the sequel.

As shown by Beaudry et al. (2012), under these assumption, the structural $MA(\infty)$ representation of the VECM can be written as

$$Y_t = \sum_{\tau=0}^{\infty} R_\tau \varepsilon_{t-\tau},$$

where $R_\tau = C_\tau Q$ and $C_\tau = B_\tau \widetilde{S}$. Then the impulse response vector of variable $Y_{i,t}$ to the structural shock j at horizon τ that corresponds to the j-th columns of R_τ, denoted $r_{i,\tau}^{(j)}$, is given by

$$r_{i,\tau}^{(j)} \equiv C_{i,\tau} q^{(j)}$$

where $C_{i,\tau}$ is the i– row of C_τ and $q^{(j)}$ i the j-th column of Q.

Our aim is to identify the transitory shock that explains the most the business cycle at the 6 to 32 quarters frequencies. We start by computing the share of volatility of variable k at 6–32 quarters frequency explained by the transitory shocks only. The spectral density of variable k from the VAR is given by

$$\widetilde{F}_k(\omega) = \frac{1}{2\pi} C_k(e^{-i\omega})\overline{C_k(e^{-i\omega})}$$

where $C_k(z)$ is the k-th row of the polynomial matrix $C(z)$ and $\overline{C_k(e^{-i\omega})}$ denotes the complex conjugate transpose of $C_k(e^{-i\omega})$. The filtered spectral density $F_Y(\omega)$ is then given by

$$F_k(\omega) = f(\omega)^2 \widetilde{F}_k(\omega)$$

where $f(\omega)$ is given by

$$f(\omega) = \begin{cases} 1 & \text{if } \frac{2\pi}{32} \leqslant \omega \leqslant \frac{2\pi}{6} \\ 0 & \text{otherwise.} \end{cases}$$

The filtered volatility of variable k, σ^2_{kk} is then obtained by the inverse Fourier transform formula.

The volatility attributable to each transitory shock in the first stage decomposition, is computed in a similar way. We first compute the filtered spectral density

$$\widehat{F}_k(\omega) = \frac{f(\omega)^2}{2\pi} \overline{\widetilde{C}_k(e^{-i\omega})} \widetilde{C}_k(e^{-i\omega})$$

where $\widetilde{C}_k(e^{-i\omega})$ only contains the column of $C_k(e^{-i\omega})$ associated with the transitory shocks. Note that the complex conjugate transpose is now in first place in order to preserve the contribution of each shock. Using the inverse Fourier transform, we obtain the vector of the volatility of variable k attributable to each transitory shock, $\widehat{\sigma}^2_{kk}$. The variance decomposition for variable k is then given by the ratio $\theta_k = \widehat{\sigma}^2_{kk}/\sigma^2_{kk}$. This calculation can be done for m variables, the problem is then to find a vector x that solve the program

$$\max_x \sum_{k=1}^{m} \alpha_k \, x \theta_k x'$$
$$\text{s.t. } xx' = 1$$

where α_k is the weight assigned to variable k in the problem. This problem is a simple eigenvalue problem that can be solved on any numerical package.

Finally, the vector q is given by $(\mathbf{0}_p, x)'$ where $\mathbf{0}_p$ is a zero vector of length p — the number of permanent shocks.

A2: A Horserace with Different Structural Shocks

In this appendix we complement the analysis of the baseline model with two exercises. First, we explore the robustness of our findings as we move to the textbook NK model. Second, we run a horserace between our confidence shocks and other structural shocks considered in the literature.

All the models in the horserace feature a single structural shock in addition to the conventional TFP shock. What differs across the models is only the nature of the additional shock. In our model, it is a shock to higher-order beliefs. In the competing models, it is one of the following: (i) an investment-specific shock; (ii) a news shock; (iii) a shock to the consumer's discount factor; and (iv) a shock to monopoly markups.

Although other kinds of shocks can be found in the literature, we focus on the aforementioned ones for the following reasons. First, investment-specific technology shocks and news shocks are empirically plausible to start with, and have already been proposed in the literature as major sources of business-cycle volatility. Second, investment-specific shocks can also mimic the effects of certain types of financial shocks on firms, while discount-factor shocks have been used as proxies for financial shocks on consumers. Third, news shocks are direct competitors to our confidence shocks in the sense that they both operate by changing expectations. Finally, although exogenous markup shocks do not seem very plausible (at least to us), they are useful for comparison purposes: they have similar real effects as our confidence shocks, but different nominal implications.

To have a fair horserace, the calibration strategy is kept the same across all models: the preference and technology parameters are fixed as in Table 3; the technology shock is assumed to be a random walk; the other shock is assumed to be an AR(1) with persistence .89; and the standard deviations of the two shocks are calibrated so as to match the volatilities of hours and output in the model to the corresponding volatilities in the data.[26] Finally, for the NK versions of these models, the monetary parameters we choose are in line with textbook treatments of the NK model: the Calvo probability of resetting prices is set to 1/4, and monetary policy is assumed to follow a simple Taylor rule with weights 1.8 on inflation and 0.1 on the output gap.

Table 12 reports the business-cycle moments generated by the NK specifications of the five models in the horserace. The RBC specifications are covered in Table 5 in the main text. Figure 11 reports the associated IRFs, using solid lines for the RBC specifications and dashed lines for the NK specifications.

As evident from Tables 5 and 12, the only competing model that fares well in comparison to our model is the one with markup shocks: the two models deliver similar moments for all the key real variables. As discussed in the main text, this is because a reduction in monopoly power has the same effects on the firm's optimal demand of all inputs and its supply of output as an increase in the expected demand for its product. Yet, as evident in Table 12 and Figure 11, the two shocks have different nominal implications: the confidence shock tends to generate pro-cyclical inflation and pro-cyclical nominal interest rates, wheres the markup shock does the opposite. It follows that our model does better than the model with markup shocks on the nominal side of the data.[27]

Focusing now on the remaining three competing models, the most noticeable failure is that they do not generate positive co-movement between consumption and either hours and output or investment. Importantly, this failure obtains whether we look at the RBC or the NK case. Additional failures include an exaggeration of the predicted investment volatility in all three models, and the prediction of counter-cyclical investment in the case of discount-factor shocks.

[26]The only exemption to the last rule is in the case of news shocks. In this case, it is impossible to match both volatilities at the same time. We thus opt to match perfectly the volatility of output and minimize the mismatch in the volatility of hours.

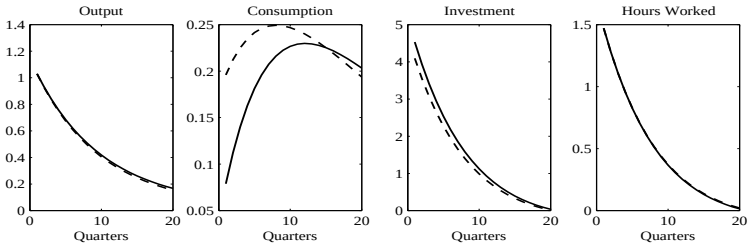
[27]Note, however, that the fit is in no way as good as the one the real side, indicating that other shocks or other mechanims are needed to capture the nominal side of the data.

