

NBER WORKING PAPER SERIES

QUANTIFYING CONFIDENCE

George-Marios Angeletos
Fabrice Collard
Harris Dellas

Working Paper 20807
<http://www.nber.org/papers/w20807>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2014

We are grateful to Stephen Morris, Kristopher Nimark, and Morten Ravn for their conference discussions of our paper as well as for their feedback. We have benefited from the comments of many colleagues; in seminars at Bocconi, BU, Harvard, LSE, MIT, UCL, the University of Zurich, the University of Oslo, the Norwegian Business School, and the Minneapolis FRB; and in presentations at the 2013 Hydra Workshop, CRETE 2013; the 2013 Bank of England/CFM conference on Macroeconomics; the 2014 LSE Conference in Honor of Robert Barro; the 2014 NBER Summer Institute; the 2014 Minnesota Workshop in Macroeconomic Theory; the 10th Dynare conference; the 2nd Rational Inattention Conference in Oxford; and the 11th DSGE Conference in Chicago. Angeletos would also like to express his gratitude to the University of Zurich, along with the UBS Center for Economics in Society, and the Excellence Foundation Zurich, for their hospitality and support during the 2013-14 academic year, during which a significant part of this project was conducted. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2014 by George-Marios Angeletos, Fabrice Collard, and Harris Dellas. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Quantifying Confidence

George-Marios Angeletos, Fabrice Collard, and Harris Dellas

NBER Working Paper No. 20807

December 2014

JEL No. E0

ABSTRACT

We enrich workhorse macroeconomic models with a mechanism that proxies strategic uncertainty and that manifests itself as waves of optimism and pessimism about the short-term economic outlook. We interpret this mechanism as variation in confidence and show that it helps account for many salient features of the data; it drives a significant fraction of the volatility in estimated models that allow for multiple structural shocks; it captures a type of fluctuations in aggregate demand that does not rest on nominal rigidities; and it calls into question existing interpretations of the observed recessions. We complement these findings with evidence that most of the business cycle in the data is captured by an empirical factor which is unlike certain structural forces that are popular in the literature but similar to the one we formalize here.

George-Marios Angeletos
Department of Economics, E18-222
MIT
77 Massachusetts Avenue
Cambridge, MA 02139
and NBER
angelet@mit.edu

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

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

1 Introduction

Recessions are often described as periods of "low confidence" and "weak aggregate demand" and, occasionally, as manifestations of "coordination failures". Yet, it is debatable what these notions mean or how to quantify them. In this paper, we develop a formalization of these notions that is grounded on the role of higher-order uncertainty in games but can be readily incorporated in workhorse macroeconomic models. We then use this to gauge the quantitative role of "confidence" and to revisit prevailing structural interpretations of the business-cycle data.

Theoretical background. Macroeconomists have long disagreed about the causes and the nature of economic fluctuations. Nonetheless, the models they use in order to formalize their preferred narratives, to interpret the data, and to guide policy tend to share one key feature: they represent the economy as essentially a complete-information game in which players attain common belief of one another's choices and, in this sense, face no frictions in coordinating their actions.

What explains the ubiquity of this modeling practice is its convenience rather than its plausibility. The notion that the millions of firms and households whose interactions determine macroeconomic outcomes can coordinate as flawlessly as presumed by the Nash equilibria of complete-information games is a gross abstraction. Macroeconomists have embraced this abstraction because it has allowed them to build a class of micro-founded, general-equilibrium models—commonly referred to as DSGE models—that are at once tractable, versatile, and quantitatively potent.

By contrast, game theorists have sought to operationalize the notion that coordination is imperfect—in their language, the notion of strategic uncertainty—by relaxing either the solution concept or the informational assumptions of any given model. Either way, a key underlying theme is that higher-order uncertainty can preclude the players of a game from reaching a common belief about one another's choices, thus also inhibiting the coordination of their actions.¹

This raises an intriguing possibility. At the aggregate level, higher-order uncertainty manifests itself as a distinct form of volatility: economic outcomes can vary around, and away from, those predicted by workhorse macroeconomic models, reflecting certain variation in the beliefs that firms and consumers form about one another's choices. To the extent that this volatility is quantitatively significant and generates realistic co-movement patterns in the key macroeconomic variables, it may call into question existing explanations of business-cycle phenomena and may lead to novel structural interpretations of the available data. The goal of this paper is to evaluate this possibility.

Empirical backdrop. We complement the preceding motivation with evidence that points towards a type of volatility in the data that differs from those formalized in standard macroeconomic models. We use a variant of dynamic factor analysis, similar to the one in Uhlig (2003), which aims at encapsulating the bulk of the observed business cycles in a single shock, or factor, without imposing a structural interpretation of it. This identifies some key regularities in the data that

¹Note that incomplete information inhibits the coordination of *actions* while maintaining the Nash or Rational-Expectations solution concepts, which themselves presume perfect coordination of *strategies*. A complementary approach relaxes the latter assumption by studying the set of rationalizable outcomes.

are at odds with many of the structural forces featured in workhorse macroeconomic models, a fact that raises doubts about existing formal interpretations of the business cycle. An integral part of our contribution is to show that a certain enrichment of beliefs in these models offers a potent structural interpretation of the identified regularities as well as of other salient features of the data.

Methodological contribution. The enriched beliefs proposed in this paper build, not only on insights from game theory, but also on a strand of the macroeconomics literature that highlights the role of higher-order beliefs. This literature goes back at least to Phelps (1971) and Townsend (1983) and has been revived recently by the influential contributions of Morris and Shin (2001, 2002) and Woodford (2002). The closest precursor to our paper is Angeletos and La’O (2013), who illustrate how higher-order uncertainty helps unique-equilibrium, rational-expectations models accommodate forces akin to “animal spirits” while maintaining the common-prior assumption.

Despite these advances, there has been limited progress on the quantitative front. This is because the introduction of higher-order uncertainty in dynamic general-equilibrium models raises two technical difficulties that hinder their solution and estimation: Kalman filtering, to deal with noisy learning; and large state spaces, to keep tract of the dynamics of higher-order beliefs.

The methodological contribution of our paper lies in the bypassing of these difficulties. This is achieved by relaxing the common-prior assumption: we develop a heterogeneous-prior specification of the information structure that allows us to engineer aggregate waves in higher-order beliefs while abstracting from noisy learning. Although this entails a systematic bias in beliefs, it can also be seen as a convenient proxy for the higher-order uncertainty that can have been sustained in common-prior settings at the expense of lower tractability. Our specification allows us to augment a large class of DSGE with a tractable form of the type of volatility alluded to earlier.

To illustrate, consider the prototypical RBC model. In that model, the equilibrium dynamics take the form $K_{t+1} = G(K_t, A_t)$, where K_t is the capital stock and A_t is the technology shock. In general, adding incomplete information to this model causes the dimension of its state space to explode. By contrast, our formulation guarantees a minimal increase in the state space: the dynamics are given by $K_{t+1} = G(K_t, A_t, \xi_t)$, where ξ_t is an exogenous variable that encapsulates aggregate variation in the beliefs that agents form about the beliefs and actions of others.

Applied contribution. In the class of RBC and NK models we study in this paper, the waves of optimism and pessimism that are triggered by transitory movements in ξ_t have three distinguishing characteristics: they introduce transient deviations from the predictions of the underlying common-knowledge versions of these models; they reflect variation in expectations of the actions of others; and they concern the short-term outlook of the economy. We interpret these waves as the manifestation of frictions in coordination and as variation in the agents’ confidence about the state of the economy. Our applied contribution rests on the characterization of the observable properties of these belief waves and on the assessment of their quantitative potential.

We start our quantitative investigation with a simple, calibrated, RBC prototype. The model features only two sources of volatility: the familiar technology shock; and a “confidence shock,”

that drives the aforementioned belief waves. In spite of its parsimony, the model has no difficulty in matching key business-cycle moments, including the strong co-movement of hours, output, investment, and consumption alongside the lack of such co-movement in labor productivity. Importantly, the observable properties of the confidence shock mirror those of the factor that, accordingly to our empirical analysis, appears to account for the bulk of the business cycle in the data.

This success is not trivial: the ability of our mechanism to capture salient co-movement patterns in the data is not shared by a variety of other structural shocks that have been deployed in the literature as proxies for shifts in either “supply” or “demand”, including investment-specific shocks, discount-rate shocks, news or noise shocks, and uncertainty shocks.

To elaborate on what drives the distinct observable implications of our mechanism, consider a negative innovation in ξ_t . This triggers pessimism about the short-term economic outlook, without affecting the medium- to long-term prospects. As firms expect the demand for their products to be low for the next few quarters, they find it optimal to lower their own demand for labor and capital. Households, on their part, experience a transitory fall in wages, capital returns, and overall income. Because this entails relatively weak wealth effects, they react by working less and by cutting down on both consumption and saving. All in all, variation in ξ_t therefore leads to the following observable properties: strong co-movement between employment, output, consumption, and investment at the business-cycle frequency, without commensurate movements in labor productivity, TFP, or the relative price of investment at *any* frequency. It is precisely these patterns that our empirical analysis distills from the data and that existing structural mechanisms have difficulty matching.

Apart from capturing salient features of the data, our mechanism also offers a formalization of the notion of fluctuations in “aggregate demand” that does not rest on either nominal rigidities or frictions in the conduct of monetary policy. Indeed, a drop in confidence has similar incentive effects on a firm’s hiring and investment decisions as a joint tax on labor and capital. Seen through the lenses of the NK framework, a drop in confidence may therefore register as an increase in measured markups and in the “output gap”. Yet, there are no nominal rigidities at work, no constraints on monetary policy, and no commensurate drop in prices.²

Notwithstanding all these points, one may wonder how potent our mechanism is once it is embedded in richer DSGE models that allow multiple structural forces to account for the variability in the data. In order to address this issue, we study two “medium-scale” DSGE models—one with flexible prices, the other with sticky prices—which we estimate on US data. In both models, our mechanism has to compete with a large set of alternative structural forces, including TFP shocks, news shocks, investment-specific and consumption-specific shocks, and fiscal and monetary shocks.

²This begs the question of applying our approach to the recent recession. Is the slow recovery the product of persistent credit frictions, nominal rigidities, and binding constraints on monetary policy? Or is it the symptom of coordination failure and lack of confidence? Because the models we study in this paper are not equipped to deal with the kind of financial frictions that appear to have played a crucial role in this recession, we leave the investigation of the aforementioned question to future research. We nevertheless wish to note that our methodology facilitates such an investigation, and that our formalization of “demand” can help explain why the Great Recession was not accompanied by “great deflation”, as one would have expected on the basis of the baseline NK framework.

Despite the presence of these competing forces, our mechanism is estimated to account for about one half of GDP volatility at business-cycle frequencies (6-32 quarters). Furthermore, the observable properties of the confidence shock continue to match those of the factor that our empirical analysis identified as the main source of volatility in the data. They are also similar across the two models, underscoring the robustness of our mechanism across RBC and NK settings—a quality not shared by other structural mechanisms. Last but not least, the model-based confidence shock tracks well empirical counterparts such as the University of Michigan Index of Consumer Sentiment—a fact that lends support to our interpretation of this structural shock.

We conclude by discussing how our approach affects the interpretation of actual recessions. To account for them, estimated NK models rely not only on nominal rigidity but also on certain kinds of adjustment costs that have dubious micro-foundations. Furthermore, these models often predict that the observed recessions would have been *booms* had the nominal rigidities been absent. We find these properties problematic. By contrast, our approach leads to a structural interpretation of the business cycle that is not unduly sensitive to either the degree of nominal rigidity or the aforementioned modeling features. Importantly, downturns are now attributed to a *non-monetary* form of “deficient demand”, potentially calling into question prevailing policy prescriptions.

Layout. The rest of the paper is organized as follows. Section 2 discusses the related literature. Section 3 draws a useful empirical backdrop. Section 4 sets up the baseline model. Section 5 explains our solution method. Section 6 explores the quantitative performance of the baseline model. Section 7 extends the analysis to our two richer, estimated models. Section 8 discusses how our approach upsets existing interpretations of the observed recessions. Section 9 concludes.

2 Related literature

The idea of using heterogeneous priors as a modeling device for enriching higher-order beliefs is not totally new. Prior applications include Yildiz and Izmalkov (2009) and Section 7 of Angeletos and La’O (2013). Our contribution is found in the application of this idea towards the development of a general methodology for enriching the higher-order beliefs in a large class of DSGE models.

In so doing, our paper adds to a blossoming literature that studies the macroeconomic effects of informational frictions.³ Important attempts to push the quantitative frontier of this literature include Nimark (2013), who studies the asymptotic accuracy of finite-state-space approximations to a class of dynamic models with private information; Rodina and Walker (2013) and Huo and Takayama (2014), who obtain analytic solutions by studying the frequency domain of such models; Melosi (2014), who estimates a version of Woodford (2002); Mackowiak and Wiederholt (2014), who study the quantitative performance of a DSGE model augmented with rational inattention; and

³This includes, not only the aforementioned work on higher-order uncertainty, but also work on sticky information (Mankiw and Reis, 2001), rational inattention (Sims, 2001, Mackowiak and Wiederholt, 2009), and other forms of costly information acquisition (Alvarez, Lippi and Pacciello, 2011, Hellwig and Veldkamp, 2009; Pavan, 2013).

David, Hopenhayn, and Venkateswaran (2014), who use firm-level data to gauge the cross-sectional misallocation caused by informational frictions on the side of firms.

We view our methodological contribution as complementary to these attempts. The benefits concern the accommodation of a flexible form of strategic uncertainty and the straightforward applicability to DSGE models. The costs are that we abstract from the endogeneity of how agents collect, exchange, and digest information and we bypass the restrictions that the common-prior assumption, potentially in combination with micro-level data, can impose on the size and dynamics of higher-order uncertainty and thereby on the observables of the theory.

An ingenious method for characterizing the set of such restrictions that are robust across *all* information structures is provided by Bergemman and Morris (2013) and Bergemman, Heumann and Morris (2014). Unfortunately, their characterization only applies to static models, providing limited guidance for the kind of dynamic settings we are interested in.

By emphasizing the role of expectations, our paper connects to the literature on news and noise shocks that followed the influential contribution of Beaudry and Portier (2006); see, inter alia, Christiano et al (2008), Jaimovich and Rebelo (2009), Lorenzoni (2009), and Barsky and Sims (2012). Yet, there are important differences. This literature formalizes a form of optimism and pessimism that is tied to signals on future technology and that concerns the medium- to long-run macroeconomic outlook. By contrast, we introduce a form of optimism and pessimism that is disconnected from TFP and concerns the short-run outlook. As we explain in due course, this is key to the distinct predictions of our theory and its superior quantitative performance.

By emphasizing the role of coordination, our paper connects to the voluminous literature on multiple equilibria and sunspot fluctuations. The latter includes seminal contributions such as Benhabib and Farmer (1994), Cass and Shell (1983), and Diamond (1982), as well as more recent work such as Benhabib, Wang and Wen (2012) and Farmer (2012). One can view our contribution, along with that of Angeletos and La'O (2013), as a bridge that extends some of the *spirit* of that literature to the class of unique-equilibrium DSGE models that dominate modern research. In fact, because a unique equilibrium can be selected in *any* model under appropriate perturbations of the payoff and information structures (Weinstein and Yildiz, 2007, Penta, 2012), higher-order uncertainty is a useful tool for modeling and quantifying the macroeconomic effects of frictions in coordination irrespectively of the particular model one may have in mind.

At a broader level, our paper relates to work that relaxes the concept of rational-expectations, such as that on robustness and ambiguity (Hansen and Sargent, 2007, Ilut and Schneider, 2014), learning (Evans and Honkapohja, 2001, Eusepi and Preston, 2011), and educative stability (Guesnerie, 1992); see also Woodford (2013) and the references therein. Although the mechanism we study and the methods we develop are distinct, we share with that literature the desire to enrich the belief dynamics of macroeconomic models.

Putting aside all these methodological connections, what most distinguishes our paper from the pertinent literature is its applied contribution: the formalization and quantification of a novel type of structural volatility, which helps explain multiple salient features of the business-cycles data.

3 Empirical Backdrop

Our exploration of the data is guided by the following question. Suppose that a 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 are the empirical properties of such a shock?