Table 12: Band-pass Filtered Moments: NK Model

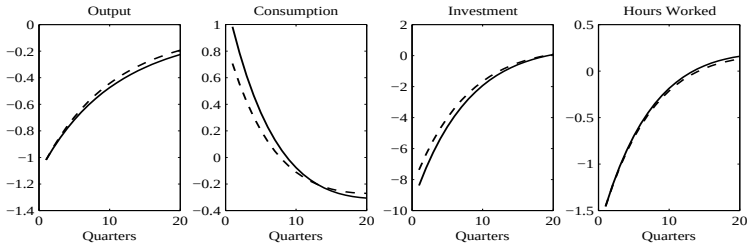
	Data	Conf.	Invt.	News	Disc.	Mkup
standard deviations						
σ_y	1.46	1.46	1.46	1.46	1.46	1.46
σ_h/σ_y	1.17	1.17	1.17	0.92	1.17	1.17
σ_c/σ_y	0.59	0.45	0.74	0.24	3.05	0.43
σ_i/σ_y	3.95	3.46	5.95	5.05	7.48	3.78
$\sigma_{y/h}/\sigma_y$	0.68	0.55	0.56	0.35	0.56	0.51
σ_{sr}/σ_y	0.86	0.58	0.61	0.63	0.59	0.63
correlations						
$\rho_{c,y}$	0.86	0.82	-0.24	-0.34	0.87	0.68
$\rho_{i,y}$	0.94	0.96	0.90	0.99	-0.69	0.95
$\rho_{h,y}$	0.88	0.89	0.87	0.99	0.87	0.89
$\rho_{y/h,y}$	0.52	0.24	0.32	0.86	0.26	0.46
$\rho_{y/h,h}$	0.07	-0.22	-0.17	0.78	-0.23	0.01
$\rho_{c,h}$	0.87	0.48	-0.68	-0.48	0.97	0.28
$\rho_{c,i}$	0.77	0.63	-0.64	-0.49	-0.95	0.43
nominal variables						
σ_π/σ_y	0.16	0.29	0.15	0.12	0.23	0.07
σ_R/σ_y	0.24	0.33	0.20	0.13	0.24	0.04
$\rho_{\pi,y}$	0.21	0.91	0.92	-0.34	0.92	-0.66
$\rho_{R,y}$	0.37	0.91	0.92	-0.26	0.92	-0.29
$\rho_{\pi,h}$	0.39	1.00	0.93	-0.41	0.99	-0.91
$\rho_{R,h}$	0.60	1.00	0.97	-0.32	0.99	-0.62

Figure 11: Horserace IRFs

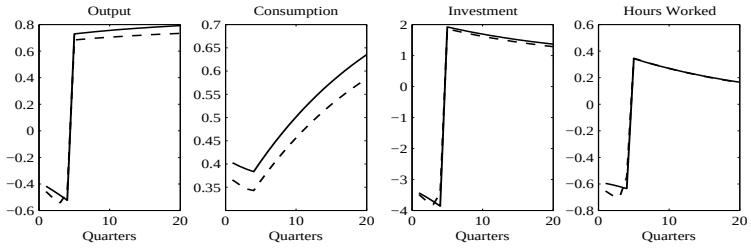
(a) Confidence shocks



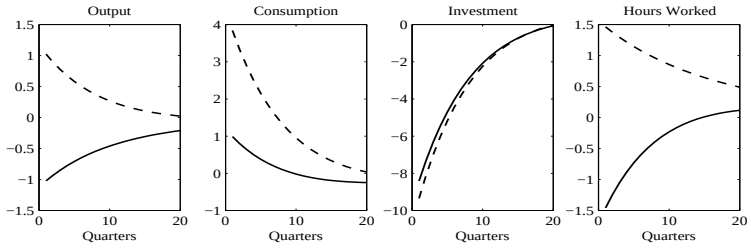
(b) Investment-Specific Shocks



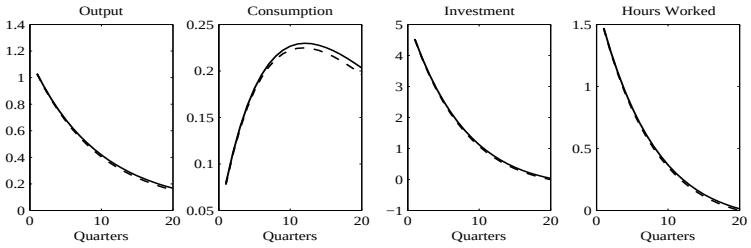
(c) News Shocks



(d) Discount-Factor Shocks



(e) Real Markup Shocks



Real Model ---- NK Model

All in all, the confidence shock therefore emerges as the winner of the horserace: it delivers a superior fit of key business-cycle moments, and it is also the only shock that mimics the empirical shock we document in Section 2.

A3: Estimated Models

In this appendix we expand on the estimation of the multi-shock RBC and NK models considered in Section 6. In particular, we first explain the estimation method; we next review the priors that we assume for the parameters and the posteriors that obtain from the estimation; we finally review a number of findings that are only briefly discussed in the main text.

Estimation. As mentioned in the main text, we estimate the model on the frequency domain, focusing on business-cycle frequencies. 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, next, 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 variable and the underlying state vector of the model; ε is the vector of the exogenous structural shocks, which is 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. The model spectral density of the the vector Y_t then writes

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

where $\omega \in [0, 2\pi]$ denotes the frequency at which the spectral density is evaluated. Then 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. Following Christiano and Vigfusson (2002) and Sala (2013), we include a weight γ_j in the computation of the likelihood in order to select the desirable frequencies: this weight is an indicator function that equals 1 when the frequency falls between 6 and 32 quarters, and 0 otherwise.

Parameters: priors and posteriors. Tables 13 and 14 report the priors and the posterior estimates for the parameters in 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 13: Estimated Parameters I

	Priors			Posteriors	
	Shape	Mean	Std. Dev.	Flexible Prices	Sticky Prices
Real Aspects					
ν	Gamma	0.500	0.250	0.647	0.793
				[0.212, 1.168]	[0.179, 1.518]
α	Beta	0.300	0.050	0.288	0.283
				[0.245, 0.341]	[0.234, 0.337]
χ	Beta	0.500	0.250	0.309	0.632
				[0.034, 0.674]	[0.277, 0.934]
ϱ	Beta	0.800	0.100	0.839	0.761
				[0.655, 0.959]	[0.508, 0.970]
ϕ	Gamma	0.300	0.200	0.278	0.119
				[0.096, 0.519]	[0.002, 0.479]
Nominal Aspects					
ρ_r	Beta	0.500	0.150		0.259
					[0.102, 0.415]
κ_π	Normal	1.700	0.300		1.304
					[1.063, 1.681]
κ_y	Normal	0.125	0.050		0.091
					[-0.016, 0.193]
ξ_π	Beta	0.660	0.100		0.672
					[0.534, 0.827]

Note: 95% HPDI into brackets. Note that the flexible price model is estimated in two steps. First we estimate the real side of the model using real data only. Then, conditionally on the estimated real side of the model, we estimate the parameters pertaining to the nominal side of the model so as to match the dynamics of inflation and the nominal interest rate.

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, σ_z, is Beta distributed around 0.5 with standard deviation 0.25.[28] The parameter governing the size of strategic complementarity, ϱ, is Beta distributed around 0.8

[28]This choice follows Smets and Wouters (2007) and nests inelastic as well as infinitely elastic utilization.