To answer this question, we employ a variant of dynamic factor analysis (Sims and Sargent, 1977, Stock and Watson, 2005). We borrow from that method the core idea of using a small number of VAR-based shocks, or factors, to span most of the variation in data. Following an approach first proposed by Uhlig (2003) and further explored in a companion paper, we then focus on a factor that maximizes the volatility of particular variables at particular frequencies.

More specifically, we run a VAR/VECM on eight macroeconomic variables over the 1960-2007 period: GDP, consumption, hours, investment, the relative price of investment, inflation, the federal funds rate, and government spending. We allow for two unit-root components that drive the long-run movements in labor productivity (the output-to-hours ratio) and the price of investment. We then identify the sought-after factor by taking the linear combination of the VAR residuals that maximizes the sum of the volatilities of hours and investment at frequencies of 6 to 32 quarters.⁴

Our empirical strategy is therefore guided by two simple principles. On the one hand, we bypass the debatable identifying restrictions employed in the Structural VAR literature. On the other hand, we let our factor concentrate on spanning the business-cycle variation in two key variables of interest, namely employment and investment, as opposed to capturing, say, the low-frequency variation in labor productivity.

Figure 1 reports the impulse response functions (IRFs) of our factor (solid lines for the baseline specification, dashed for the alternative). Table 1 reports its variance contribution at different frequencies. Inspection of these results reveals the following patterns.

First, our identified factor captures one-half or more of the volatility of output, hours, and investment at business-cycle frequencies. It also gives rise to a clear business cycle, with consumption and the aforementioned variables all moving in tandem. This justifies why we can think of this factor as the “main business-cycle shock in the data”.

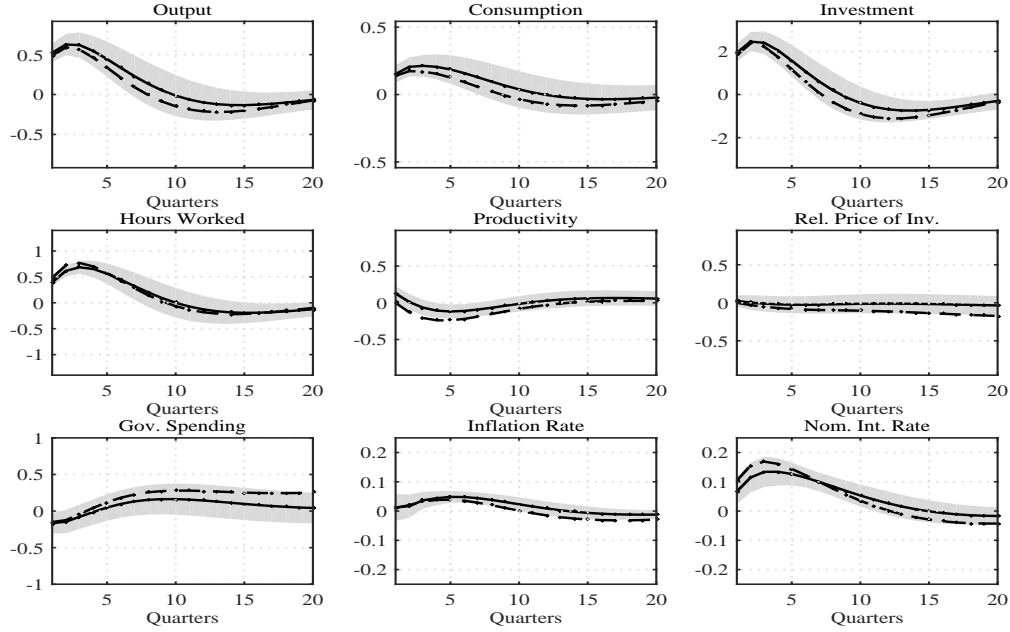
Second, the factor is associated with only mild movements in labor productivity and the price of investment at *all* frequencies. Furthermore, our companion paper obtains a similar picture when we add Fernald’s (2014) TFP measure: the factor is largely uncorrelated with TFP at all frequencies.

Finally, the factor is associated with modest pro-cyclical movements in inflation, and with pro-cyclical movements in both the nominal and the real interest rate.

The robustness of these findings and their connection to the pertinent empirical literature are explored in ongoing companion work. For the purposes of the present paper, we ask the reader to take these findings for granted. The DSGE estimations carried out in Section 7 are reassuring: they favor a structural shock that happens to be a look-alike of the factor we have documented here.

⁴See Appendix A for a description of the data and the empirical strategy. We drop the post-2007 data because we wish to abstract from the financial phenomena that have played a distinct role in the recent recession.

Figure 1: IRFs to Main Business-Cycle Factor (US Data)



Note: — Baseline specification, with permanent components excluded; ■ 68% HPDI; - - - Variant specification, with permanent components included.

Table 1: Variance Contribution of Main Business-Cycle Factor

	Y	I	h	C	Y/h	P_i	G	π	R
<i>Baseline specification, with permanent components excluded.</i>									
6–32 quarters	49.62	55.70	49.22	24.34	15.03	5.92	11.41	17.74	31.33
32–80 quarters	21.49	28.52	28.19	8.89	6.44	4.24	5.07	14.63	31.38
80–∞ quarters	0.00	0.00	3.37	0.00	0.00	0.00	0.00	2.08	5.15
<i>Variant specification, with permanent components included.</i>									
6–32 quarters	47.97	55.87	58.97	21.45	23.23	4.96	12.85	15.87	44.39
32–80 quarters	17.27	25.01	26.55	9.46	12.89	6.22	6.70	15.86	43.44
80–∞ quarters	6.67	6.67	7.26	6.66	6.67	6.62	6.97	6.68	9.52

Note: This table reports the contribution of our identified factor to the volatility of the variables at different frequencies. The definition of the factor is discussed in the main text.

Alternative ways of dissecting the data therefore appear to lead to essentially the same picture. In our view, this picture poses a challenge for the state of the art.

First of all, our identified factor is unlike the exogenous technology shocks, whether neutral or investment-specific, that have played a prominent role in RBC and NK models alike. It also does not fit in with models that emphasize either news or noise shocks: in these models, the business-cycle movements in output, employment, and investment are driven by signals of future technology and are therefore themselves correlated with movements in labor productivity and/or the price of investment at the same or lower frequencies—a prediction not supported by our evidence.

Our identified factor is also inconsistent with models that tie the business cycle to *endogenous* movements in TFP and labor productivity. This is true whether these movements are themselves triggered by sunspots, as in Benhabib and Farmer (1994); by preference shocks, as in Bai, Ríos-Rull, and Storesletten (2012); or by uncertainty shocks, as in Bloom (2009) and Bloom et al (2012).

Following the tradition of Blanchard and Quah (1989), it seems attractive to interpret our identified factor as a shock to “aggregate demand”. We are sympathetic to this notion, but we contend that a satisfactory formalization of it remains elusive. Monetary shocks do not square well with the counter-cyclical interest-rate movements seen in Figure 1. Furthermore, both monetary and fiscal shocks have too small quantitative effects to be able to explain the magnitudes seen in Table 1. Finally, other formalizations of aggregate-demand fluctuations within the NK framework, such as those based on consumption- and investment-specific shocks, have difficulty generating the strong co-movement between consumption, investment, and hours seen in Figure 1.⁵

The evidence presented here may not point to a unique “right” theory. First, the evidence contradicts the hypothesis that one of the aforementioned mechanisms explains the *bulk* of the observed business cycles, but leaves room for any of these mechanisms to play a non-trivial role. Second, the evidence could be consistent with a model in which *one* of the aforementioned structural shocks propagates in very different ways than those encapsulated in workhorse macroeconomic models. Finally, the evidence may also be consistent with a model where the joint contribution of *multiple* shocks happens to deliver the right observable patterns even though none of these shocks individually looks like our empirical factor. Notwithstanding these qualifications, we believe that the evidence presented in this section provides a useful gauge for the evaluation of competing theories of the business cycle. An integral part of our contribution is then to develop a structural interpretation of this evidence that is at once parsimonious and robust across baseline and richer, “medium-scale”, versions of workhorse macroeconomic models. What facilitates this is a particular enrichment of the belief structure of these models, which we describe next.

⁵We provide formal support to these claims in Appendix C by embedding examples of the aforementioned structural shocks to canonical RBC and NK models, and contrasting their IRFs to those seen in Figure 1. For further discussion, see Sections 6 and 7.

4 An RBC Prototype with a tractable form of higher-order beliefs

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

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

Time is discrete, indexed by $t \in \{0, 1, \dots\}$, and each period contains two stages. The labor and capital markets of each island operate in stage 1. At this point, the firm decides how much labor and capital to demand—and, symmetrically, the household decides how much of these inputs to supply—on the basis of incomplete information regarding the concurrent level of economic activity on other islands. In stage 2, the centralized markets for the intermediate and the final goods operate, the actual level of economic activity is publicly revealed, and the households make their consumption and saving decisions on the basis of this information.

Households. Consider the household on island i . 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 $\nu > 0$ is the inverse of the Frisch elasticity of labor supply. The household's budget constraint is $P_t c_{it} + P_t i_{it} = w_{it} n_{it} + r_{it} k_{it} + \pi_{it}$, where P_t is the price of the final good, i_{it} is investment, w_{it} is the local wage, r_{it} is the local rent on capital, and π_{it} is the profit of the local firm. Finally, the law of motion for capital is $k_{i,t+1} = (1 - \delta)k_{it} + i_{it}$, where $\delta \in (0, 1)$ is the depreciation rate.

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

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

where A_t is aggregate TFP, k_{it} is the local capital stock, and u_{it} is utilization. The latter is chosen in stage 2 and entails a cost equal to $\Psi(u_{it})k_{it}$ units of the final good. The firm's profit is $\pi_{it} = p_{it}y_{it} - w_{it}n_{it} - r_{it}k_{it} - P_t\Psi(u_{it})k_{it}$. We let $\Psi(u) = \psi_0 u^{\frac{1}{1-\psi}}$, with $\psi_0, \psi > 0$ and $\psi \in (0, 1)$.

Final-good sector. The final good is produced with a Cobb-Douglas technology. It follows that its quantity is given by $\log Y_t = \int_0^1 \log y_{it} di$ and the demand for the island goods satisfies

$$\frac{p_{it}}{P_t} = \frac{Y_t}{y_{it}}. \quad (1)$$

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

A tractable form of higher-order uncertainty. We open the door to strategic uncertainty by removing common knowledge of A_t in stage 1 of period t : each island i observes in that stage only a private signal of the form $x_{it} = \log A_t + \varepsilon_{it}$, where ε_{it} is an island-specific error. We then engineer the desired variation in higher-order beliefs by departing from the common-prior assumption and letting each island believe that the signals of others are biased: for every i , the prior of island i is that $\varepsilon_{it} \sim \mathcal{N}(0, \sigma^2)$ and that $\varepsilon_{jt} \sim \mathcal{N}(\xi_t, \sigma^2)$ for all $j \neq i$, where ξ_t is a random variable that becomes commonly known in stage 1 of period t and that represents the perceived bias in one another’s signals. These priors are commonly known: the agents “agree to disagree”. Finally, we assume that the actual signals are unbiased and focus on the limit case in which $\sigma = 0$.⁶

These modeling choices strike a balance between the need for tractability and the desire to enrich the dynamics of beliefs. Under a common prior, large and persistent divergences between first- and higher-order beliefs are possible only if the former are ridden with idiosyncratic noise and learning is imperfect. As a result, accommodating higher-order uncertainty in dynamic models typically comes at the cost of technical complications, including Kalman filtering and large state spaces. By contrast, our heterogeneous-prior specification allows for higher-order beliefs to diverge from first-order beliefs *even* when the noise in the latter vanishes. The variable ξ_t then permits us to engineer the desirable aggregate belief dynamics, while letting $\sigma \rightarrow 0$ guarantees that agents act in equilibrium as if they were perfectly informed about the state of the economy, thus abstracting from noisy learning and bypassing the aforementioned technical difficulties.

We close the model by assuming that ξ_t follows an $AR(1)$ process:

$$\xi_t = \rho \xi_{t-1} + \zeta_t,$$

where $\rho \in (0, 1)$ and ζ_t is drawn from a Normal distribution with mean 0 and variance σ_ξ^2 . This specification is in line with standard DSGE practice. More importantly for our purposes, it helps capture within our framework the basic idea that learning is likely to reduce strategic uncertainty over time—and therefore that the belief waves we are after are likely to be short-lived.

Remark 1. Although we model ξ_t as an exogenous shock, we think of it as proxying a mechanism whose deeper causes we abstract from for the sake of making progress in understanding its consequences. As already explained, this mechanism regards the uncertainty that firms and consumers face about one another’s choices, and thereby about the short-term outlook of the economy, as a result of frictions in coordination. In richer settings, the nature of these frictions and the resulting uncertainty can be endogenous to the network of markets or to other social interactions; to the details of how agents collect, digest, and exchange information; to a variety of forces that may shape expectations of one another’s actions. In this paper, we take for granted the variation

⁶We have in mind a sequence of games in which first- and higher-order beliefs converge to Dirac measures as $\sigma \rightarrow 0$. But instead of studying the equilibria of these games, we only study the equilibrium of the game with $\sigma = 0$.

in this kind of expectations, which we henceforth interpret as variation in the level of “confidence” about the state of the economy, and focus on gauging its quantitative implications. In so doing, we help formalize and quantify a certain kind of waves of optimism and pessimism that can be thought of as the symptom of frictions in coordination, but we remain agnostic about the deeper determinants and the triggers of these waves.

Remark 2. Our heterogeneous-prior formulation entails a systematic bias in the beliefs that agents form about the equilibrium impact of ξ_t on macroeconomic outcomes: although they predict correctly the *sign* of this impact, they tend to overestimate its *magnitude*.

We do not wish to take this property too seriously. Common-prior settings such as those studied in Angeletos and La’O (2013), Nimark (2013), Rodina and Walker (2013), and Huo and Takayama (2014) can accommodate similar waves in higher-order beliefs and in equilibrium outcomes as our approach, without invoking a systematic bias in beliefs. In effect, what is “bias” in our approach becomes “rational confusion” in those settings.

Having said that, our formulation does away with some of the restrictions that the common-prior assumption, perhaps in combination with appropriate data, can impose on the magnitude of the belief waves we are engineering. Accordingly, we also wish to invite the following broader, but complementary, interpretation of our modeling approach: we are departing from the standard rational-expectations solution concept in order to accommodate a certain kind of “mistakes” in the beliefs that agents form about one another’s actions, but do not necessarily wish to take a stand on whether these mistakes are “fully” rational.

These mistakes introduce stochastic deviations between the observable outcomes of our belief-augmented RBC prototype and the observable outcomes of the standard RBC model. Our analysis takes the existence of such deviations for granted and, in this sense, it only *assumes* the presence of an additional source of volatility. Nonetheless, unlike the case, say, of arbitrary trembles or measurement error, these deviations are not entirely free. Instead, they are disciplined by two requirements. First, all the additional variation in the observables of the model is spanned by the variation in beliefs, because there is no change in the payoff structure (preferences, technology, resource constraints, etc) of the model. And second, the beliefs agents form about all kinds of economic outcomes (output, employment, consumption, investment, wages, interest rates, etc) and at all horizons (current quarter, next quarter, 5 years later on, etc) are ultimately anchored to a particular form of higher-order uncertainty.

The fluctuations we engineer in this paper are therefore disciplined by certain “cross-equation restrictions”, which themselves manifest as certain co-movement patterns in the observables of the model. This property is akin to any other form of structural volatility: in the absence of direct measures of the underlying structural shocks, the “testability” of macroeconomic models rests on such cross-equation restrictions. Accordingly, the applied contribution of our paper rests on spelling out the observable patterns of the particular kind of belief-related volatility we have introduced, and on contrasting them with those of other structural shocks that are popular in the literature.

5 Equilibrium characterization and solution method

In this section we develop a recursive representation of the equilibrium. This serves two goals. First, it clarifies how the belief enrichment we propose in this paper enters the equilibrium determination and the restrictions this entails on the observables of the theory. Second, it illustrates the logic behind the solution method developed in Appendix E, which contains our broader methodological contribution and facilitates the quantitative evaluations in the rest of the paper. To simplify the exposition, we momentarily abstract from utilization and, without any loss, normalize $P_t = 1$.