Table 14: Estimated Parameters II

	Priors			Posteriors	
	Shape	Mean	Std. Dev.	Flex. Price	Sticky Price
Shocks: persistence					
ρ_a	Beta	0.500	0.100	0.524	0.501
				[0.347, 0.711]	[0.305, 0.672]
ρ_i	Beta	0.500	0.100	0.463	0.466
				[0.271, 0.652]	[0.271, 0.662]
ρ_g	Beta	0.500	0.100	0.596	0.594
				[0.415, 0.768]	[0.396, 0.760]
ρ_m	Beta	0.500	0.100	–	0.346
					[0.208, 0.497]
ρ_ξ	Beta	0.500	0.100	0.629	0.796
				[0.429, 0.794]	[0.619, 0.915]
Shocks: volatilities					
$\sigma_{a,p}$	Inv. Gamma	1.000	∞	0.515	0.475
				[0.302, 0.761]	[0.305, 0.680]
$\sigma_{a,t}$	Inv. Gamma	1.000	∞	0.531	0.520
				[0.283, 0.827]	[0.294, 0.776]
$\sigma_{a,n}$	Inv. Gamma	1.000	∞	0.529	0.636
				[0.235, 0.976]	[0.312, 1.072]
$\sigma_{i,p}$	Inv. Gamma	1.000	∞	0.577	0.672
				[0.395, 0.782]	[0.479, 0.895]
$\sigma_{i,t}$	Inv. Gamma	1.000	∞	0.359	0.378
				[0.217, 0.553]	[0.216, 0.596]
$\sigma_{i,n}$	Inv. Gamma	1.000	∞	0.440	0.702
				[0.235, 0.789]	[0.385, 1.090]
σ_g	Inv. Gamma	1.000	∞	2.121	2.165
				[1.443, 2.955]	[1.523, 3.004]
σ_m	Inv. Gamma	1.000	∞	–	0.363
					[0.251, 0.501]
σ_ξ	Inv. Gamma	1.000	∞	1.387	2.694
				[0.758, 2.252]	[1.257, 5.069]

Note: 95% HPDI into brackets.

with standard deviation 0.1.[29] Finally, the parameter governing the size of investment adjustment costs, φ, is Gamma distributed around 0.3 with a standard deviation 0.2; the mean value of 0.3 departs from the much higher values used in the DSGE literature, but is in line with the micro-level estimates provided by Eberly, Rebelo and Vincent (2012); the relatively high standard deviation reflects our own uncertainty about this modeling feature, but also permit the estimation to accommodate the much higher point values required by the pertinent DSGE literature, in case this improves the empirical performance of the model by a significant margin. Finally, as far as the monetary side of the model is concerned, we impose the following priors, which are in line with the literature: the Calvo probability of not resetting prices is Beta distributed around 0.66 with standard deviation 0.1; the persistence parameter in the Taylor rule is Beta distributed around 0.5 with standard deviation 0.15; the reaction coefficient on inflation is Normally distributed around 1.7 with standard deviation 0.3; and the reaction coefficient on the output gap is also Normally distributed with mean 0.125 and standard deviation 0.05.

IRFs and Covariance Decompositions. Figures 12 through 14 illustrate the IRFs of our estimated models with respect to the various structural shocks. (The solid lines correspond to the flexible-price model; the dashed ones to the sticky-price model.) Tables 15 and 16 present the co-variance decompositions of the business-cycle volatiles in our estimated models.

Table 15: Covariance Decomposition (6–32 Quarters): Flexible-Price model

	$\text{Cov}(Y,h)$	$\text{Cov}(Y,I)$	$\text{Cov}(Y,C)$	$\text{Cov}(h,I)$	$\text{Cov}(h,I)$	$\text{Cov}(I,C)$
Belief Shock	58.04	41.00	49.00	53.19	139.81	206.99
Technology shocks						
Permanent	4.80	5.85	55.05	1.45	27.61	79.36
Transitory	12.85	22.69	31.40	11.88	29.93	128.03
News	13.73	19.09	-14.59	18.89	-39.75	-150.66
Investment shocks						
Permanent	1.42	0.65	0.85	0.79	2.52	1.02
Transitory	0.80	1.18	-2.04	1.53	-5.80	-19.92
News	7.50	10.56	-17.56	13.55	-48.95	-161.22
Fiscal Shock	0.85	-1.02	-2.11	-1.28	-5.36	16.41

[29] As in our calibrated version, this parameter is not tightly identified: varying this parameter changes the estimates of other parameters, most notably that of σ_ξ, without affecting the empirical fit of the model.

Figure 12: Theoretical IRFs to three types of neutral technology shocks

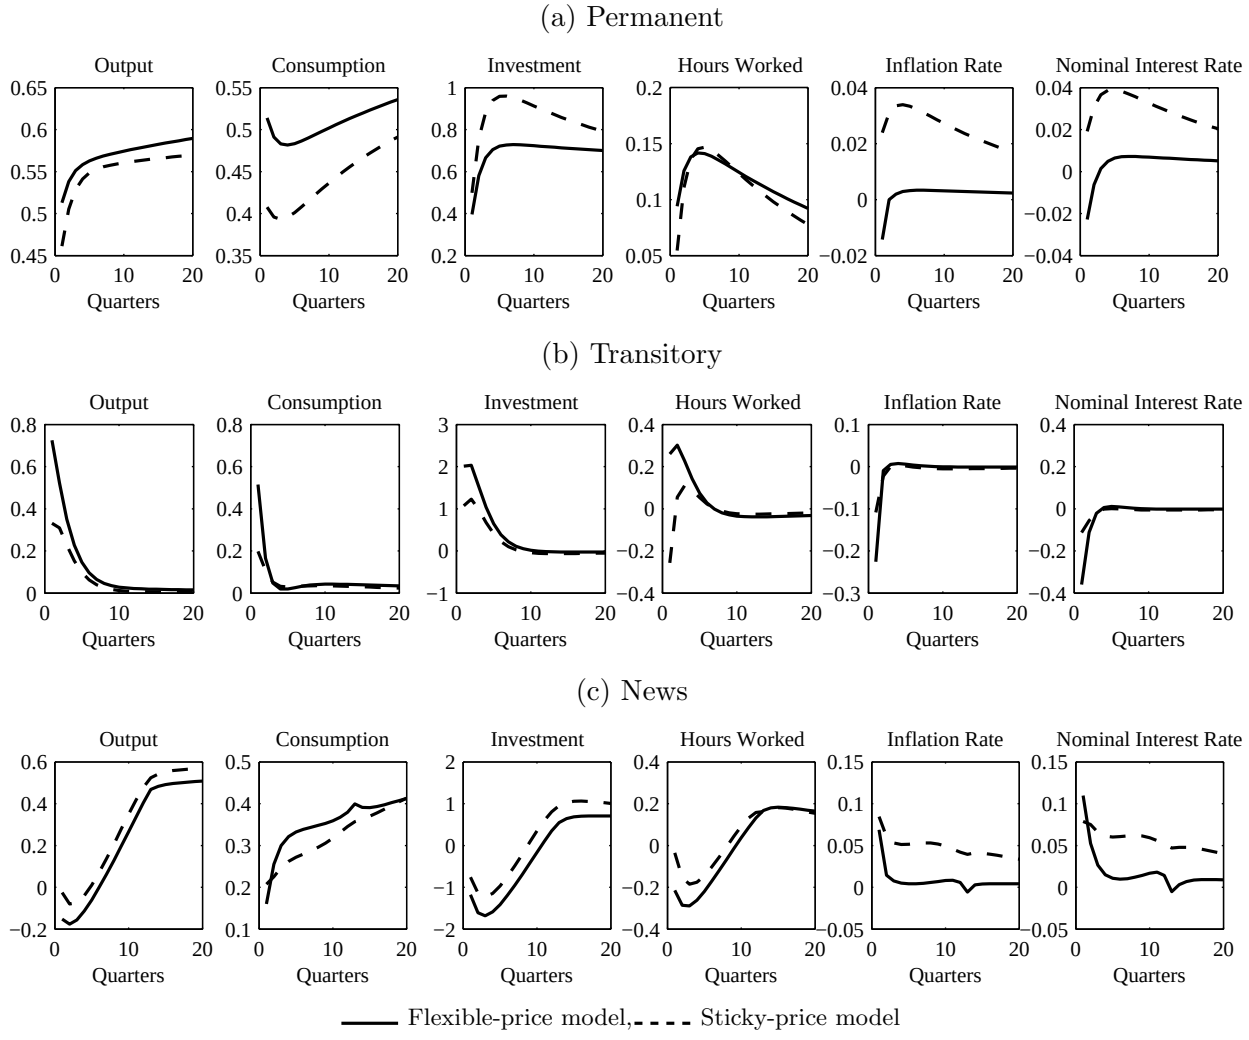
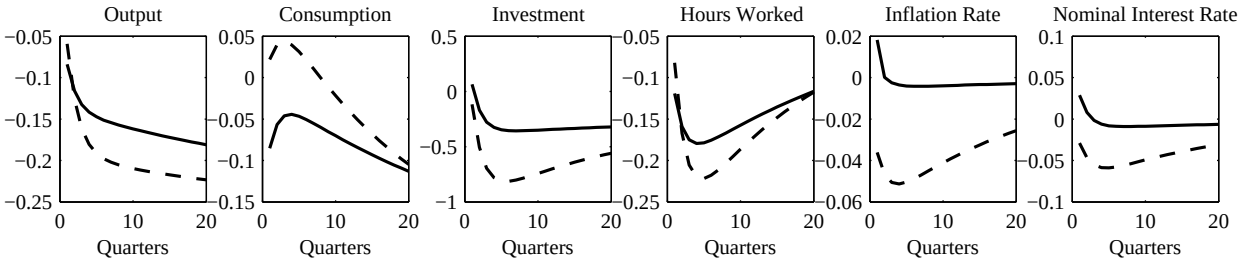
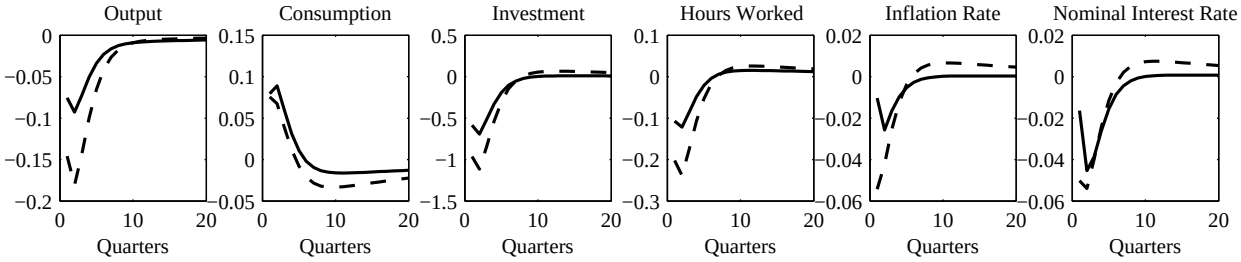


Figure 13: Theoretical IRFs to three types of investment shocks

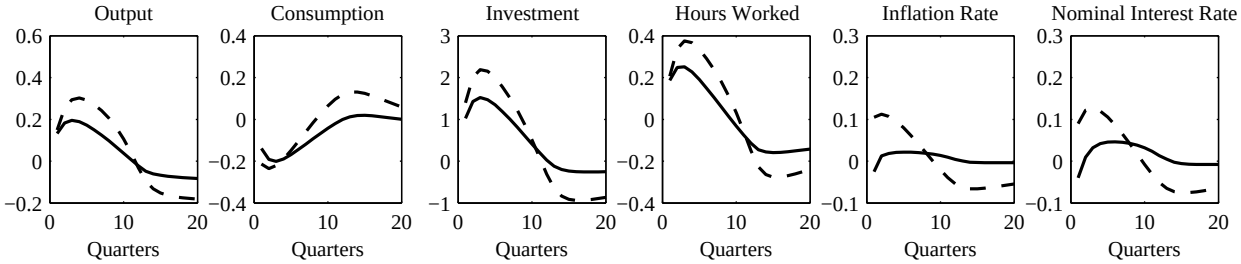
(a) Permanent



(b) Transitory



(c) News



Flexible-price model, Sticky-price model

Figure 14: Theoretical IRFs to two types of policy shocks

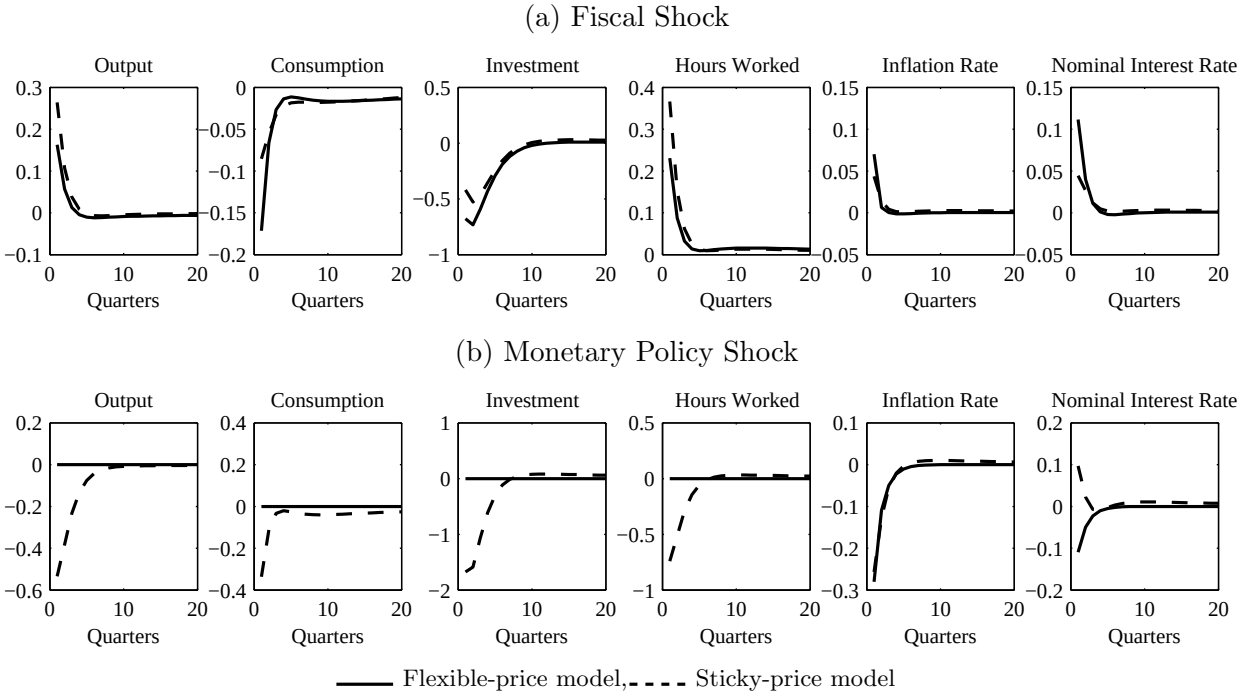


Table 16: Covariance Decomposition (6–32 Quarters): Sticky Price model

	$\text{Cov}(Y,h)$	$\text{Cov}(Y,I)$	$\text{Cov}(Y,C)$	$\text{Cov}(h,I)$	$\text{Cov}(h,I)$	$\text{Cov}(I,C)$
Belief Shock	65.66	54.21	85.85	60.91	161.85	305.64
Technology shocks						
Permanent	2.67	4.59	43.54	1.09	15.62	65.30
Transitory	0.26	4.71	8.36	0.44	-1.71	28.76
News	4.83	8.34	-8.78	6.16	-15.56	-70.12
Investment shocks						
Permanent	1.32	1.31	-2.13	1.43	-4.48	-9.57
Transitory	1.82	2.35	-3.74	2.65	-7.41	-22.51
News	12.19	17.47	-36.63	19.40	-74.29	-255.42
Fiscal Shock	1.36	-0.74	-2.04	-0.81	-3.72	6.23
Monetary Shock	9.88	7.75	15.57	8.73	29.71	51.70

A4: 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.