Recursive equilibrium. Our formulation implies that an island’s hierarchy of beliefs is pinned down by two objects: the local signal, which itself coincides with A_t , the true TFP shock; and the perceived bias in the signals of others, which is given by ξ_t . This indicates that, relative to the standard RBC model, the state space of our model has been extended only by the ξ_t variable. We thus study a recursive equilibrium in which the aggregate state vector is given by (A, ξ, K) .

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

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

Note that this constraint depends on p_{it} , and thereby on the economy’s aggregate output, both of which are taken as given by the fictitious local planner. This dependence represents the type of aggregate-demand externalities that is at the core of many modern DSGE models.

To make his optimal decisions, the aforementioned planner must form beliefs about the evolution of p_{it} (or of Y_t) over time. These beliefs encapsulate the beliefs that the local firm forms about the evolution of the demand for its product and of the costs of its inputs, as well as the beliefs that the local consumer forms about the dynamics of local income, wages, and capital returns. The fact that all these kinds of beliefs are tied together underscores the cross-equation restrictions that discipline the belief enrichment we are considering in this paper: if expectations were “completely” irrational, the beliefs of different endogenous objects would not have to be tied together. The observable implications of this kind of restrictions, both in our baseline model and in richer DSGE models, are discussed below. For now, we emphasize that ξ_t matters for equilibrium outcomes because, and only because, it triggers co-movement in this kind of expectations.

Finally, note that beliefs of p_{it} (or of Y_t) are themselves tied to beliefs of how the aggregate economy evolves over time. In a recursive equilibrium, the latter kind of beliefs are encapsulated in the law of motion of the aggregate capital stock, the endogenous state variable. We thus define a recursive equilibrium as a collection of functions \mathbf{P} , \mathbf{G} , V_1 , and V_2 , such that the following is true:

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

- V_1 and V_2 solve the following Bellman equations:

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

$$\left. \begin{aligned} V_2(m; A, \xi, K) &= \max_{\{c, k'\}} \log c + \beta \int V_1(k'; A', \xi', K') df(A', \xi' | A, \xi) \\ \text{s.t. } c + k' &= m \\ K' &= \mathbf{G}(A, \xi, K) \end{aligned} \right\} \quad (4)$$

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

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

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

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

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

To interpret condition (5), recall that, in stage 1, each island believes that, with probability one, TFP satisfies $A = x$ and the signals of all other islands satisfy $x' = A + \xi = x + \xi$. Together with the fact that all islands make the same choices in equilibrium and that the function \mathbf{y} captures their equilibrium production choices, this implies that the local beliefs of local and aggregate output are

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

Summing up, an equilibrium is given by a fixed point between the Bellman equations (3)-(4) and the consistency conditions (5)-(6). In principle, one can obtain the global, non-linear solution of this fixed-point problem with numerical methods. As in most of the DSGE literature, however, we find it useful to concentrate on the log-linear approximation of the solution around the steady state. Once we do this, we obtain the equilibrium dynamics of our model as a tractable transformation of the equilibrium dynamics of the standard RBC model.

Log-linear solution. The linearized equilibrium dynamics of our model satisfy the following:

$$(\tilde{Y}_t, \tilde{N}_t, \tilde{I}_t, \tilde{C}_t; \tilde{K}_{t+1}) = \Gamma_K \tilde{K}_t + \Gamma_A \tilde{A}_t + \Gamma_\xi \xi_t \quad (7)$$

where the *tilde* symbol indicates the log-deviation of a variable from its steady-state value and where Γ_K , Γ_A and Γ_ξ are vectors that are pinned down by exogenous parameters and that regulate the model's impulse response functions. Importantly, Γ_K and Γ_A are the same as those in the standard RBC model, whereas Γ_ξ is obtained by solving an equation that contains Γ_K and Γ_A .

In Appendix E, we explain how the type of solution obtained in (7) generalizes to a large class of DSGE models. This facilitates the simulation, calibration, and estimation of the type of belief-augmented macroeconomic models we are interested in as in the case of workhorse DSGE models. The results developed in that appendix are therefore instrumental for the broader methodological contribution of our paper, as well as for the quantitative exercises we conduct next.

Remark. The fact that Γ_ξ solves an equation that itself depends on Γ_K and Γ_A underscores a more general principle: the dynamic effects of higher-order beliefs in any given model are tightly connected to the payoff-relevant shocks and the propagation mechanisms that are embedded in the complete-information version of that model. This in turn explains our choice to focus on higher-order uncertainty of TFP as opposed to higher-order uncertainty of, say, discount rates: the former has a better chance to generate realistic waves of optimism and pessimism within the RBC framework, because TFP shocks in the first place do a better job in generating realistic business cycles than discount-rate shocks. We would thus like to invite a judicious application of our methodology: whether higher-order uncertainty should be tied to one type of fundamental or another depends both on the phenomena the researcher is after and the particular model at hand.

6 Quantitative evaluation

Notwithstanding the theoretical motivation behind our paper, all that one ultimately sees in (7) is a DSGE model with two structural shocks. From this perspective, our applied contribution ultimately rests on assessing the observable properties of the confidence shock and on contrasting them with other structural shocks proposed in the literature. In this section, we take a first pass at this task by studying the empirical fit of a calibrated version of our baseline model.

Table 2: Parameters

Parameter	Role	Value
β	Discount Rate	0.99
ν	Inverse Elasticity of Labor Supply	0.50
α	Capital Share in Production	0.30
δ	Depreciation Rate	0.015
ψ	Inverse Elasticity of Utilization	0.30
σ_a	St.Dev. of Technology Shock	0.67
σ_ξ	St.Dev. of Confidence Shock	2.65
ρ	Persistence of Confidence Shock	0.75

Calibration. Table 2 reports the values we select for the parameters of the model. The preference and technology parameters of the model are set at conventional values: the discount factor is 0.99; the Frisch elasticity of labor supply is 2; the capital share in production is 0.3; the elasticity of utilization is 3; and the depreciation rate is 0.015. These values guarantee that our quantitative exercise is directly comparable to the literature as well as that the steady state values of our model are consistent with the long-run patterns in the data.

Since the technology shock follows a random walk, three parameters remain to be set in order to complete the parameterization of the model: the standard deviations of the two shocks and the persistence of the confidence shock. For the latter, we set $\rho = 0.75$. This choice is somewhat arbitrary, but it is motivated by the following considerations. In our setting, ρ pins down the persistence of the deviations between first- and higher-order beliefs. In common-prior settings, such deviations cannot last for ever, but can be quite persistent insofar as learning is slow. By setting $\rho = 0.75$, we assume that the half-life of these deviations is less than 2.5 quarters, which does not sound implausibly large. Furthermore, to the extent that ξ_t is a look-alike of the empirical factor documented in Section 3 (a property that remains to be shown), our parameterization of ρ is broadly consistent with the persistence of the fluctuations seen in Figure 3.

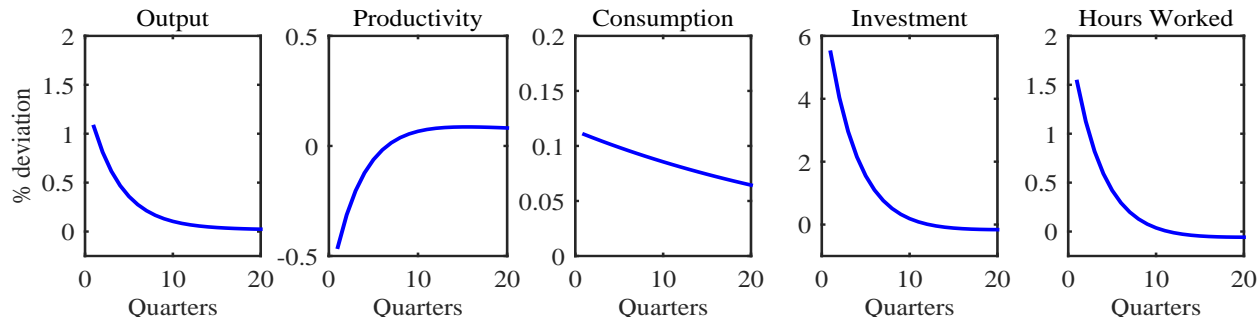
Turning to σ_a and σ_ξ , the standard deviations of the two shocks, we set them so as to minimize the distance between the volatilities of output, consumption, investment and hours found in the data and those generated by the model.⁷ This yields $\sigma_a = 0.67$ and $\sigma_\xi = 2.64$. Clearly, there is

⁷By “volatilities” we always refer to the Bandpass-filtered variances at frequencies corresponding to 6-32 quarters.

no compelling empirical justification for this parameterization. Furthermore, the relatively high value for σ_ξ begs the question of whether the belief fluctuations we accommodate in this paper are perhaps too large to be reconcilable with the common-prior assumption. We revisit this issue in Section 7 and expand on it in Appendix D. Notwithstanding these points, we believe that the chosen calibration strategy is useful because it facilitates the comparison of our mechanism to competing structural mechanisms in the literature. Furthermore, although the predicted moments of the model are sensitive to the chosen values for σ_a and σ_ξ , the shapes of the IRFs are entirely invariant to these values, and therefore the ability of our mechanism to generate the co-movement patterns that are seen in the data is also invariant to them.

The effects of the confidence shock. To reveal the observable properties of our mechanism, Figure 2 reports the IRFs of the model’s key variables to a positive innovation in ξ_t .⁸ As is evident from this figure, an increase in “confidence” is associated with 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.

Figure 2: Impulse responses to a positive confidence shock



These co-movement patterns encapsulate key testable predictions of our theory. Importantly, these patterns mirror the empirical regularities documented in Section 3 and are not shared by other structural mechanisms proposed in the literature. We next elaborate on the economic forces that shape the aforementioned IRFs and drive the distinct empirical implications of our theory.

In our setting, a positive innovation in ξ_t signals a *transitory* boom in aggregate output and, in this sense, captures optimism about aggregate demand in the short run. In response to this particular kind of optimism, firms find it optimal to raise their demand for both labor and capital, which in turn pushes up wages and capital returns. Other things being equal, this motivates the households to supply more labor as well as to invest more, because of the familiar static and intertemporal substitution effects. A countervailing wealth effect is also at work, because the household experiences a boom in its income. However, because the boom is transitory, the wealth

Also, in the minimization objective, each of the model-based volatilities is weighted by the precision of its estimator.

⁸The IRFs to the technology shock are the same as in the standard RBC model and are thus omitted. Also note that the shape of the IRFs, and therefore the co-movement regularities we document for the confidence shock do not depend on the values of σ_a and σ_ξ . These values matter only when we compute the model’s second moments.

effect is weak relative to the substitution effects, guaranteeing that households raise the labor supply and split their additional income between consumption and saving. All in all, our mechanism therefore induces employment, output, consumption and investment to move in tandem.

Narratives and mechanisms. The arguments above help explain, not only why the confidence shock is a look-alike of the empirical factor we documented in Section 3, but also why this success is not shared by other structural mechanisms that have been proposed in the literature in order to explain the data and support certain narratives of the observed business cycles.

Consider, in particular, the case of news and noise shocks. The literature has used such shocks to capture the idea that optimism about the future growth prospects of the economy could lead to a boom in the present. This idea, however, does not square well with the canonical RBC framework. In this framework, positive news or noise shocks raise the consumers’ expectations of “permanent income”. In response to this, consumers raise their demand, not only for goods, but also for leisure. In general equilibrium, this leads to a drop in hours, output, and investment, and therefore to *negative* co-movement between consumption and any of these variables.

Jaimovich and Rebelo (2006) seek to overturn this negative co-movement by considering two modifications of the baseline RBC model: a particular form of internal habit, which implies that the anticipation of higher consumption in the future *increases* the supply of labor today for any given wage; and an adjustment cost in investment, which makes the investment today increase in anticipation of higher investment in the future. The first feature, which amounts in effect to a perverse short-run income effect on labor supply, helps undo the negative response in employment, while the second features helps undo the negative response of investment. Lorenzoni (2009), on the other hand, proposes a resolution within the NK framework. The resolution relies on a monetary policy that “accommodates” consumer optimism, in the sense of letting the shift in expectations of permanent income induce pro-cyclical deviations from the underlying flexible-price allocations; and on shutting down, or dampening enough, the countervailing behavior of the latter.⁹

Whether one finds these alternative mechanisms to be plausible or not, they do not square well with the evidence in Section 3: the factor that spans most of the business-cycle movements in employment and investment in the data does not contain a signal of future productivity movements.¹⁰ Furthermore, our mechanisms has no difficult in generating realistic co-movement patterns within either RBC or NK settings for a very simple reason: it captures beliefs about aggregate demand in the short run, as opposed to beliefs about productivity and income in the medium to long run.

⁹Lorenzoni’s model assumes away capital, guaranteeing that employment and output move mechanically in the same direction as consumption. As illustrated in Appendix C, this is not an innocuous simplification: signals of future TFP cause consumption to move in the opposite direction from employment, output and investment in a calibrated NK model with capital. This is because the accommodating role of monetary policy is not sufficiently strong to offset the countervailing dynamics of the underlying flexible-price allocations once investment is free to adjust. Blanchard et al (2013) seek to fix this problem by adding significant adjustment costs to investment and assuming a sufficiently accommodative monetary-policy response. See Barsky and Sims (2011) for a complementary discussion.

¹⁰Complementary in this respect is the evidence in Barsky and Sims (2011) and Barsky, Basu, and Lee (2014) regarding identified news shocks.

Consider, next, the case of discount-rate shocks. Such shocks are often deployed in the literature in order to account for the impact of financial constraints on consumer spending. The storyline is as follows: tighter credit leads consumers to cut down on their spending, which reduces “aggregate demand” and triggers employment losses. We believe that neither the RBC nor the NK model can capture this storyline in a satisfactory way. In the RBC model, discount-rate shocks lead to opposite movements in investment and hours than in consumption: as resources are freed up by the drop in consumption, investment increases; and as higher discounting makes the consumers more eager to work, employment also goes up. This negative co-movement may be mitigated in NK settings via the combination of nominal rigidity, pro-cyclical output gaps, and certain forms of adjustment costs in investment and consumption. Nevertheless, as illustrated in Section 7, discount-rate shock are still unable either to generate realistic co-movement or to capture a significant fraction of the business-cycle volatility in an estimated NK that embeds these features. In contrast, our mechanism has no difficulty in generating realistic business-cycle patterns. It can therefore also help reconcile the aforementioned storyline with *either* the RBC or the NK framework insofar as the drop in consumer spending arises from, or comes with, a drop in “confidence”.¹¹

Finally, consider the recent work on uncertainty shocks. The class of linear models we study in this paper cannot accommodate this kind of shocks. Nevertheless, it is worth noting that existing formalizations of the macroeconomic effects of these shocks, such as those in Bloom (2009) and Bloom et al (2013), appear to hinge on strong pro-cyclical movements in aggregate TFP, a property that is not shared by our mechanism and that also seems at odds with the evidence of Section 3.

In short, the contribution of this paper rests, not on just telling a certain narrative about the business cycle, but rather on identifying a structural mechanism whose ability to capture salient features of the macroeconomic data does not appear to be shared by competing mechanisms. We provide additional support to this claim in the sequel, in Appendix C, and in Section 7.

Business-Cycle Moments. Complementing the preceding discussion, we now evaluate our model’s ability to match standard business-cycle moments. Table 3 reports some key moments in the data (column 1), in our model (column 2), and in four competing models (columns 3-6). Each of the competing models replaces the confidence shock with one of the following alternative structural shocks, which have been considered in the literature: a news shock; a discount-rate shock; an investment-specific shock; and a transitory TFP shock.¹²

¹¹A refined version of the aforementioned storyline is that financial frictions depress the demand, not only for consumption, but also for investment. This version, which seems plausible in the context of the Great Recession, is easier to reconcile with the NK model: a combination of a discount-rate and an investment-specific shock can generate a realistic recession, even if each one the shocks by itself would not. Note, however, that the recession would have to come with a commensurate deflation episode, a prediction that seems prima-facia inconsistent with the US experience. For a possible resolution that adds a countervailing inflationary shock, see Christiano et al (2014).

¹²To ensure a fair horserace, the parameterization of the competing models is in line with that of our model; see Appendix C for details. Also, the moments are computed on Bandpass-filtered series, at frequencies 6-32 quarters. This filter is preferable to the simpler HP filter because it removes not only low-frequency trends but also high-frequency “noise” such as seasonal fluctuations and measurement error; see Stock and Watson (1999).

Table 3: Bandpass-filtered Moments

	Data	Our Model	Alternative Two-Shock Models			
			TFP	Invt	Disc	News
<i>Standard deviations</i>						
$\text{stddev}(y)$	1.42	1.42	1.54	1.16	1.16	1.32
$\text{stddev}(h)$	1.56	1.52	1.04	1.05	1.05	1.09
$\text{stddev}(c)$	0.76	0.76	0.73	0.95	0.95	0.82
$\text{stddev}(i)$	5.43	5.66	6.76	6.94	6.94	7.15
<i>Correlations</i>						
$\text{corr}(c, y)$	0.85	0.77	0.68	0.22	0.22	0.30
$\text{corr}(i, y)$	0.94	0.92	0.93	0.79	0.79	0.86
$\text{corr}(h, y)$	0.88	0.85	0.94	0.79	0.79	0.88
$\text{corr}(c, h)$	0.84	0.34	0.39	-0.41	-0.41	-0.17
$\text{corr}(i, h)$	0.82	0.99	0.99	0.99	0.99	0.99
$\text{corr}(c, i)$	0.74	0.47	0.38	-0.42	-0.42	-0.21
<i>Correlations with productivity</i>						
$\text{corr}(y, y/h)$	0.08	0.15	0.84	0.44	0.45	0.56
$\text{corr}(h, y/h)$	-0.41	-0.37	0.60	-0.18	-0.18	0.12
$\text{corr}(y, sr)$	0.82	0.85	0.98	0.90	0.90	0.93
$\text{corr}(h, sr)$	0.47	0.47	0.86	0.47	0.47	0.66

Note: The first column reports moments in the US data, bandpass-filtered, over the 1960-2007 period. The second column reports the moments in our model. The rest of the table reports the moments in the four competing models discussed in the text. Red color indicates a significant difference between a model's predicted moment and the corresponding moment in the data.

Our model does a very good job in matching the relevant moments in the data. The only shortcoming is that it underestimates the correlations of consumption with output and hours. As explained in detail in Appendix B, the overall fit owes to a delicate balance between the contribution of the two structural shocks: if we shut down either shock, the model misses multiple moments at a time. But once the two shocks are combined, all the moments fall in place, as if by magic.

Importantly, this “magic” is not a trivial consequence of adding a second shock to the RBC prototype: none of the competing two-shock models is able to replicate the empirical fit of our model. By trying to attribute the variation in the data to structural shocks that do not exhibit the appropriate co-movement patterns, these models end up missing, not only certain key correlations, but also the relative volatilities of certain variables.

Appendix C shows that a similar property characterizes baseline versions of the NK framework. We interpret this as a further indication of the superior ability of the structural mechanism we propose in this paper to capture salient features of the data.

Output gaps, markups, and aggregate demand. Within NK models, the notion of fluctuations in “aggregate demand” has been tied to deviations from the model’s underlying flexible-price allocations. These deviations manifest as variation in markups and in measured “output gaps”, and come together with commensurate movements in inflation.

In our setting, there are no nominal rigidities. Nevertheless, because firms make their input choices prior to observing the demand for their products, a drop in confidence manifests as an increase in the realized markup. Furthermore, the resulting recession will register as a negative “output gap” insofar as the latter is measured relative to the underlying RBC benchmark.

In this regard, the notion of aggregate demand formalized here with the help of higher-order uncertainty has a flavor and empirical content that are similar as those in the NK framework. There is, however, an important difference: in our setting, the fluctuations in the “output gap” may arise without any movements in inflation. Our formalization therefore bypasses the “inflation puzzles” of the NK framework and may help explain, *inter alia*, why the severe contraction in output and employment during the recent recession were not accompanied by severe deflation.¹³ We discuss additional distinguishing aspects of our formalization in Section 8.

7 Extension and estimation

The preceding analysis showed that the introduction of a particular kind of waves of optimism and pessimism in an RBC prototype offers a parsimonious account of multiple salient features of the data, most notably the co-movement patterns documented in Section 2. In this section, we extend the analysis to a pair of richer, “medium scale”, DSGE models, which we estimate on US data.

The two models differ on whether they incorporate nominal rigidity, but are similar in that they accommodate a multitude of structural shocks. They therefore permit us to address the following questions. First, does our mechanism account for a significant fraction of the observed business-cycle volatility once multiple other structural forces are allowed to drive the business cycle? Second, do the structural estimations of these models deliver a picture that is consistent with the one that emerged from the empirical analysis in Section 3? Finally, how does our formalization of “aggregate demand” compare to the one already embedded in the NK framework?

In what follows, we first describe briefly the ingredients of two models and their estimation. We then review the key results and the answers to the above questions. Various details and additional results are delegated to Appendix D.

The two models. The two models we study in this section share the same backbone as our baseline model, but add a number of competing structural shocks, along with habit persistence in consumption, adjustment costs in investment, and, in one of two models, nominal rigidity.

¹³For a discussion of the “inflation puzzles” faced by NK models and an alternative resolution, see Beaudry and Portier (2013). That paper develops a formalization of non-monetary “demand shocks” that rests on the interaction of incomplete markets and news shocks. The quantitative potential of this formalization remains unexplored.

To accommodate price-setting behavior, we now let each island contain a large number of monopolistic firms, each of which produces a differentiated commodity. These commodities are combined through a CES aggregator into an island-specific composite good, which in turn enters the production of the final good in the mainland. In one of the two models we study in this section, firms are free to adjust their price in each and every period, after observing the realized demand for their product; we refer to this model as the flexible-price model. In the other model, firms can instead adjust prices only infrequently, in the familiar Calvo fashion; we refer to this model as the sticky-price model and we close it by adding a conventional Taylor rule for monetary policy.

To accommodate the possibility that the business cycle is explained by multiple structural forces, and to let these forces compete with our mechanism, we consider the following shocks in addition to our confidence shock: a permanent and a transitory TFP shock; a permanent and a transitory investment-specific shock; a news shock regarding future productivity; a transitory discount-rate shock; a government-spending shock; and, in the sticky-price model, a monetary shock.

This menu of shocks is motivated by various considerations. First, previous research has argued that investment-specific technology shocks are at least as important as neutral, TFP shocks (Fischer, 2006). Second, as already noted, news shocks are obvious competitors to our mechanism and have been on center stage in recent business-cycle research. Third, these kinds of shocks, as well as monetary, fiscal, and transitory discount-rate or investment-specific shocks, have been proposed as formalizations of the notion of “aggregate demand shocks” within the NK framework. Fourth, transitory TFP, investment-specific, or discount-rate shocks are often used as proxies for financial frictions that lead to, respectively, misallocation, a wedge in the firm’s investment decisions, or a wedge in the consumer’s saving decisions.¹⁴ Fifth, the introduction of multiple transitory shocks, whatever their interpretation, maximizes the chance that these shocks, rather than our confidence shock, will pick up the transitory fluctuations in the data. All in all, although the menu we have considered does not exhaust all the shocks that have appeared in the literature, we believe it permits us to embed a variety of mechanisms that seem a priori plausible and have also been found to be quantitatively significant in prior structural estimations.

Finally, we allow for the kind of adjustment costs in investment (IAC) and habit persistence in consumption (HP) popularized by Christiano, Eichenbaum and Evans (2005) and Smets and Wouters (2007). These modeling devices lack compelling micro-foundations, but have played an important dual role in the DSGE literature: as sources of persistence; and as mechanisms that help improve the empirical performance of certain shocks, including monetary, investment-specific, discount-rate, and news shocks. Our preceding analysis has already established that these features are *not* needed for our mechanism to deliver realistic fluctuations. Here, we incorporate them in order to keep our exercise as close as possible to standard DSGE practice, as well as to give a better chance to the aforementioned competing shocks to outperform the confidence shock.

¹⁴See Christiano et al (2014) for a recent example of using these shocks as proxies for financial shocks, and Buera and Moll (2012) for a careful analysis of how different types of financial frictions map to different wedges.

Estimation. We estimate our models with Bayesian maximum likelihood in the frequency domain, as in Christiano and Vigfusson (2002) and Sala (2013); see Appendix D for details. The advantage of this method, relative to estimation in the time domain, is that it guides the estimation of a model on the basis of its performance at frequencies that correspond to business-cycle phenomena (between 6 and 32 quarters) as opposed to medium- or long-run trends.¹⁵

The data used in the estimation include GDP, consumption, investment, hours worked, the inflation rate, and the federal fund rate for the period 1960Q1 to 2007Q4.¹⁶ Our sticky-price model is estimated on the basis of all these six variables. By contrast, our flexible-price model is estimated on the basis of real quantities only (GDP, consumption, investment, and hours). The rationale is that this kind of model is not designed to capture the properties of nominal data. For certain purposes, however, it is useful to obtain nominal predictions from this model. In particular, this is necessary when we seek to replicate the empirical strategy of Section 3 on artificial data from the two models. To this goal, we augment our estimated flexible-price model with a simple monetary policy rule that stabilizes inflation without the exogenous random monetary disturbances.¹⁷

The priors used in the estimation are reported in Tables 10 and 11 in Appendix D. The priors for the preference, technology, and monetary parameters are in line with the pertinent literature. The priors for the shocks are guided by two principles: first, we do not a priori favor the confidence shock; and second, we impose a tight prior only on the persistence of the transitory shocks, in order to help the estimation disentangle them from the permanent shocks.

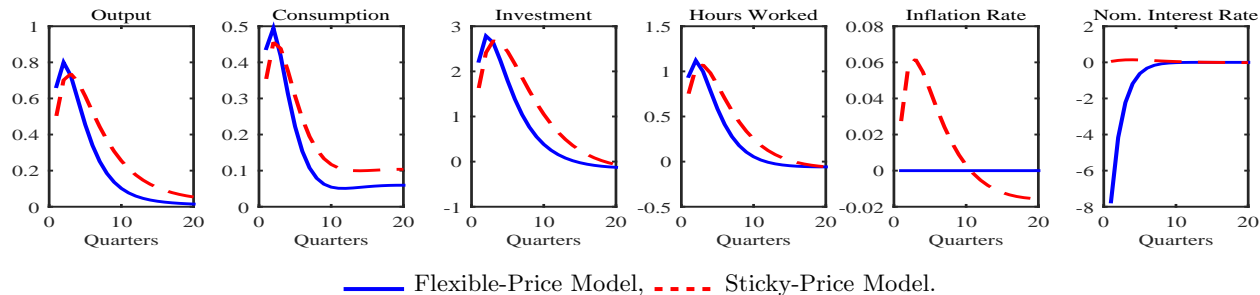
Posterior distributions were obtained with the MCMC algorithm. The resulting estimates of the parameters are reported in the aforementioned tables. The estimated values of the preference, technology, and monetary parameters are close to previous estimates in the literature. As for the shock parameters, the following point is worth making. To the extent that the ξ_t shock represents higher-order uncertainty, its estimated size should not be unreasonably high relative to the estimated size of the payoff-relevant uncertainty. As we discuss in Appendix D, this does not appear to be a serious problem: the level of strategic uncertainty we have accommodated in our estimations with the help of a heterogeneous-prior formulation does not appear to be implausibly large, especially so in the sticky-price model. Having said this, it would certainly be desirable to further discipline the type of quantitative exercises we introduce in this paper by taking specific stands on the micro-structure of how agents update their beliefs and interact with one another and/or by utilizing survey evidence on expectations of economic activity. We leave this to future research.

¹⁵A side-benefit of this approach is that it dispenses with some ad-hoc features that, e.g., Smets and Wouters (2007) need to assume in order to accommodate the observed low-frequency movements in inflation and hours.

¹⁶The first four variables are in logs and linearly de-trended. Following Justiniano, Primiceri and Tambalotti (2010), we do not include the price of investment in our data so as to accommodate both supply-side and demand-side interpretations of either the permanent or the transitory investment shock.

¹⁷This policy is equivalent to allowing the intercept in the Taylor rule to track the underlying natural rate. Furthermore, when this rule is appended to the sticky-price model, the allocations of the latter coincide with those of the flexible-price model. Hence, one can readily re-interpret our flexible-price model as the sticky-price model augmented with the aforementioned policy rule instead of the more standard Taylor rule, whose intercept is a constant.

Figure 3: Theoretical IRFs to Confidence Shock



The confidence shock. Figure 3 reports the estimated IRFs to a positive confidence shock. As far as real quantities are concerned, the IRFs are similar across the two models, as well as similar to those in our baseline model. The introduction of investment-adjustment costs and consumption habit adds a hump-shaped property, but does not alter the co-movement patterns found in the baseline model. This underscores the robustness of the key positive implications our mechanism as we move between RBC and NK settings, or as we add the aforementioned modeling ingredients.

What differs, however, is the behavior of inflation and interest rates. In response to the confidence shock, as well as to other shocks, the flexible-price model predicts implausibly large movements in the real interest rate, due to the inclusion of the particular types of investment-adjustment costs and habit persistence. This compromises the model’s performance vis-a-vis inflation and interest rates.¹⁸ By contrast, because nominal rigidity permits the actual real interest to deviate from its natural level, the sticky-price model is able to accommodate simultaneously modest movements in both the real and the nominal interest rate, as well as in inflation.

Therefore, if we interpret the confidence shock as a formal counterpart of the empirical factor we documented in Section 3, both models capture quite well the co-movement of real quantities, but only the sticky-price model does a good job vis-a-vis inflation and the interest rate.

Table 4: Variance Contribution of Confidence Shock (6–32 Quarters)

	Y	C	I	h	π	R
Flexible Price	50.98	43.72	54.63	76.04	0.00	99.15
Sticky Price	47.73	40.89	44.24	65.66	11.95	32.64

Table 4 turns to the estimated contribution of the confidence shock to the volatility of the key macroeconomic variables at business-cycle frequencies (6–32 quarters). Despite all the competing shocks, the confidence shock emerges as the single most important source of volatility in real

¹⁸Since we have augmented our flexible-price model with a monetary policy that stabilizes inflation, the large volatility of the real rate manifests fully in the nominal rate. But even if we had assumed a different monetary policy, the model not possibly accommodate modest movements in both inflation and the nominal interest rate.

quantities. For example, the confidence shock accounts for 51% of the business-cycle volatility in output in the flexible-price model, and for 48% in the sticky-price model.

Table 5 completes the picture by looking at the estimated contribution of the confidence shock to the *covariances* of output, hours, investment, and consumption. The confidence shock is, by a significant margin, the main driving force behind the co-movement of all these variables.

Table 5: Contribution of Confidence Shock to co-movements (6–32 Quarters)

	$\text{Cov}(Y, h)$	$\text{Cov}(Y, I)$	$\text{Cov}(Y, C)$	$\text{Cov}(h, I)$	$\text{Cov}(h, C)$	$\text{Cov}(I, C)$
Flexible Price	75.80	60.06	56.34	75.67	96.53	84.75
Sticky Price	68.53	53.23	58.40	62.64	106.30	107.41

These findings are not driven by the priors in the estimation: the variance contribution of the confidence shock at the priors is less than 3% for output in either model. Rather, what seems to explain these findings is, first, that the data favor a mechanism that triggers strong procyclical movements in hours, investment, and consumption without commensurate movements in labor productivity, TFP, inflation, and interest rates and, second, that our mechanism is well positioned to generate this kind of co-movement within workhorse macroeconomic models.

The other shocks. To economize on space, the estimated role of the other shocks is reported in Appendix D. Two findings are nevertheless worth reporting. First, the IRFs of all the other shocks are unlike those of the empirical factor we documented in Section 3. And second, the co-movement implications of investment-specific, discount-rate, and news shocks change substantially depending on whether we allow or shut down the particular forms of adjustment costs to investment and habit persistence in consumption that are popular in the NK literature. This explains why the estimated contribution of these shocks, whether in our own models or in the existing literature (e.g., Justiniano et al, 2010, Blanchard et al, 2014), depends heavily on the inclusion of these modeling features. By contrast, neither the co-movement implications of the confidence shock nor its estimated contribution are unduly sensitive to the inclusion or exclusion of these modeling ingredients, as well as to the presumed degree of nominal rigidity—a kind of robustness that we view as an advantage of our formalization of the notion of fluctuations in aggregate demand.

Business-cycle moments. We now evaluate the empirical fit of our estimated models with regard to business-cycle moments. Table 6 reports some key moments of the data (first column); those predicted by our two estimated models (second and third column); and, for comparison purposes, those predicted by the Smets and Wouters (2007) (fourth column) and our own baseline model (last column). Inspecting of these results leads to the following conclusions.