The log-linear version of our baseline RBC model. Before we consider the general case, it is useful to review the linearized version of our baseline RBC model. This will help fix ideas.

Thus consider the FOCs of the local planning problem and log-linearize them around the deterministic steady state. Let a hat denote the log-deviation of a variable from its steady-state value.[30] The log-linearized equilibrium conditions take the following form:

$$\widehat{y}_{i,t} = \widehat{a}_t + \alpha \widehat{k}_{i,t} + (1-\alpha)\widehat{n}_{i,t} \tag{9}$$

$$(1+\nu)\widehat{n}_{i,t} = (1-\varrho)\widehat{y}_{i,t} + \varrho \mathbb{E}_{i,t}\widehat{Y}_t - \mathbb{E}_{i,t}\widehat{c}_{i,t} \tag{10}$$

$$\widehat{k}_{i,t+1} = \delta \widehat{i}_{i,t} + (1-\delta)\widehat{k}_{i,t} \tag{11}$$

$$\mathbb{E}'_{i,t}\widehat{c}_{i,t+1} - \widehat{c}_{i,t} = (1-\beta(1-\delta))\mathbb{E}'_{i,t}\left[(1-\varrho)\widehat{y}_{i,t+1} + \varrho\widehat{Y}_{t+1} - \widehat{k}_{i,t+1}\right] \tag{12}$$

$$(1-\varrho)\widehat{y}_{i,t} + \rho\widehat{Y}_t = (1-s)\widehat{c}_{i,t} + s\widehat{i}_{i,t} + \alpha(1-\rho)\widehat{z}_{i,t} \tag{13}$$

The interpretation of these equations should be familiar: (9) is the local production function; (10) is the optimality condition for labor; (11) is the law of motion for local capital; (12) is the Euler condition; and (13) is the local resource constraint, with s denoting the saving rate in steady state. The only novelty in the above equilibrium conditions is the presence of island-specific expectation operators. In particular, $\mathbb{E}_{i,t}$ denotes the local expectation in stage 1 of period t, whereas $\mathbb{E}'_{i,t}$ denotes the updated expectation in stage 2. The difference between these two expectation operators derives from the fact that islands do not know Y_t in stage 1 but observe it in stage 2. This explains why the first expectation shows up in the optimality condition for labor, while the second shows up in the optimality conditions for consumption and investment.

The general case. We henceforth consider an economy that can be represented by the following system:

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

[30] To be precise, since we have allowed TFP to have a drift $\gamma > 0$, we consider a log-linearization around the associated deterministic balance-growth path.

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 captures backward-looking variables (such as capital in our model); x^f is captures forward-looking variables that are chosen in stage 2 of each period (such as consumption and investment in our model); y captures the variables that are instead chosen in stage 1 (such as employment in our model); s captures the shocks to payoff (such as technology); and finally ξ is meant to capture the shocks to higher-order beliefs. Accordingly, we make the following assumptions about beliefs, which mirror the assumptions we made in our baseline model:

$$\begin{aligned} \mathbb{E}_{it} s_t = s_t, \quad \mathbb{E}_{it}\mathbb{E}_{jt} s_t = s_t + \Delta\xi_t \\ \mathbb{E}'_{it} s_t = s_t, \quad \mathbb{E}'_{it}\mathbb{E}'_{jt} s_t = s_t \end{aligned}$$

where Δ is a loading matrix.

The precise interpretation of the aforementioned variables depends on the model under consideration. Similarly, the matrices that show up in the above system depend on the primitives of the model. In what follows, we show how to obtain a tractable solution for arbitrary such variables and matrices.

But let us first consider the complete-information, representative-agent counterpart of the general model we have introduced above:

$$\begin{aligned} Y_t &= M_X X_t + M_{EY} Y_t + M_F \mathbb{E}_t X_t^f + M_s \mathbb{E}_{it} s_t \\ X_{t+1} &= N_X X_t + N_Y Y_t + N_F X_t^f + N_s s_t \\ (P_{f0} - P_{F0})\mathbb{E}'_t X_{t+1}^f &= P_{F1} X_t^f + P_{Y0}\mathbb{E}'_{it} Y_{t+1} + P_X X_t + P_{Y1} Y_t + P_s s_t \\ s_t &= R s_{t-1} + \varepsilon_t \\ \xi_t &= Q\xi_{t-1} + \nu_t \end{aligned}$$

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:

$$X_t = \Gamma_X X_t^b + \Gamma_s s_t + \Gamma_\xi \xi_t,$$

where $X_t = (X_t^f, Y_t, X_{t+1}^b)$ collects all the variables, Γ_X and Γ_s are the same matrices as those that appear in the solution of the representative-agent model, and Γ_ξ is a new matrix, which identifies the effect of higher-order beliefs.

Solving the general case. We now show how to obtain the aforementioned solution. The basic strategy behind our solution is familiar: we start by making a guess for the policy rules followed by all islands expect one (equivalently, a guess for the aggregate outcomes); we next obtain the policy rules that represent the best response of the latter island; we finally look for the fixed point. The

only novelty, and complication, pertains in the treatment of expectations: we must keep track of the difference between first- and higher-order beliefs.

Preliminary steps and educated guesses. The system can be rewritten as

$$y_{it} = M_x(x^b_{it} - X_t) + M_X X_t + M_{EY}\mathbb{E}_{it}Y_t + M_f\mathbb{E}_{it}(x^f_{it} - X^f_t) + M_F X^f_t + M_s\mathbb{E}_{it}s_t \quad (14)$$

$$x^b_{it+1} = N_x(x^b_{it} - X_t) + N_X X_t + N_y(y_{it} - Y_t) + N_Y Y_t + N_f(x^f_{it} - X^f_t) + N_F X^f_t + N_s s_t \quad (15)$$

$$P_{f0}\mathbb{E}'_{it}x^f_{it+1} = P_{f1}(x^f_{it} - X^f_t) + P_{F0}\mathbb{E}'_{it}X^f_{t+1} + P_{F1}X^f_t + P_x(x^b_{it} - X_t) + P_X X_t + \quad (16)$$

$$+ P_{y0}(\mathbb{E}'_{it}y_{it+1} - \mathbb{E}'_{it}Y_{t+1}) + P_{Y0}\mathbb{E}'_{it}Y_{t+1} + P_{y1}(y_{it} - Y_t) + P_{Y1}Y_t + P_s s_t \quad (17)$$

where

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

Let us formulate the following **Dynamic Guess** for the x^f variables.

$$x^f_{it} = \Gamma^f_x(x^b_{it} - X_t) + \Gamma^f_X X_t + \Gamma^f_s\mathbb{E}_{it}s_t + \Gamma^f_\xi\xi_t$$

where we ignore the endowment shock as we already showed it will not be present in the solution.