First, both of our estimated models do a good job on the real side of the data. Perhaps this is not surprising given that our baseline model had attained a good fit under a more rigid theoretical structure. But this only underscores the good empirical performance of our mechanism.

Table 6: Band-pass Filtered Moments

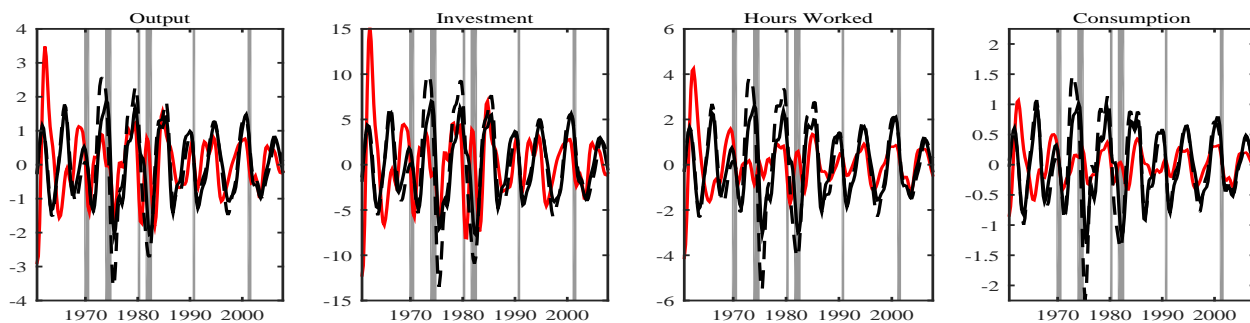
	Data	FP	SP	SW	Baseline	Data	FP	SP	SW	Baseline
	<i>Standard Deviations</i>					<i>Correlations with Output</i>				
y	1.42	1.46	1.43	1.42	1.42					
i	5.43	5.12	5.64	4.86	5.66	0.94	0.87	0.84	0.74	0.92
h	1.56	1.73	1.87	0.97	1.52	0.88	0.82	0.81	0.81	0.85
c	0.76	0.92	0.91	1.11	0.76	0.85	0.82	0.72	0.67	0.77
y/h	0.75	0.97	1.08	0.84	0.79	0.08	0.03	-0.10	0.74	0.15
π	0.23	0.22	0.27	0.34	–	0.21	0.01	0.40	0.13	–
R	0.35	6.42	0.36	0.35	–	0.34	-0.59	0.45	0.06	–
	<i>Correlations with Hours</i>					<i>Correlations with Investment</i>				
i	0.82	0.84	0.83	0.67	0.98					
c	0.84	0.60	0.47	0.59	0.34	0.74	0.55	0.34	0.30	0.47
y/h	-0.41	-0.53	-0.65	0.22	-0.37	0.08	-0.19	-0.33	0.47	-0.24
π	0.44	0.00	0.66	0.23	–	0.09	0.01	0.47	0.18	–
R	0.61	-0.70	0.67	0.21	–	0.24	-0.57	0.54	0.23	–

Note: FP and SP: our estimated flexible- and sticky-price models, SW: the model in Smets and Wouters (2007), Baseline: the calibrated RBC prototype studied in Section 6.

Second, in comparison to Smets and Wouters (2007), our sticky-price model does a good job in matching, not only the real, but also the nominal side of the data. In this regard, the inclusion of our mechanism in NK models does not seem to interfere with their ability to match the nominal side of the data. Nonetheless, as we elaborate below, it does call into question the structural interpretation that existing versions of these models offer for the observed recessions.

Back to the VAR evidence. We have already seen that the IRFs to the confidence shock in our models mirror the IRFs to the factor identified in Section 3. Figure 4 shows that this resemblance carries over to the estimated paths of the two types of shocks.

Figure 4: Confidence Shock (in the models) vs Main Business-Cycle Factor (in the data)



Note: — the fluctuations in the data as predicted by the factor we identified in Section 3, — confidence-driven fluctuations in the flexible-price model, - - - confidence-driven fluctuations in the sticky-price model.

The solid black lines in the figure report the counterfactual paths of output, hours, investment and consumption that obtain from our flexible-price model when it is fed with the estimated path of the confidence shock only. The dashed black lines give the same paths for the sticky-price model. Finally, the solid red lines give the counterfactual paths that obtain when we feed the VAR in Section 3 with our identified factor only.

Two lessons emerge from this picture. First, the confidence-driven fluctuations in the two models are nearly identical. This underscores, once again, the robustness of our mechanism across RBC and NK settings. Second, the confidence-driven fluctuations in our models track quite well the fluctuations in the US data that are attributed to the identified factor of Section 3. Importantly, this match is specific to the confidence shock: if we consider any of the other structural shocks in the model, there is no resemblance to our identified factor.

There are, however, two potential caveats in interpreting the good match between the confidence shock and the empirical factor of Section 3 as a validation of our theory. First, as emphasized by Chari, Kehoe and McGrattan (2008), comparing the theoretical IRFs of a model to the IRFs of a VAR in the data can be misleading insofar as there is no exact mapping between the former and the latter. Second, the identified factor can be a convoluted combination of multiple structural shocks, even under the assumption that our model is the true data-generating process.

To address these issues and shed further light on the empirical performance of our theory, we consider the following exercise. We generate artificial data from our two models and subject them the same empirical strategy as the one applied on actual US data. The idea is that we stop giving any particular interpretation to the evidence of Section 3 and, instead, use empirical methodology and the identified factor merely as a filter: we ask whether applying this filter to our models gives the same picture as applying this filter to US data.¹⁹

Figure 5 answers this question in terms of IRFs to the identified factor.²⁰ Both models pass the test with regard to the real quantities: the IRFs obtained when we apply our empirical strategy to artificial data generated by either model mirror quite well those obtained in actual US data. The only significant difference between the two models is that the sticky-price model outperforms the flexible-price one with regard to interest rates, for the reasons discussed earlier.

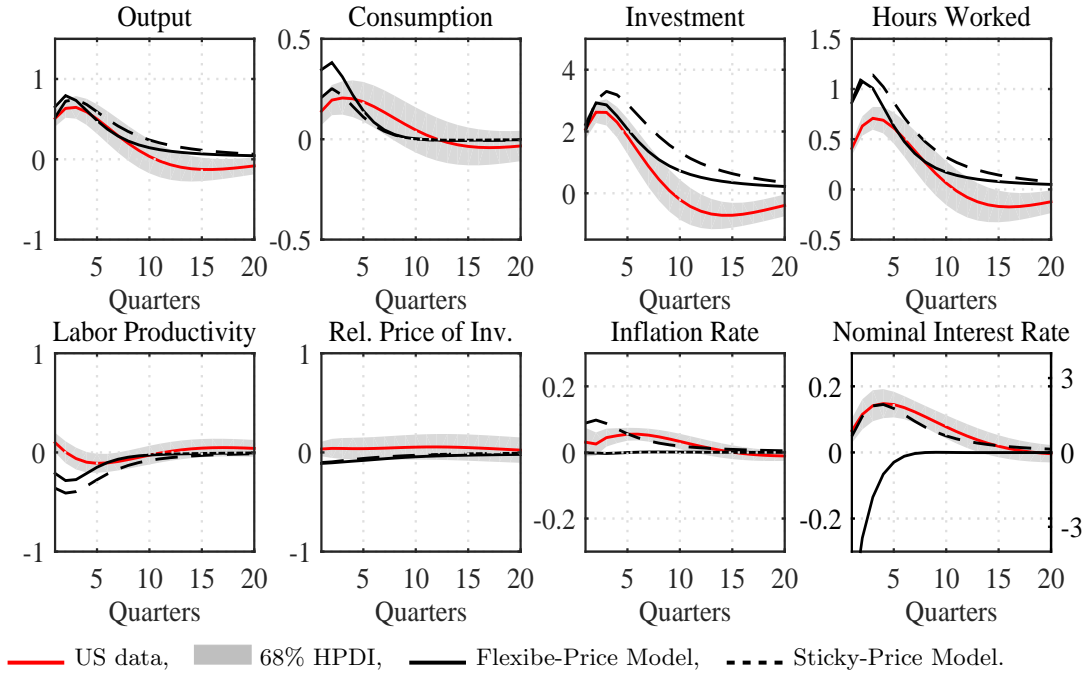
Shocks vs empirical proxies. To further corroborate our theory, we now consider the following question: do the technology and confidence shocks in our theory have any resemblance to empirical measures of, respectively, TFP and “market psychology” in the real world?

We address this question in Figure 6 by comparing the estimated series of the technology and

¹⁹There is a minor difference between the VAR specification considered here and the one in Section 3, namely, the latter does not contain government spending. This omission is due to the fact that in our models the sum of consumption, investment, and government spending coincides with GDP: if we were to include all these variables in the VAR, we would face a multi-collinearity problem. For consistency, in this section we also drop government spending from the VAR we run on the US data. As is evident from a comparison of Figures 1 and 5, this makes no essential difference for the picture that emerges from the data.

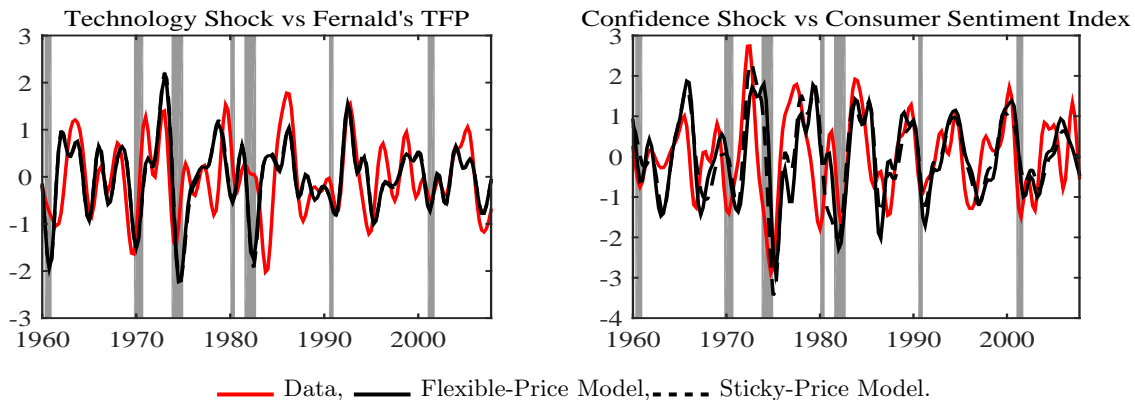
²⁰See Appendix D for a similar exercise in terms of variance decompositions.

Figure 5: IRFs to Identified Factor from the Same VAR on Data and Models



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

Figure 6: Shocks vs Empirical Proxies



8 Nominal rigidities, aggregate demand, and the US recessions

We conclude our analysis by asking the following question: what explains the apparent deficiency in aggregate demand during recessions?

To address this question, one must first define the “potential” level of economic activity, relative to which the aforementioned “deficiency” is to be measured. Our approach shares with the NK framework the same perspective on what that potential is: it is the level of economic activity predicted by the underlying RBC benchmark, which abstracts from both nominal rigidities and imperfections in beliefs. The two approaches differ, however, with regard to the friction that creates “output gaps”: in the NK tradition, the key friction is a nominal rigidity; in our approach, it is a flaw in the agents’ coordination of their beliefs and actions.

The two approaches are not mutually inconsistent. Our view is that *both* types of frictions are important for understanding business cycles. Furthermore, the NK mechanism appears to complement our mechanism by enhancing the impact of higher-order beliefs in the following manner: although the estimated contribution of our mechanism to the volatility of macroeconomic outcomes is similar across the two models studied in the previous section, the estimated variance of the confidence shock is much smaller in the sticky-price model than that in the flexible-price model.

Notwithstanding this point, we now evaluate the relative significance of the two mechanisms, when each one operates in isolation. We do this by computing the posterior odds that the data are generated by either of three possible models: the variant of our flexible-price model that shuts down the confidence shock; the variant of our sticky-price model that also shuts down the confidence shock; and finally the flexible-price model that contains the confidence shock. The first model serves as a benchmark, capturing the RBC core of the other two models. The second model adds a nominal rigidity to the first model, isolating the NK mechanism. The third model replaces the nominal rigidity with a confidence shock, isolating our proposed mechanism. To assess their fit vis-a-vis real economic activity, all three models are estimated on the basis of real quantities only. We then compute the posterior odds that the data are generated by the sticky-price model rather than either of the flexible-price models, starting with an uniform prior over the three models. These odds provide a metric of how well a given model captures the data relative to another model.²¹

Consider first the pair-wise comparison between the sticky-price model and the standard flexible-price model. In this case, the sticky-price model wins: the posterior odds that the data are generated by that model are 90%. Consider next how this comparison is affected once the flexible-price model is augmented with the confidence shock. The odds are now completely reversed: the probability that the data are generated by the sticky-price model are only 2%. By this metric, nominal rigidity is essential for the ability of the theory to match the real data when our mechanism is absent, but not once it is present: our mechanism appears to be *more* potent than the NK mechanism when their relative performance is evaluated in terms of likelihood, as described above.

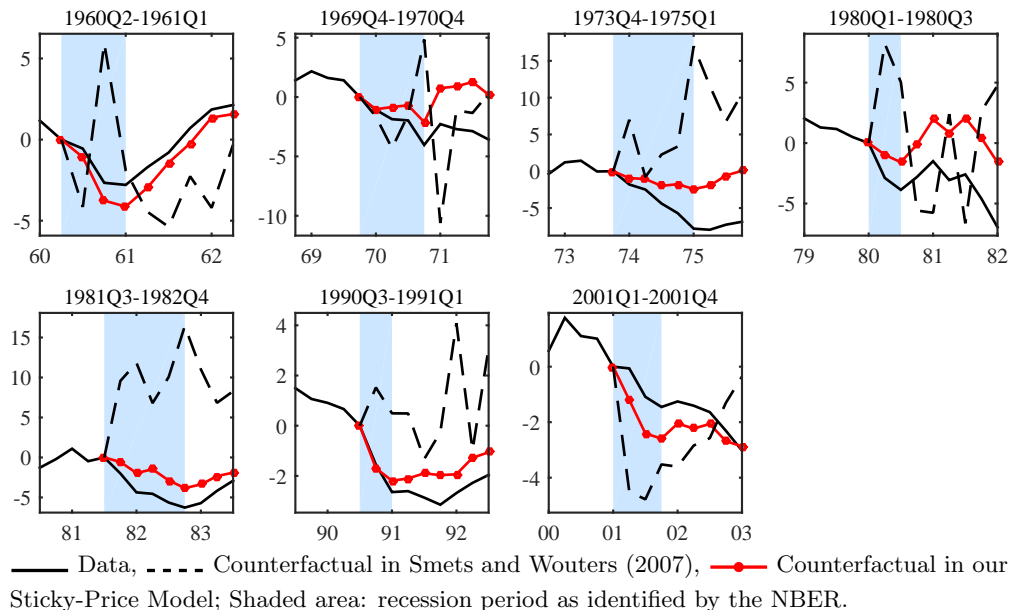
²¹See Table 15 in Appendix C for a tabulation of the posterior odds of the aforementioned three models along with those of the model that combines our mechanism with the NK mechanism.

We interpret these results as providing further confirmation, not only of our mechanism’s ability to capture salient aspects of the data, but also to operationalize the notion of fluctuations in aggregate demand. Our approach can serve as a potent *substitute* to the NK approach—but it can also *complement* it by providing what, in our view, is a more appealing structural interpretation of the observed business-cycle phenomena.

To illustrate this last point, we now take a closer look at the structural interpretation of the US recessions offered by two alternative models: our preferred version of the NK model, which contains the confidence shock; and the “canonical” NK model of Smets and Wouters (2007). Note that both models give a prominent role to nominal rigidity, but differ in the way they formalize and quantify the fluctuations in aggregate demand.

In each of these models, we ask the following counterfactual: how would the historical recessions have looked if nominal rigidities were shut down?²² Figure 7 answers this question in terms of the dynamics of output: the red dotted lines in the figure give the counterfactual path of output in our model; the blue dashed lines give the counterfactual path in the model of Smets and Wouters (2007); the black solid lines give the actual data. Note that, by construction, both models match perfectly the actual data when the nominal rigidity is at work. The gap between the aforementioned counterfactuals and the actual data therefore reveals the precise role that nominal rigidity plays within each model.