Consider now the prior expectation of x^f_{it} and X^f_t. Using the above guess, we have

$$\begin{aligned}
X^f_t &= \Gamma^f_X X_t + \Gamma^f_s\mathbb{E}_{it}s_t + \Gamma^f_\xi\xi_t\\
\mathbb{E}_{it}X^f_t &= \Gamma^f_X X_t + \Gamma^f_s\mathbb{E}_{it}s_t + (\Gamma^f_\xi + \Gamma^f_s\Delta)\xi_t\\
\mathbb{E}_{it}x^f_{it} &= \Gamma^f_x(x^b_{it} - X_t) + \Gamma^f_X X_t + \Gamma^f_s\mathbb{E}_{it}s_t + \Gamma^f_\xi\xi_t
\end{aligned}$$

which also implies that

$$\mathbb{E}_{it}(x^f_{it} - X^f_t) = \Gamma^f_x(x^b_{it} - X_t) - \Gamma^f_s\Delta\xi_t$$

Plugging this in (14), the equation for y_{it}, we get

$$\begin{aligned}
y_{it} &= M_x(x^b_{it} - X_t) + M_X X_t + M_{EY}\mathbb{E}_{it}Y_t + M_f\mathbb{E}_{it}(x^f_{it} - X^f_t) + M_F\mathbb{E}_{it}X^f_t + M_s\mathbb{E}_{it}s_t \\
&= M_x(x^b_{it} - X_t) + M_X X_t + M_{EY}\mathbb{E}_{it}Y_t + M_f\left[\Gamma^f_x(x^b_{it} - X_t) - \Gamma^f_s\Delta\xi_t\right] \\
&+ M_F\left[\Gamma^f_X X_t + \Gamma^f_s\mathbb{E}_{it}s_t + (\Gamma^f_\xi + \Gamma^f_s\Delta)\xi_t\right] + M_s\mathbb{E}_{it}s_t \\
&= (M_x + M_f\Gamma^f_x)(x^b_{it} - X_t) + (M_X + M_F\Gamma^f_X)X_t + M_{EY}\mathbb{E}_{it}Y_t \\
&+ (M_s + M_F\Gamma^f_s)\mathbb{E}_{it}s_t + \left(M_F(\Gamma^f_\xi + \Gamma^f_s\Delta) - M_f\Gamma^f_s\Delta\right)\xi_t
\end{aligned}$$

Let us then make the following **Static Guess** for the y variables:

$$y_{it} = \Lambda^y_x(x^b_{it} - X_t) + \Lambda^y_X X_t + \Lambda^y_s\mathbb{E}_{it}s_t + \Lambda^y_\xi\xi_t$$

We then have, by aggregation,

$$\begin{aligned}
Y_t &= \Lambda^y_X X_t + \Lambda^y_s\bar{s}_t + \Lambda^y_\xi\xi_t \\
\mathbb{E}_{it}Y_t &= \Lambda^y_X X_t + \Lambda^y_s\mathbb{E}_{it}s_t + (\Lambda^y_\xi + \Lambda^y_s\Delta)\xi_t
\end{aligned}$$

Plugging these results in the equation we obtained for y_{it}, we get

$$\begin{aligned}
y_{it} &= (M_x + M_f\Gamma^f_x)(x^b_{it} - X_t) + (M_X + M_F\Gamma^f_X)X_t + M_{EY}\left(\Lambda^y_X X_t + \Lambda^y_s\mathbb{E}_{it}s_t + (\Lambda^y_\xi + \Lambda^y_s\Delta)\xi_t\right) \\
&+ (M_s + M_F\Gamma^f_s)\mathbb{E}_{it}s_t + \left(M_F(\Gamma^f_\xi + \Gamma^f_s\Delta) - M_f\Gamma^f_s\Delta\right)\xi_t
\end{aligned}$$

For this to be consistent with our static guess for y_{it}, we must have

$$\begin{aligned}
\Lambda^y_x &= M_x + M_f\Gamma^f_x \\
\Lambda^y_X &= (I - M_{EY})^{-1}(M_X + M_F\Gamma^f_X) \\
\Lambda^y_s &= (I - M_{EY})^{-1}(M_s + M_F\Gamma^f_s) \\
\Lambda^y_\xi &= (I - M_{EY})^{-1}(M_F\Gamma^f_\xi + (M_F - M_f)\Gamma^f_s\Delta + M_{EY}\Lambda^y_s\Delta)
\end{aligned}$$

We have thus verified the static guess and we have obtained the Ψ and Λ coefficients as functions of primitives and the Γ coefficients in our dynamic guess.

Beliefs about the future. The agent's beliefs about aggregate next-period output:

$$\begin{aligned}
Y_{t+1} &= \Lambda^y_X X_{t+1} + \Lambda^y_s\bar{s}_{t+1} + \Lambda^y_\xi\xi_{t+1} \\
\mathbb{E}_{it+1}Y_{t+1} &= \Lambda^y_X X_{t+1} + \Lambda^y_s\mathbb{E}_{it+1}s_{t+1} + (\Lambda^y_\xi + \Lambda^y_s\Delta)\xi_{t+1} \\
\mathbb{E}'_{it}Y_{t+1} &= \Lambda^y_X X_{t+1} + \Lambda^y_s Rs_t + (\Lambda^y_\xi + \Lambda^y_s\Delta)Q\xi_t
\end{aligned}$$

His beliefs about his own next-period output:

$$\begin{aligned}
y_{it+1} &= \Lambda^y_x(x^b_{it+1} - X_{t+1}) + \Lambda^y_X X_{t+1} + \Lambda^y_s\mathbb{E}_{it+1}s_{t+1} + \Lambda^y_\xi\xi_{t+1} \\
\mathbb{E}'_{it}y_{it+1} &= \Lambda^y_x(x^b_{it+1} - X_{t+1}) + \Lambda^y_X X_{t+1} + \Lambda^y_s Rs_t + \Lambda^y_\xi Q\xi_t
\end{aligned}$$

His beliefs about his next-period forward variable:

$$\mathbb{E}_{it+1}x^f_{it+1} = \Gamma^f_x(x^b_{it+1} - X_{t+1}) + \Gamma^f_X X_{t+1} + \Gamma^f_s \mathbb{E}_{it+1}s_{t+1} + \Gamma^f_\xi \xi_{t+1}$$
$$\mathbb{E}'_{it}x^f_{it+1} = \Gamma^f_x(x^b_{it+1} - X_{t+1}) + \Gamma^f_X X_{t+1} + \Gamma^f_s R s_t + \Gamma^f_\xi Q \xi_t$$

and for the aggregate forward variables

$$\mathbb{E}_{it+1}X^f_{t+1} = \Gamma^f_X X_{t+1} + \Gamma^f_s \mathbb{E}_{it+1}s_{t+1} + (\Gamma^f_\xi + \Gamma^f_s \Delta)\xi_{t+1}$$
$$\mathbb{E}'_{it}X^f_{t+1} = \Gamma^f_X X_{t+1} + \Gamma^f_s R s_t + (\Gamma^f_\xi + \Gamma^f_s \Delta) Q \xi_t$$

From these, we get

$$y_{it} - Y_t = \Lambda^y_x(x^b_{it} - X_t) + \Lambda^y_s(\mathbb{E}_{it}s_t - \bar{s}_t) = \Lambda^y_x(x^b_{it} - X_t)$$
$$x^f_{it} - X^f_t = \Gamma^f_x(x^b_{it} - X_t) + \Gamma^f_s(\mathbb{E}_{it}s_t - \bar{s}_t) = \Gamma^f_x(x^b_{it} - X_t)$$

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

$$\begin{aligned} x^b_{it+1} &= N_x(x^b_{it} - X_t) + N_X X_t + N_y(y_{it} - Y_t) + N_Y Y_t + N_f(x^f_{it} - X^f_t) + N_F X^f_t + N_s s_t \\ &= N_x(x^b_{it} - X_t) + N_X X_t + N_y\left\{\Lambda^y_x(x^b_{it} - X_t)\right\} + N_Y\left\{\Lambda^y_X X_t + \Lambda^y_s \bar{s}_t + \Lambda^y_\xi \xi_t\right\} \\ &\quad + N_f\left\{\Gamma^f_x(x^b_{it} - X_t)\right\} + N_F\left\{\Gamma^f_X X_t + \Gamma^f_s \bar{s}_t + \Gamma^f_\xi \xi_t\right\} + N_s s_t \end{aligned}$$