Figure 7: Counterfactuals from Removing Nominal Rigidity



²²More precisely, the counterfactual is constructed as follows. For each of the two models, we fix the estimated paths of all the structural shocks, as well as of all the estimated preference and technology parameters. We also maintain the estimated nominal parameters until the onset of the recession under consideration. We then compute the counterfactual path of output that obtains when the nominal rigidity is shut down from that point on.

In our version of the NK model, which contains the confidence shock, recessions look qualitatively similar whether the nominal rigidity is shut down or not. In particular, our model predicts that the nominal rigidity exacerbated most of the recessions, which seems consistent with some economists’ priors, but attributes the bulk of the recessions to forces that would have remained potent even if the nominal rigidity were absent. By contrast, the model of Smets and Wouters (2007) attributes the recessions to forces that rely on nominal rigidity so heavily that most of the recessions are predicted to turn into *booms* when the nominal rigidity is shut down.

There is no obvious way to test these counterfactuals, or to know which model is “right”. We nevertheless find it hard to accept a structural interpretation of the business cycle that hinges on the prediction that the observed recessions would turn into booms in the absence of nominal rigidities. Instead, we favor an interpretation that attributes a role to “coordination failures” and “lack of confidence”, while also accommodating a meaningful role for monetary policy.

9 Conclusion

By assuming a particular solution concept together with complete information, standard macroeconomic models impose a rigid structure on how agents form beliefs about endogenous economic outcomes and how they coordinate their actions. In this paper, by contrast, we introduced a certain relaxation of this structure and evaluated its quantitative implications. In particular, we augmented DSGE models with a tractable form of higher-order belief dynamics that proxies the aggregate effects of strategic uncertainty and captures a certain kind of waves of optimism and pessimism about the short-term outlook of the economy. We believe that this adds to our understanding of business-cycle phenomena along the following dimensions:

- It offers a parsimonious explanation of salient features of the macroeconomic data, including standard business-cycle moments and the dynamic patterns detected in our empirical analysis.
- It appears to outperform alternative structural mechanisms that are popular in the literature, calling into question, or refining, prevailing explanations of business-cycle phenomena.
- It offers a potent formalization of the notion of fluctuations in “aggregate demand” that can either substitute or complement the NK formalization of this notion.
- It leads to a structural interpretation of the observed recessions that is not unduly sensitive to the degree of nominal rigidity (whose bite at the aggregate level remains debatable) and that attributes a potentially significant role to forces that can be interpreted as “coordination failures” or “lack of confidence”.

These findings beg the question of what triggers the drop in confidence during a recession, or more generally the waves of optimism and pessimism in the agents’ beliefs about one another’s actions. Having treated the “confidence shock” as exogenous, we can not offer a meaningful answer

to this question. This limitation, however, is not specific to what we do in this paper: any formal model must ultimately attribute the business cycle to some exogenous trigger, whether this is a technology shock, a discount-rate shock, a financial shock, or even a sunspot. Therefore, although we have to remain agnostic about the “micro-foundations” of the aforementioned belief waves, we hope to have provided a useful gauge of their potential quantitative importance along with a re-evaluation of prevailing theories of the business cycle.

We conclude with the following note. Our investigation of the data in Section 3 indicated that a single transitory force can account for the bulk of the observed business-cycle fluctuations. The structural estimations we conducted in Section 7 approached the data with a different empirical strategy but produced essentially the same result: they attributed the bulk of the observed business-cycle volatility to a particular structural shock that resembled the empirical factor documented in Section 3. Admittedly, this coincidence does not prove that the theory we proposed in this paper offers the most relevant, or appealing, formal interpretation of the data. Nonetheless, we think that our results set a useful reference point, which could invite future research either in the direction of further quantitative evaluations of the macroeconomic effects of strategic uncertainty, or in the direction of developing novel structural interpretations of the business-cycle data.

APPENDICES

A. Data and Identification

In this appendix we describe the data we use in this paper and the identification of the empirical factor that we introduce in Section 3.

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

Table 7 summarizes information about the data. GDP, Y , is measured by the seasonally adjusted GDP. Consumption, C , is measured by the sum of personal consumption expenditures in nondurable goods (CND) and services (CS). Investment, I , is measured by the sum of personal consumption expenditures on durable goods (CD), fixed private investment (FPI) and changes in inventories (DI). Government Spending, G , is measured by government consumption expenditures (GCE). Hours worked, H , are measured by hours of all persons in the nonfarm business sector. Labor productivity, Y/H , is measured by Real Output Per Hour of All Persons in the nonfarm business sector. The inflation rate, π , is the log-change in the implicit GDP deflator. The nominal interest rate, R , is the effective federal funds rate measured on a quarterly basis. Given that the effective federal funds rate is available at the monthly frequency, we use the average over the quarter (denoted FEDFUNDS).

The relative price of investment was built in the same way as in Benati (2014), and we thank him for his help in doing so. This involves constructing a chained price index for the sum of gross private investment and consumption of durables, along with a chained price index for the sum of consumption of non-durables and services. The relative price of investment is then given by the ratio of these two quantities. For a detailed description, see <http://economics.mit.edu/files/10307>.

Identification. Our VAR/VECM admits the following representation:

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

where Y_t is a N -dimensional random vector, $A(L)$ is a matrix polynomial, $E(u_t u_t') = \Sigma$ and $C(0) = I$. The number of lags, 1, was selected according to the Bayesian Information Criterion (BIC). We assume that there exists a linear mapping between the VAR innovations u_t and some underlying shocks ε_t :

$$u_t = S\varepsilon_t$$

for some matrix S . We normalize $\mathbb{E}(\varepsilon\varepsilon') = I$, so that the matrix S satisfies $SS' = \Sigma$.

Table 7: Description of the Data

Description of the Data	
Data	Formula
GDP	$Y = \text{GDP} / (\text{GDPDEF} \times \text{CNP160V})$
Consumption	$C = (\text{CND} + \text{CS}) / (\text{GDPDEF} \times \text{CNP160V})$
Investment	$I = (\text{CD} + \text{FPI} + \text{DI}) / (\text{GDPDEF} \times \text{CNP160V})$
Government Spending	$G = \text{GCE} / (\text{GDPDEF} \times \text{CNP160V})$
Hours Worked	$H = \text{HOANBS} / \text{CNP160V}$
Labor Productivity	GDP / H
Inflation Rate	$\pi = \log(\text{GDPDEF}) - \log(\text{GDPDEF})_{-1}$
Nominal Interest Rate	$R = \text{FEDFUNDS} / 4$
Relative Price of Investment	see http://economics.mit.edu/files/10307
Mnemonic	Source
GDP	http://research.stlouisfed.org/fred2/series/GDP
CND	http://research.stlouisfed.org/fred2/series/PCND
CD	http://research.stlouisfed.org/fred2/series/PCEDG
CS	http://research.stlouisfed.org/fred2/series/PCESV
FPI	http://research.stlouisfed.org/fred2/series/FPI
DI	http://research.stlouisfed.org/fred2/series/CBI
GCE	http://research.stlouisfed.org/fred2/series/GCE
HOANBS	http://research.stlouisfed.org/fred2/series/HOANBS
GDPDEF	http://research.stlouisfed.org/fred2/series/GDPDEF
FEDFUNDS	http://research.stlouisfed.org/fred2/series/FEDFUNDS
CNP160V	http://research.stlouisfed.org/fred2/series/CNP160V

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

As in Beaudry et al. (2012), the $MA(\infty)$ 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 $C_{\tau} = B_{\tau}\tilde{S}$. Then, the impulse response vector of variable $Y_{i,t}$ to the 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 $q^{(j)}$ is the j -th column of Q .

Our aim is to identify a factor—a linear combination of the VAR innovations or, equivalently, the ε shocks—that captures most of the business cycle. To this goal, 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

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

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

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

where $f(\omega)$, $\omega \in [-\pi; \pi]$ is given by

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

The filtered volatility of variable k , σ_{kk}^2 , 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

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

where $\tilde{C}_k(e^{-i\omega})$ 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, $\hat{\sigma}_{kk}^2$. The variance decomposition for variable k is then given by the ratio $\theta_k = \hat{\sigma}_{kk}^2 / \sigma_{kk}^2$. This calculation can be done for m variables, the problem is then to find a vector x that solves the problem

$$\begin{aligned} \max_x \quad & \sum_{k=1}^m \alpha_k x \theta_k x' \\ \text{s.t.} \quad & xx' = 1 \end{aligned}$$

where α_k is the weight assigned to variable k . This problem is a simple eigenvalue problem and can easily be solved numerically.

The vector q is then given by $(\mathbf{0}_p, x)'$ where $\mathbf{0}_p$ is a zero vector of length p (the number of unit-root components).

B. An Anatomy of our Baseline Model

Here we elaborate on the distinct quantitative roles that the technology and the confidence shock play in our RBC prototype. To this goal, Table 8 reports the business-cycle moments predicted by the versions of our model that shut down either of the two shocks (see the last two columns) and compares them to those in the data (first column) and in our full model (second column).

Table 8: Bandpass-filtered Moments

	Data	Our Model	A only	ξ only
<i>Standard deviations</i>				
$\text{stddev}(y)$	1.42	1.42	1.46	1.18
$\text{stddev}(h)$	1.56	1.52	0.49	1.70
$\text{stddev}(c)$	0.76	0.76	1.13	0.14
$\text{stddev}(i)$	5.43	5.66	3.01	6.08
<i>Correlations</i>				
$\text{corr}(c, y)$	0.85	0.77	0.99	0.84
$\text{corr}(i, y)$	0.94	0.92	0.99	0.99
$\text{corr}(h, y)$	0.88	0.85	0.99	0.99
$\text{corr}(c, h)$	0.84	0.34	0.99	0.80
$\text{corr}(i, h)$	0.82	0.99	0.99	0.99
$\text{corr}(c, i)$	0.74	0.47	0.99	0.81
<i>Correlations with productivity</i>				
$\text{corr}(y, y/h)$	0.08	0.15	0.99	-0.98
$\text{corr}(h, y/h)$	-0.41	-0.37	0.99	-0.99
$\text{corr}(y, sr)$	0.82	0.85	0.99	0.98
$\text{corr}(h, sr)$	0.47	0.47	0.99	0.99

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

Consider now the fourth column, which isolates the confidence shock.²³ The key failures are now the counterfactually high volatility in hours, the perfectly negative correlation between labor

²³This version is similar to the model studied in Section 7 of Angeletos and La'O (2013).

productivity and either hours or output, and the counterfactually low volatility in consumption. The first two properties follow directly from the fact that technology is fixed and exhibits diminishing returns in labor, while the last property is driven by the transitory nature of the confidence shock.

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

C. Competing Structural Shocks: A Horserace

This appendix studies the “horserace” introduced in Section 6, when we compared the empirical fit of our baseline model to that obtained in other two-shock models. It also provides a formal basis for the claims we make in Section 3 regarding the inability of some popular structural shocks to capture the empirical factor we identified in the data, at least insofar as we stay within the context of canonical RBC and NK models.

Recall that the competing models were variants of our baseline model that replace the confidence shocks with one of the following: a transitory technology shock; a transitory investment-specific shock; a transitory discount-factor shocks; or a news shocks about future TFP. Also recall that the analysis was limited to an RBC model. Here, we complement that analysis in two ways: we revisit the horserace in the context of a NK model; and we add a monetary shock to the set of shocks contained in the horserace.

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

Table 9 reports the same kind of information as Table 3, that is, the business-cycle moments of the different models in the horserace, but now for the NK versions of these models. Figure 8 reports the IRFs for both the RBC and the NK versions of the models.

Four key lessons emerge. First, the superior empirical performance of our baseline model survives when we introduce sticky prices. Second, with the exception of the monetary shock, none of the aforementioned competing structural shocks is able to generate the kind of co-movement patterns in the real data that was documented in Section 3, whether one considers the RBC or the NK version of the models. Third, the monetary shock can match these patterns quite well, but only at

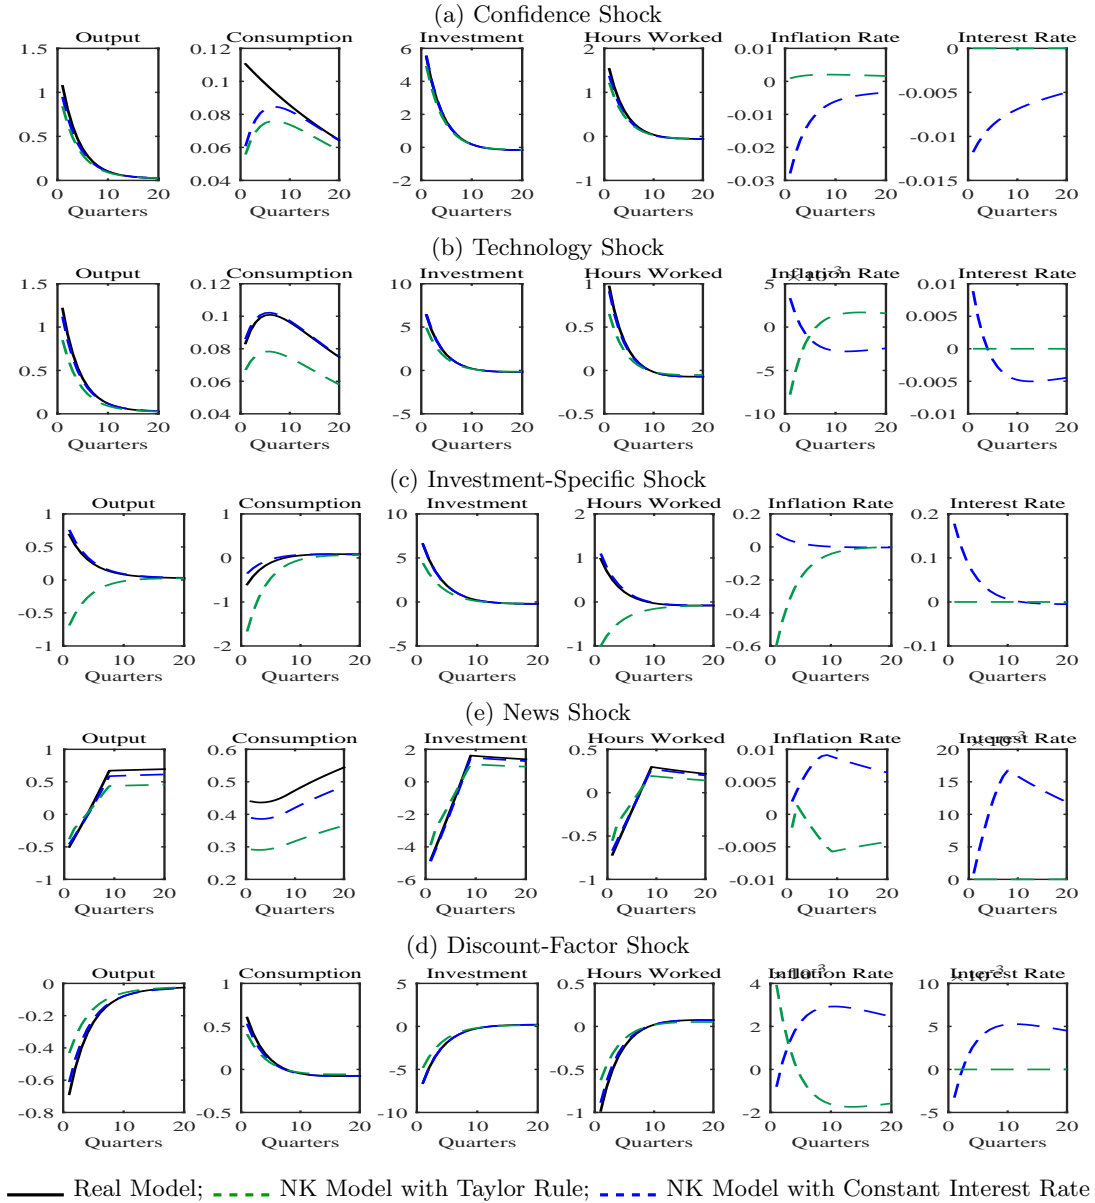
the expense of requiring an implausibly large contribution of monetary shocks to the overall business cycle and of having counterfactual movements on the nominal side (strongly counter-cyclical interest rates and strongly pro-cyclical inflation). Finally, the similarity between the real effects of the confidence shock and those of the monetary shock provide further justification for reinterpreting the confidence shock as some short of “aggregate demand shock”.

Table 9: Bandpass-filtered Moments

	Data	Our Model	Alternative Two-Shock Models				
			TFP	Invt	Disc	News	Mon.
<i>Standard deviations</i>							
$\text{stddev}(y)$	1.42	1.38	1.48	1.24	1.15	1.29	1.37
$\text{stddev}(h)$	1.56	1.45	1.00	1.18	0.97	1.02	1.44
$\text{stddev}(c)$	0.76	0.77	0.75	0.86	0.95	0.84	0.77
$\text{stddev}(i)$	5.43	6.10	6.98	7.03	7.04	7.24	6.20
$\text{stddev}(y/h)$	0.75	0.79	0.67	0.76	0.73	0.65	0.80
$\text{stddev}(sr)$	1.25	1.03	1.30	1.06	1.04	1.16	1.14
<i>Correlations</i>							
$\text{corr}(c, y)$	0.85	0.75	0.71	0.42	0.37	0.43	0.73
$\text{corr}(i, y)$	0.94	0.90	0.92	0.82	0.75	0.84	0.90
$\text{corr}(h, y)$	0.88	0.84	0.93	0.80	0.77	0.86	0.84
$\text{corr}(c, h)$	0.84	0.29	0.40	-0.19	-0.29	-0.07	0.24
$\text{corr}(i, h)$	0.82	0.99	1.00	1.00	1.00	1.00	0.99
$\text{corr}(c, i)$	0.74	0.40	0.38	-0.17	-0.33	-0.13	0.35
<i>Correlations with productivity</i>							
$\text{corr}(y, y/h)$	0.08	0.18	0.83	0.37	0.54	0.61	0.20
$\text{corr}(h, y/h)$	-0.41	-0.36	0.56	-0.24	-0.10	0.13	-0.36
$\text{corr}(y, sr)$	0.82	0.69	0.98	0.92	0.92	0.94	0.94
$\text{corr}(h, sr)$	0.47	0.21	0.85	0.52	0.49	0.65	0.61
<i>nominal variables</i>							
σ_π	0.23	0.03	0.01	0.08	0.01	0.01	0.06
σ_R	0.35	0.02	0.02	0.18	0.02	0.02	0.02
$\rho_{\pi, y}$	0.21	-0.42	0.82	0.71	0.85	0.64	0.81
$\rho_{R, y}$	0.34	0.23	0.88	0.69	0.89	0.75	0.61

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

Figure 8: Horseshoe IRFs (I)



D. Estimated Models

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

The details of the two models. As mentioned in the main text, the two models we study in Section 7 share the same backbone as our baseline model, but add a number of structural shocks along with certain forms of habit persistent in consumption and adjustment costs in investment, as in Christiano et al. (2005) and Smets and Wouters (2007). To accommodate monopoly power and sticky prices, we also introduce product differentiation within each island.

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

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

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

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

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

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

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

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

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

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

where ε_{t-j}^n is the component of the period- t innovation in TFP that becomes known in period $t-j$.²⁴ In line with our baseline model, the confidence shock is now modeled as a shock to higher-order beliefs of a_t^p .

²⁴We have experimented with alternative forms of diffusion, as well as with specifications such as $\zeta_t^n = \varepsilon_{t-4}^n$, and we have found very similar results.

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

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

where ζ_t^c is a transitory preference shock, modeled as an AR(1), $b \in (0, 1)$ is a parameter that controls for the degree of habit persistence, and C_{t-1} denotes the aggregate consumption in the last period.²⁵

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

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

where Z_t denotes the relative price of investment, G_t denotes government spending (the cost of which is assumed to equally spread across the islands), and $\exp(Z_t)\Psi(u_{it})$ denotes the resource cost of utilization per unit of capital. The latter is scaled by $\exp(Z_t)$ in order to transformed the units of capital to units of the final good, and thereby also guaranteed a balanced-growth path. Z_t is modeled as a random walk: $Z_t = Z_{t-1} + \varepsilon_t^z$. Literally taken, this represents an investment-specific *technology* shock. But since our estimations do not include data on the relative price of invest, this shock can readily be re-interpreted as a demand-side shock. Government spending is given by $G_t = \bar{G} \exp(\tilde{G}_t)$, where \bar{G} is a constant and

$$\tilde{G}_t = \zeta_t^g + \frac{1}{1-\alpha}a_t - \frac{\alpha}{1-\alpha}Z_t.$$

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

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

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

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

The above description completes the specification of the flexible-price model of Section 7. The sticky-price model is then obtained by embedding the Calvo friction and a Taylor rule form monetary policy. In particular, the probability that any given firm resets its price in any given period is given by $1 - \chi$, with $\chi \in (0, 1)$. As for the Taylor rule, the reaction to inflation is given by $\kappa_\pi > 1$,

²⁵Note that we are assuming that habit is external. We experimented with internal habit, as in Christiano et al (2007), and the results were virtually unaffected.

the reaction to the output gap is given by $\kappa_y > 0$, and the parameter that controls the degree of interest-rate smoothing is given by $\kappa_R \in (0, 1)$; see condition (18) below.

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

$$\tilde{\lambda}_{it} = \zeta_t^c - \frac{1}{1-b}\tilde{c}_{it} + \frac{b}{1-b}\tilde{C}_{t-1} \quad (8)$$

$$\zeta_t^c + \nu\tilde{n}_{it} = \mathbb{E}_{it} \left[\tilde{\lambda}_{it} + \tilde{s}_{it} + \tilde{Y}_t - \tilde{n}_{it} \right] \quad (9)$$

$$Z_t + \frac{1}{1-\psi}\tilde{u}_{it} = \tilde{s}_{it} + \tilde{Y}_t - \tilde{k}_{it} \quad (10)$$

$$\tilde{y}_{it} = a_t + \alpha(\tilde{u}_{it} + \tilde{k}_{it}) + (1-\alpha)\tilde{n}_{it} \quad (11)$$

$$\tilde{Y}_t = s_c\tilde{c}_{it} + (1-s_c-s_g)(Z_t + \tilde{v}_{it}) + \tilde{G}_t + \alpha\tilde{u}_{it} \quad (12)$$

$$\tilde{k}_{it+1} = \delta\tilde{v}_{it} + (1-\delta)\tilde{k}_{it} \quad (13)$$

$$\tilde{q}_{it} = (1+\beta)\varphi\tilde{v}_{it} - \varphi\tilde{v}_{t-1} - \beta\varphi\mathbb{E}'_{it}\tilde{v}_{it+1} + Z_t - \zeta_t^i \quad (14)$$

$$\tilde{R}_t = \tilde{\lambda}_{it} - \mathbb{E}'_{it} \left[\tilde{\lambda}_{it+1} - \tilde{\pi}_{it+1} \right] \quad (15)$$

$$\tilde{\lambda}_{it} + \tilde{q}_{it} = \mathbb{E}'_{it} \left[\tilde{\lambda}_{it+1} + (1-\beta(1-\delta)) \left(\tilde{s}_{it+1} + \tilde{Y}_{t+1} - \tilde{u}_{it+1} - \tilde{k}_{it+1} \right) + \beta(1-\delta)\tilde{q}_{it+1} \right] \quad (16)$$

$$\tilde{X}_t = s_c\tilde{C}_t + (1-s_c-s_g)(Z_t + \tilde{I}_t) + s_g\tilde{G}_t \quad (17)$$

$$\tilde{R}_t = \kappa_R\tilde{R}_{t-1} + (1-\kappa_R) \left(\kappa_\pi\tilde{\pi}_t + \kappa_y(\tilde{X}_t - \tilde{X}_t^F) \right) + \zeta_t^m \quad (18)$$

$$\tilde{\pi}_{it} = (1-\chi)(1-\beta\chi)\mathbb{E}_{it}\tilde{s}_{it} + (1-\chi)\mathbb{E}_{it}\tilde{\Pi}_t + \beta\chi\mathbb{E}_{it}\tilde{\pi}_{it+1} \quad (19)$$

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

The interpretation of the above system is familiar. Condition (8) gives the marginal utility of consumption. Conditions (9) and (10) characterizes the equilibrium employment and utilization levels. Condition (11) and (12) give the local output and the local resource constraint. Conditions (13) and (14) give the local law of motion of capital and the equilibrium investment decision. Conditions (15) and (16) are the two Euler equations: the first corresponds to optimal bond holdings and gives the relation between consumption growth and the interest rates, while the second corresponds to optimal capital accumulation and gives the evolution of Tobin's Q. Condition (17) gives the measured aggregate GDP. Conditions (18) gives the Taylor rule for monetary policy. Finally, condition (19) gives the inflation rate in each island; aggregating this condition across islands gives our model's New Keynesian Phillips Curve. The only essential novelty in all the above is the presence of the subjective expectation operators in the conditions characterizing the local equilibrium outcomes of each island.

Finally, the flexible-price allocations are obtained by the same set of equations, modulo the following changes: we set $s_{it} = 0$, meaning that the realized markup is always equal to the optimal

markup; we restate the Euler condition (15) in terms of the real interest rate; and we drop the nominal side of this system, namely conditions (18) and (19).

Estimation. As mentioned in the main text, we estimate the model on the frequency domain. This method amounts to maximizing the following posterior likelihood function:

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

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

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

Here, Y_t and X_t denote, respectively, the vector of observed variables and the underlying state vector of the model; ε is the vector of the exogenous structural shocks, drawn from a Normal distribution with mean zero and variance-covariance matrix $\Sigma(\theta)$; $M_y(\theta)$ and $M_x(\theta)$ are matrices whose elements are (non-linear) functions of the underlying structural parameters θ ; and finally M_e is a selection matrix that describes how each of the structural shocks impacts on the state vector. The model spectral density of the the vector Y_t is

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

where $\omega \in [0, 2\pi]$ denotes the frequency at which the spectral density is evaluated. The likelihood function is asymptotically given by

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

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 1 when the frequency falls between 6 and 32 quarters, and 0 otherwise.

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

The inverse labor supply elasticity, ν , is Gamma distributed around 0.5 with standard deviation 0.25. The capital share, α , is Beta distributed around 0.3 with standard deviation 0.05. The utilization elasticity parameter, ψ , is Beta distributed around 0.3 with standard deviation 0.25. The parameter governing the size of investment adjustment costs, φ , is Gamma distributed around

Table 10: Estimated Parameters, Part I

	Priors			Posteriors	
	Shape	Mean	Std. Dev.	FP	SP
<i>Real Aspects</i>					
ν	Gamma	0.500	0.250	0.491 [0.106, 0.941]	0.627 [0.193, 1.066]
α	Beta	0.300	0.050	0.295 [0.255, 0.336]	0.330 [0.286, 0.374]
ψ	Beta	0.300	0.150	0.365 [0.113, 0.647]	0.406 [0.143, 0.656]
φ	Gamma	2.000	1.000	1.720 [0.849, 2.810]	1.965 [0.614, 3.697]
b	Beta	0.500	0.250	0.702 [0.604, 0.799]	0.667 [0.518, 0.805]
<i>Nominal Aspects</i>					
χ	Beta	0.500	0.250	–	0.703 [0.617, 0.786]
κ_R	Beta	0.750	0.100	–	0.434 [0.291, 0.568]
κ_π	Normal	1.500	0.250	–	1.415 [0.999, 1.946]
κ_y	Normal	0.125	0.050	–	0.123 [0.039, 0.210]

Note: 95% HPDI into brackets.

Table 11: Estimated Parameters, Part II

	Priors			Posteriors	
	Shape	Mean	Std. Dev.	FP	SP
<i>Shocks: persistence</i>					
ρ_a	Beta	0.500	0.100	0.484 [0.294, 0.673]	0.496 [0.301, 0.702]
ρ_i	Beta	0.500	0.100	0.507 [0.316, 0.693]	0.522 [0.342, 0.702]
ρ_c	Beta	0.500	0.100	0.499 [0.303, 0.683]	0.496 [0.299, 0.697]
ρ_g	Beta	0.500	0.100	0.581 [0.394, 0.766]	0.586 [0.394, 0.768]
ρ_m	Beta	0.500	0.100	–	0.412 [0.258, 0.570]
ρ_ξ	Beta	0.500	0.100	0.557 [0.370, 0.737]	0.761 [0.707, 0.809]
<i>Shocks: volatilities</i>					
$\sigma_{a,p}$	Inv. Gamma	1.000	4.000	0.579 [0.358, 0.831]	0.567 [0.354, 0.781]
$\sigma_{a,t}$	Inv. Gamma	1.000	4.000	0.448 [0.235, 0.687]	0.442 [0.234, 0.671]
$\sigma_{a,n}$	Inv. Gamma	1.000	4.000	0.553 [0.253, 0.925]	0.604 [0.326, 0.892]
$\sigma_{i,p}$	Inv. Gamma	1.000	4.000	1.135 [0.256, 2.356]	0.963 [0.234, 2.056]
$\sigma_{i,t}$	Inv. Gamma	1.000	4.000	1.286 [0.232, 2.920]	3.198 [0.758, 6.798]
σ_c	Inv. Gamma	1.000	4.000	0.907 [0.283, 1.628]	0.565 [0.236, 1.007]
σ_g	Inv. Gamma	1.000	4.000	1.906 [1.243, 2.659]	1.824 [1.130, 2.576]
σ_m	Inv. Gamma	1.000	4.000	–	0.282 [0.209, 0.369]
σ_ξ	Inv. Gamma	1.000	4.000	3.353 [1.300, 6.152]	0.586 [0.274, 1.021]

Note: 95% HPDI into brackets.

2 with a standard deviation 1; the relatively high standard deviation reflects our own uncertainty about this modeling feature, but also allows the estimation to accommodate the higher point estimates required by the pertinent DSGE literature, in case that would improve the empirical performance of the model. The Calvo probability of not resetting prices, χ , is Beta distributed around 0.5 with standard deviation 0.25. The persistence parameter in the Taylor rule, κ_R , is Beta distributed around 0.75 with standard deviation 0.1; the reaction coefficient on inflation, κ_π , is Normally distributed around 1.5 with standard deviation 0.25; and the reaction coefficient on the output gap, κ_y , is also Normally distributed with mean 0.125 and standard deviation 0.05. Finally, the following three parameters are fixed: the discount factor β is 0.99; the depreciation rate δ is 0.015; and the CES parameter η is such that the monopoly markup is 15%.

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

The estimated σ_ξ . To the extent that we want to think of the confidence shock as a proxy for the type of higher-order uncertainty that can obtain in common-prior settings, one would like to have a theorem that provides a tight bound on the size of this higher-order uncertainty as a function of the underlying payoff uncertainty. By comparing the estimated value of σ_ξ to that bound, we could then judge the empirical plausibility of the presumed level of strategic uncertainty even if we don't know anything about the details of the underlying information structure.

In static settings, such a bound can be obtained with the methods developed in Bergeman and Morris (2013) and Bergeman, Haumman and Morris (2014). Unfortunately, analogous methods are not available for dynamic settings: one would like to have bounds on the volatility generated by higher-order uncertainty at different frequencies, but it is unclear how to obtain such bounds.

Having said this, the following seems a reasonable guess. Due to the 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.

In want of a better alternative, we propose the following rough metric. 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 0.26 when $T = 8$ years and 0.07 when $T = \infty$. For the sticky-price model, the corresponding ratios are 0.08 and 0.02. 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 we view the estimated higher-order uncertainty as modest relative to the estimated payoff uncertainty.