Using the fact that $\bar{s}_t = s_t = \mathbb{E}'_{it}s_t$, we get

$$x^b_{it+1} = \Omega_x(x^b_{it} - X_t) + \Omega_X X_t + \Omega_s s_t + \Omega_\xi \xi_t$$

and hence

$$X_{t+1} = \Omega_X X_t + \Omega_s s_t + \Omega_\xi \xi_t$$

where

$$\begin{aligned} \Omega_x &= N_x + N_y\Lambda^y_x + N_f\Gamma^f_x \\ \Omega_X &= N_X + N_Y\Lambda^y_X + N_F\Gamma^f_X \\ \Omega_s &= N_s + N_Y\Lambda^y_s + N_F\Gamma^f_s \\ \Omega_\xi &= N_Y\Lambda^y_\xi + N_F\Gamma^f_\xi \end{aligned}$$

It follows that

$$\begin{aligned} \mathbb{E}'_{it}x^f_{it+1} &= \Gamma^f_x(x^b_{it+1} - X_{t+1}) + \Gamma^f_X X_{t+1} + \Gamma^f_s R s_t + \Gamma^f_\xi Q \xi_t \\ &= \Gamma^f_x \Omega_x(x^b_{it} - X_t) + \Gamma^f_X\left\{\Omega_X X_t + \Omega_s s_t + \Omega_\xi \xi_t\right\} + \Gamma^f_s R s_t + \Gamma^f_\xi Q \xi_t \end{aligned}$$

or equivalently

$$\mathbb{E}'_{it}x^f_{it+1} = \Phi_x(x^b_{it} - X_t) + \Phi_X X_t + \Phi_s s_t + \Phi_\xi \xi_t$$

where

$$\begin{aligned}
\Phi_x &= \Gamma_x^f \Omega_x = \Gamma_x^f \left\{ N_x + N_y \Lambda_x^y + N_f \Gamma_x^f \right\} \\
\Phi_X &= \Gamma_X^f \Omega_X = \Gamma_X^f \left\{ N_X + N_Y \Lambda_X^y + N_F \Gamma_X^f \right\} \\
\Phi_s &= \Gamma_s^f R + \Gamma_X^f \Omega_s = \Gamma_s^f R + \Gamma_X^f \left\{ N_s + N_Y \Lambda_s^y + N_F \Gamma_s^f \right\} \\
\Phi_\xi &= \Gamma_\xi^f Q + \Gamma_X^f \Omega_\xi = \Gamma_\xi^f Q + \Gamma_X^f \left\{ N_Y \Lambda_\xi^y + N_F \Gamma_\xi^f \right\}
\end{aligned}$$

With the above steps, we have calculated all the objects that enter the Euler condition (16). We can thus proceed to solve for the dynamic guess.

Dynamic fixed point. Recall that the Euler condition (14) is given by

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

Using our preceding results, we can rewrite this condition as follows:

$$\begin{aligned}
P_{f0} \left\{ \Phi_x(x^b_{it} - X_t) + \Phi_X X_t + \Phi_s s_t + \Phi_\xi \xi_t \right\} &= P_{f1} \Gamma_x^f (x^b_{it} - X_t) + P_{F0} \left\{ \Phi_X X_t + \Phi_s s_t + \Phi_\xi \xi_t \right\} \\
&+ P_{F1} \left\{ \Gamma_X^f X_t + \Gamma_s^f s_t + \Gamma_\xi^f \xi_t \right\} \\
&+ P_x(x^b_{it} - X_t) + P_X X_t + P_{y1} \Lambda_x^y (x^b_{it} - X_t) \\
&+ P_{Y1} \left\{ \Lambda_X^y X_t + \Lambda_s^y s_t + \Lambda_\xi^y \xi_t \right\} \\
&+ P_{y0} \left\{ \Lambda_x^y \Omega_x (x^b_{it} - X_t) - \Lambda_s^y \Delta Q \xi_t \right\} \\
&+ P_{Y0} \Lambda_X^y \left[\Omega_X X_t + \Omega_s s_t + \Omega_\xi \xi_t \right] \\
&+ P_{Y0} \left\{ \Lambda_s^y R s_t + (\Lambda_\xi^y + \Lambda_s^y \Delta) Q \xi_t \right\} + P_s s_t
\end{aligned}$$

For this to hold of all realizations, the following Γ coefficients must solve the following system:

$$\begin{aligned}
(P_{f0} - P_{F0}) \Phi_X &= P_{F1} \Gamma_X^f + P_X + P_{Y1} \Lambda_X^y + P_{Y0} \Lambda_X^y \Omega_X \\
(P_{f0} - P_{F0}) \Phi_s &= P_{F1} \Gamma_s^f + P_{Y1} \Lambda_s^y + P_{Y0} \left\{ \Lambda_X^y \Omega_s + \Lambda_s^y R \right\} + P_s \\
(P_{f0} - P_{F0}) \Phi_\xi &= P_{F1} \Gamma_\xi^f + P_{Y1} \Lambda_\xi^y + P_{Y0} \Lambda_\xi^y Q + P_{Y0} \Lambda_X^y \Omega_\xi + (P_{Y0} - P_{y0}) \Lambda_s^y \Delta Q
\end{aligned}$$

where the matrices Φ, Λ, Ω are themselves transformations of the Γ matrices.

One can readily verify that $\Gamma_e^f = 0$ solves the last equation. This means that we don't have to solve the model with "endowment shocks"

The key question then is if the top two questions are the ones from the standard representative-agent model. I bet they are.

Provided that this is true, we can proceed to solve for Γ_ξ as a function of the primitives and the representative-agent policy rules.

Solving for Γ_ξ. From the aforementioned system, the condition that pins down Γ_ξ is the following:

$$(P_{f0} - P_{F0})\,\Phi_\xi = P_{F1}\Gamma^f_\xi + P_{Y1}\Lambda^y_\xi + P_{Y0}\Lambda^y_\xi Q + P_{Y0}\Lambda^y_X \Omega_\xi + (P_{Y0} - P_{y0})\Lambda^y_s \Delta Q$$

Using the facts that

$$\begin{aligned}\Omega_\xi &= N_Y\Lambda^y_\xi + N_F\Gamma^f_\xi \\ \Phi_\xi &= \Gamma^f_\xi Q + \Gamma^f_X\left\{N_Y\Lambda^y_\xi + N_F\Gamma^f_\xi\right\}\end{aligned}$$

we can express the aforementioned condition as

$$\begin{aligned}(P_{f0} - P_{F0})\left\{\Gamma^f_\xi Q + \Gamma^f_X\left(N_Y\Lambda^y_\xi + N_F\Gamma^f_\xi\right)\right\} = P_{F1}\Gamma^f_\xi + P_{Y1}\Lambda^y_\xi + P_{Y0}\Lambda^y_\xi Q + P_{Y0}\Lambda^y_X\left(N_Y\Lambda^y_\xi + N_F\Gamma^f_\xi\right) \\ + (P_{Y0} - P_{y0})\Lambda^y_s\Delta Q\end{aligned}$$

or equivalently

$$\begin{aligned}\left((P_{f0} - P_{F0})\Gamma^f_\xi - P_{Y0}\Lambda^y_\xi\right)Q + \left((P_{f0} - P_{F0})\Gamma^f_X N_Y - P_{Y1} - P_{Y0}\Lambda^y_X N_Y\right)\Lambda^y_\xi \\ + \left((P_{f0} - P_{F0})\Gamma^f_X N_F - P_{F1} - P_{Y0}\Lambda^y_X N_F\right)\Gamma^f_\xi = (P_{Y0} - P_{y0})\Lambda^y_s\Delta Q\end{aligned}$$

Next, replacing

$$\Lambda^y_\xi = (I - M_{EY})^{-1}M_F\Gamma^f_\xi + (I - M_{EY})^{-1}\left[(M_F - M_f)\Gamma^f_s\Delta + M_{EY}\Lambda^y_s\Delta\right]$$

we obtain

$$A\Gamma^f_\xi Q + B\Gamma^f_\xi = C \tag{18}$$

with

$$\begin{aligned}A =& (P_{f0} - P_{F0}) - P_{Y0}(I - M_{EY})^{-1}M_F \\ B =& \left((P_{f0} - P_{F0})\Gamma^f_X N_Y - P_{Y1} - P_{Y0}\Lambda^y_X N_Y\right)(I - M_{EY})^{-1}M_F + (P_{f0} - P_{F0})\Gamma^f_X N_F - P_{F1} - P_{Y0}\Lambda^y_X N_F \\ C =& \left((P_{Y0} - P_{y0})\Lambda^y_s + P_{Y0}(I - M_{EY})^{-1}\left[(M_F - M_f)\Gamma^f_s + M_{EY}\Lambda^y_s\right]\right)\Delta Q \\ & - \left((P_{f0} - P_{F0})\Gamma^f_X N_Y - P_{Y1} - P_{Y0}\Lambda^y_X N_Y\right)(I - M_{EY})^{-1}\left[(M_F - M_f)\Gamma^f_s + M_{EY}\Lambda^y_s\right]\Delta\end{aligned}$$

with $\Lambda^y_X = (I - M_{EY})^{-1}(M_X + M_F\Gamma^f_X)$ and $\Lambda^y_s = (I - M_{EY})^{-1}(M_s + M_F\Gamma^f_s)$.

Note that A, B, C are determined by primitives plus the policy rules of the representative-agent model. This verifies that we can calculate the "distortion" matrix Γ_ξ on the basis of the solution of the representative agent model.

We conclude that the policy rule for our economy as follows:

$$X^f_t = \Gamma^f_X X_t + \Gamma^f_s s_t + \Gamma^f_\xi \xi_t,$$

where Γ^f_X and Γ^f_s are the same matrices as in the representative agent model and Γ^f_ξ is the solution to (18).

Computation of forecasts. We have

$$X_{t+1} = \Omega_X X_t + \Omega_s \overline{s}_t + \Omega_\xi \xi_t$$

We want to compute $\mathbb{E}_{it} X_{t+1}$. We know that $\mathbb{E}_{it} \overline{s}_t = \overline{s}_t + \Delta \xi_t$, such that

$$\mathbb{E}_{it} X_{t+1} = \Omega_X X_t + \Omega_s \overline{s}_t + (\Omega_\xi + \Omega_s \Delta) \xi_t$$

Likewise, we have

$$\mathbb{E}_{it} X_{t+2} = \Omega_X \mathbb{E}_{it} X_{t+1} + \Omega_s \mathbb{E}_{it} \overline{s}_{t+1} + (\Omega_\xi + \Omega_s \Delta) \mathbb{E}_{it} \xi_{t+1}$$

which, using the properties of (i) expectations and (ii) the shocks, leads to

$$\mathbb{E}_{it} X_{t+2} = \Omega_X \mathbb{E}_{it} X_{t+1} + \Omega_s R \overline{s}_t + (\Omega_\xi Q + \Omega_s \Delta Q) \xi_t$$

Plugging the expected value for X_{t+1}, we get

$$\mathbb{E}_{it} X_{t+2} = \Omega_X^2 X_t + (\Omega_X \Omega_s + \Omega_s R) \overline{s}_t + [\Omega_X (\Omega_\xi + \Omega_s \Delta) + (\Omega_\xi + \Omega_s \Delta) Q] \xi_t$$

Keeping iterating, we get that

$$\mathbb{E}_{it} X_{t+j} = \Omega_X^j X_t + \sum_{k=0}^{j-1} \Omega_X^{j-1-k} \Omega_s R^k \overline{s}_t + \sum_{k=0}^{j-1} \Omega_X^{j-1-k} (\Omega_\xi + \Omega_s \Delta) Q^k \xi_t$$

Defining the vector $Z_t = (X_t \ \overline{s}_t \ \xi_t)'$ and the matrix

$$M_Z = \begin{pmatrix} \Omega_X & \Omega_s & \Omega_\xi + \Omega_s \Delta \\ 0 & R & 0 \\ 0 & 0 & Q \end{pmatrix}$$

it can be verified, by iterating on the matrix M_Z that

$$\mathbb{E}_{it} X_{t+j} = \Xi M_Z^j Z_t$$

where $\Xi = (I \ 0 \ 0)'$ is a selection matrix that fetches element X in the vector Z.

Our problem is actually to compute $\mathbb{E}_{it} Y_{t+j}$:

$$\mathbb{E}_{it} Y_{t+j} = \Lambda_X^y \mathbb{E}_{it} X_{t+j} + \Lambda_s^y \mathbb{E}_{it} s_{t+j} + \Lambda_\xi^y \mathbb{E}_{it} \xi_{t+j}$$

We have the first term from the previous calculations, and we have

$$\Lambda_s^y \mathbb{E}_{it} s_{t+j} + \Lambda_\xi^y \mathbb{E}_{it} \xi_{t+j} = \Lambda_s^y R^j \overline{s}_t + (\Lambda_\xi^y + \Lambda_s^y \Delta) Q^j \xi_t$$

Then it comes immediately that

$$\mathbb{E}_{it} Y_{t+j} = \begin{pmatrix} \Lambda_X^y & \Lambda_s^y & \Lambda_s^y \Delta + \Lambda_\xi^y \end{pmatrix} M_Z^j Z_t = \widetilde{M}_{Y,j} Z_t$$

Likewise, we have

$$Y_{t+j} = M_Y M_X^j Z_t + M_Y \sum_{k=0}^{j-1} M_X^k M_E \varepsilon_{t+j-k} = M_{Y,j} Z_t + M_Y \sum_{k=0}^{j-1} M_X^k M_E \varepsilon_{t+j-k}$$

Hence

$$\eta_{t,t+j} = (M_{Y,j} - \widetilde{M}_{Y,j}) Z_t + M_Y \sum_{k=0}^{j-1} M_X^k M_E \varepsilon_{t+j-k}$$

It then comes that

$$\mathbb{V}(\eta_{t,t+j}) = (M_{Y,j} - \widetilde{M}_{Y,j}) \Sigma_{ZZ} (M_{Y,j} - \widetilde{M}_{Y,j})' + M_Y \left(\sum_{k=0}^{j-1} M_X^k \Sigma_{EE} M_X^{k'} \right) M_Y'$$

where $\Sigma_{EE} = M_E \Sigma M_E'$.

$$\mathbb{C}ov(\eta_{t,t+j}, Y_t) = (M_{Y,j} - \widetilde{M}_{Y,j}) \Sigma_{ZZ} M_Y'$$

Simple, but tedious, algebra shows that the first order autocovariance of $\eta_{t,t+j}$ is finally given by

$$\begin{aligned} \mathbb{C}ov(\eta_{t,t+j}, \eta_{t-1,t-1+j}) =& (M_{Y,j} - \widetilde{M}_{Y,j}) M_X \Sigma_{ZZ} (M_{Y,j} - \widetilde{M}_{Y,j})' + (M_{Y,j} - \widetilde{M}_{Y,j}) \Sigma_{EE} M_X^{j-1'} M_Y' \\ &+ M_Y M_X \left(\sum_{k=0}^{j-2} M_X^k \Sigma_{EE} M_X^{k'} \right) M_Y' \end{aligned}$$

References

[To be completed]