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

Variance/Covariance Decompositions. Tables 12 and 13 report the estimated contribution of the shocks to, respectively, the variances and the co-variances of the key variables at business-cycle frequencies. (The confidence shock is omitted here, because its contributions were reported in the main text.) For comparison purposes, we also include the estimated contributions that obtain in the variants of the models that remove the confidence shock. Three findings are worth mentioning:

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

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

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

Figure 9: Theoretical IRFs, Part I

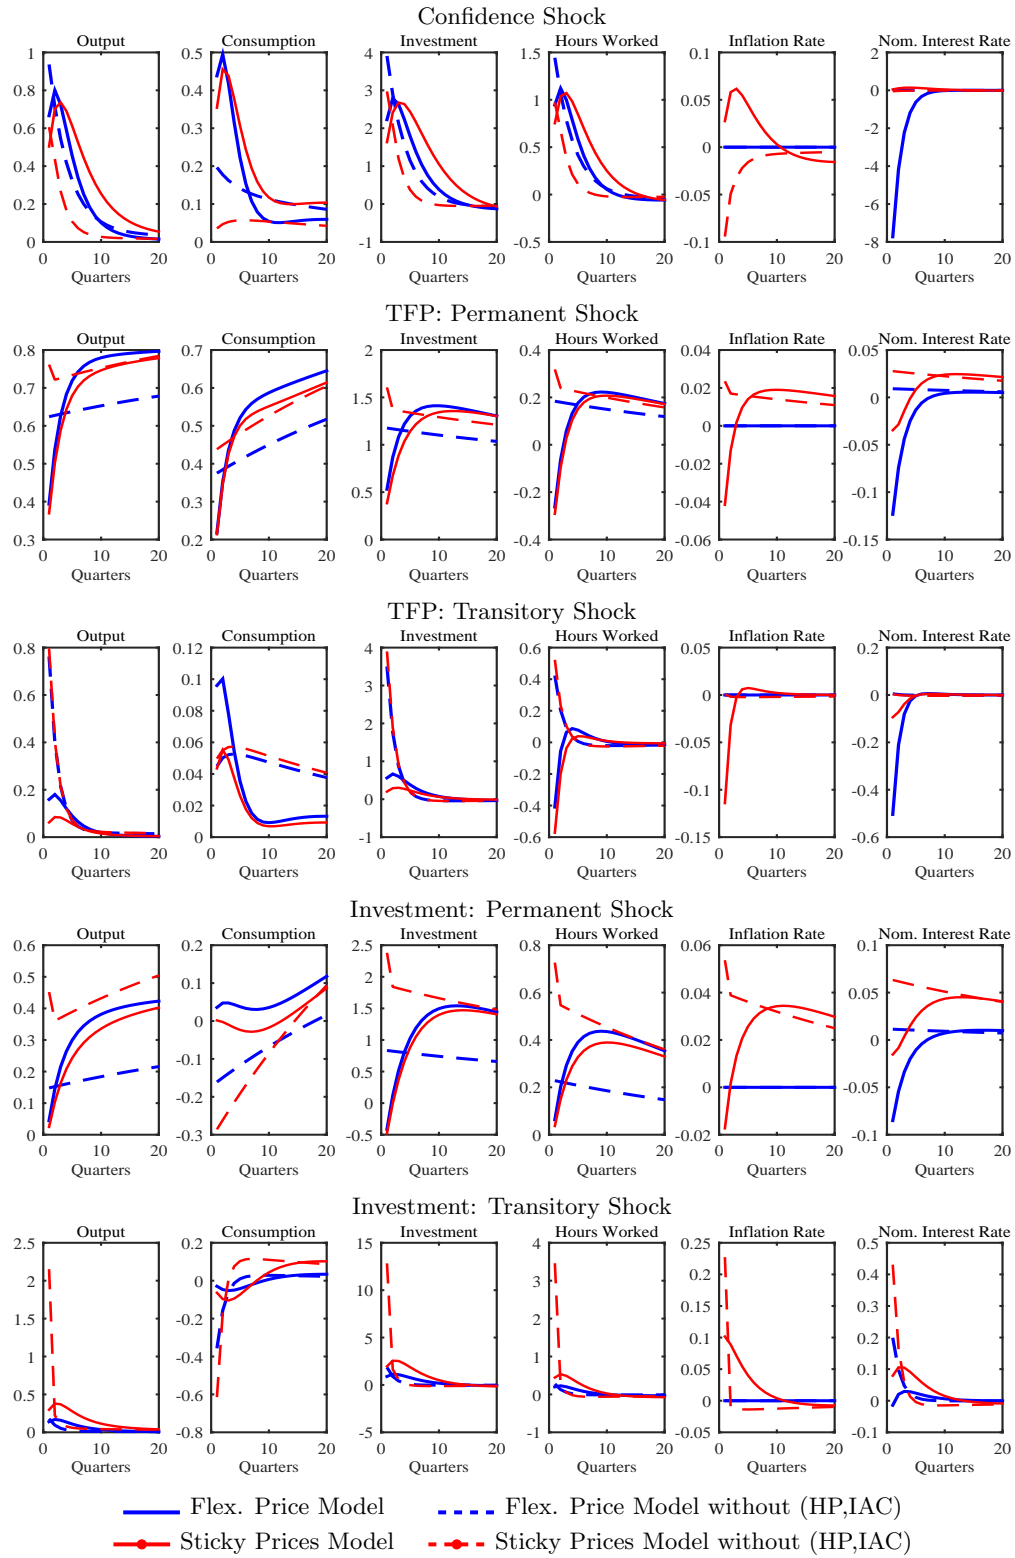


Figure 10: Theoretical IRFs, Part II

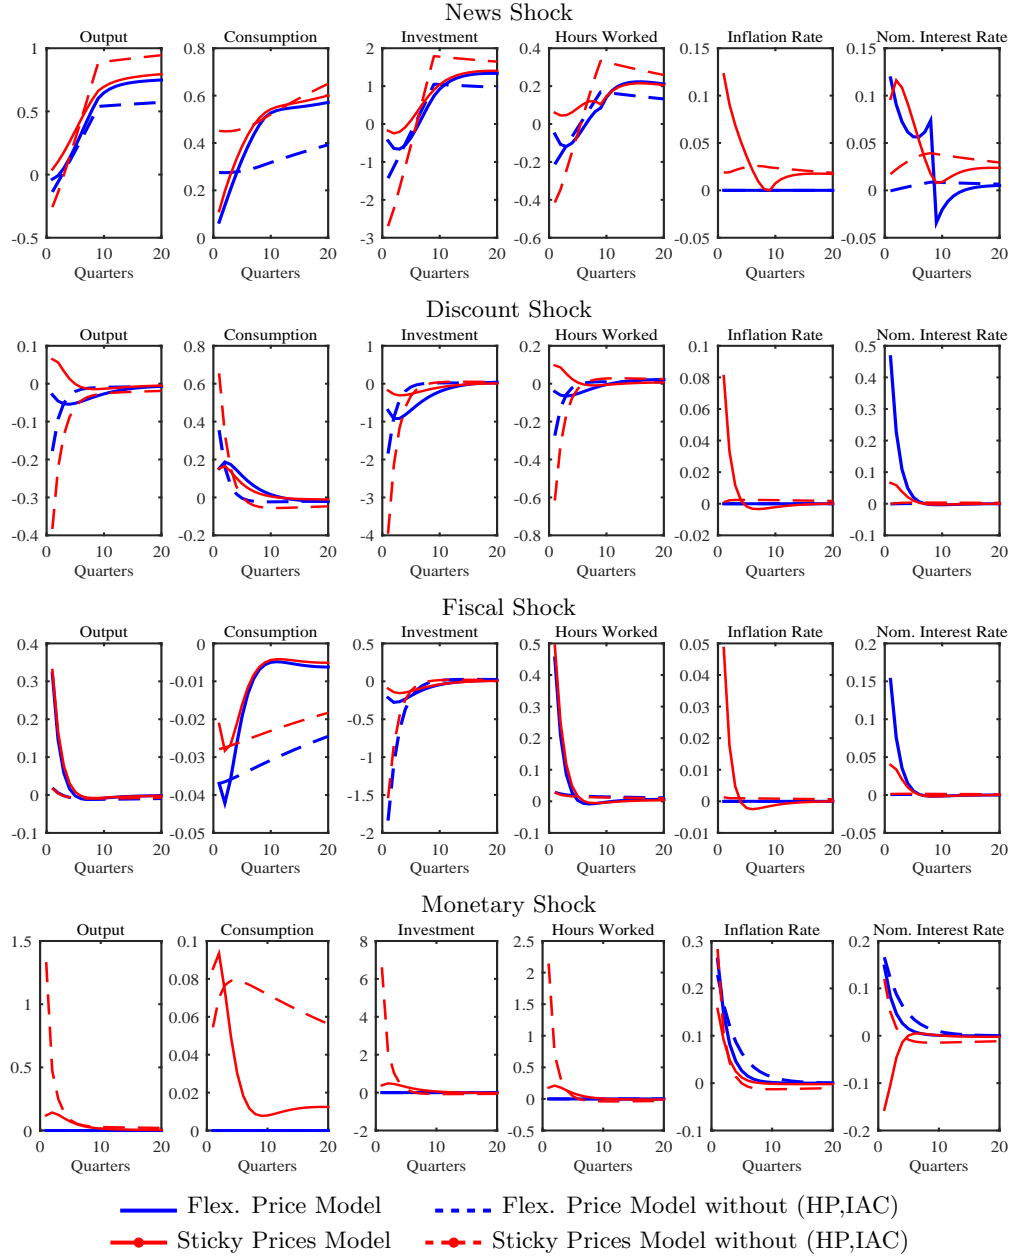


Table 12: Contribution of Shocks to Volatilities (6–32 Quarters)

	Y	C	I	h	π	R	Y	C	I	h	π	R
Permanent TFP Shock												
Flexible Price	23.73	28.95	6.55	4.26	0.00	0.03	13.43	16.12	12.09	3.36	0.00	0.05
Sticky Price	20.99	27.93	3.85	3.67	2.54	2.14	11.13	18.08	4.31	0.30	22.41	18.79
<i>In the absence of Belief shock</i>												
Flexible Price	5.05	15.28	0.81	9.55	0.98	0.99	13.17	9.00	20.29	11.95	19.78	19.86
Sticky Price	34.19	55.05	4.21	6.78	1.29	1.96	14.98	23.26	5.01	1.59	28.72	44.62
Transitory TFP Shock												
Flexible Price	2.45	1.72	2.95	2.72	0.00	0.34	0.26	7.77	6.40	0.36	0.00	0.33
Sticky Price	0.61	0.56	0.54	4.62	8.30	6.82	0.31	5.50	0.58	0.36	5.19	4.33
<i>In the absence of Belief shock</i>												
Flexible Price	0.91	0.05	4.70	0.31	2.99	2.99	0.04	0.04	1.98	0.04	1.51	1.53
Sticky Price	1.60	1.82	0.32	4.30	8.80	12.65	5.72	0.64	0.13	12.17	1.75	2.58
Permanent Investment Shock												
Flexible Price	4.04	0.43	7.56	5.93	0.00	0.02	2.82	0.32	0.56	4.01	0.00	0.04
Sticky Price	2.67	0.45	5.53	3.54	2.09	1.71	3.26	0.16	0.15	4.27	1.74	1.44
<i>In the absence of Belief shock</i>												
Flexible Price	66.55	65.11	53.22	64.33	62.65	64.02	9.74	7.20	12.93	9.45	8.68	8.68
Sticky Price	17.53	3.73	21.07	29.29	3.83	6.00	0.28	2.33	0.35	2.04	38.25	6.16
Monetary Policy Shock												
Flexible Price	2.29	0.98	9.26	3.33	0.00	0.00	0.00	0.00	0.00	0.00	100.00	0.04
Sticky Price	11.76	4.91	39.49	15.47	20.58	17.23	1.53	1.52	1.30	2.10	25.20	14.90
<i>In the absence of Belief shock</i>												
Flexible Price	4.54	3.32	6.07	4.38	1.79	1.78	0.00	0.00	0.00	0.00	1.62	0.14
Sticky Price	25.59	8.20	68.50	43.02	8.19	12.75	0.10	4.97	0.40	0.83	9.16	13.27

Table 13: Contribution of Shocks to Comovements (6–32 Quarters)

	(Y, h)	(Y, I)	(Y, C)	(h, I)	(h, C)	(I, C)		(Y, h)	(Y, I)	(Y, C)	(h, I)	(h, C)	(I, C)
	Covariance							Covariance					
Permanent TFP Shock													
Flexible Price	5.26	14.12	31.56	3.46	9.94	24.82		5.93	11.36	13.67	7.48	1.54	4.91
Sticky Price	3.62	10.12	33.13	2.56	8.49	27.93		1.50	5.76	16.19	0.91	2.18	4.94
<i>In the absence of Belief shock</i>													
Flexible Price	6.98	2.11	9.05	2.99	12.20	4.02		12.54	16.80	10.58	16.03	10.21	13.43
Sticky Price	15.42	14.16	109.07	5.86	285.22	107.50		7.39	8.57	38.57	3.38	69.88	17.71
Transitory TFP Shock													
Flexible Price	-0.69	3.06	2.43	-0.44	-1.44	3.85		0.36	1.38	-1.51	1.78	-2.72	-12.58
Sticky Price	-0.89	0.67	0.78	-0.60	-2.06	1.40		0.41	-0.24	1.56	-0.28	2.63	-4.22
<i>In the absence of Belief shock</i>													
Flexible Price	0.52	2.15	-0.08	1.22	-0.02	-0.27		0.04	-0.24	0.04	-0.24	0.04	-0.21
Sticky Price	-3.56	0.79	4.25	-0.27	-58.67	5.16		15.64	-0.55	-4.30	-0.73	-74.37	1.98
Permanent Investment Shock													
Flexible Price	5.92	5.59	0.10	6.62	0.42	-1.25		4.10	-0.83	-0.89	-1.04	-1.49	0.72
Sticky Price	3.72	4.01	-1.39	4.34	-2.41	-3.95		4.57	-0.39	-0.73	-0.45	-1.28	0.39
<i>In the absence of Belief shock</i>													
Flexible Price	65.79	62.03	67.80	62.64	65.38	66.96		9.65	11.69	8.62	11.83	8.34	10.96
Sticky Price	4.76	20.14	-16.86	22.89	119.30	-26.51		1.10	0.41	1.25	0.96	63.94	6.13
Monetary Policy Shock													
Flexible Price	3.33	5.26	-1.71	6.47	-2.78	-5.23		0.00	0.00	0.00	0.00	0.00	0.00
Sticky Price	16.33	25.21	-9.98	28.94	-17.53	-37.49		2.20	1.63	2.03	1.92	3.68	3.58
<i>In the absence of Belief shock</i>													
Flexible Price	4.48	5.47	4.00	5.53	3.85	5.11		0.00	0.00	0.00	0.00	0.00	0.00
Sticky Price	58.76	56.53	-33.34	68.13	-559.56	-203.03		0.49	-0.05	1.36	-0.23	54.25	-8.94

Connecting to the evidence of Section 3. Table 14 complements Figure 5 in the main text by comparing the variance decompositions that obtain when we apply our empirical strategy on artificial data from our two models with the one that we obtained in Section 3 when we run the same strategy on the US data (see Table 1). These decompositions, just like the IRFs in Figure 5, do not have a structural interpretation, because there is no one-to-one mapping between the theoretical shocks in our models and the identified factor. They nevertheless provide further evidence of the ability of our models to capture the type the volatility that is present in the data.

Table 14: Variance Contribution of Our Identified Factor: Models vs Data

	Y	C	I	h	Y/h	P_i	G	π	R
Identified Main Shock									
Data	49.66	21.90	61.26	49.58	13.24	6.33	18.01	34.08	29.05
Flexible Price	47.69	39.96	54.97	62.00	20.64	4.76	6.95	4.50	31.69
Sticky Price	42.58	18.52	51.28	58.11	37.48	4.46	8.35	36.98	39.01
<i>Without Filtering Permanent Shocks</i>									
Data	59.16	32.24	65.16	64.29	22.80	11.91	20.74	51.12	26.49
Flexible Price	77.51	55.30	76.02	62.33	14.82	1.93	27.90	3.07	33.45
Sticky Price	74.67	42.97	73.87	60.75	12.60	1.99	3.32	23.81	27.04
All Transitory Shocks									
Data	71.57	49.56	80.00	69.78	74.52	58.12	60.53	76.37	70.96
Flexible Price	58.69	62.27	67.22	72.87	68.71	19.76	61.03	92.73	84.68
Sticky Price	53.19	39.20	60.20	72.20	76.45	89.72	35.51	85.99	83.48

Posterior odds. Table 15 reports the posterior odds of four alternative models, starting from a uniform prior and estimating them on the real data only. The models differ on whether they assume flexible or sticky prices, and on whether they contain a confidence shock or not. Since we have dropped the nominal data for this exercise, the nominal parameters of the sticky-price models are now well identified. We have thus chosen to fix these parameters at the values that obtained when the models were estimated on both real and nominal data. We nevertheless re-estimate the preference and technology parameters and the shock processes in order to give a fair chance to each model to match the real data.

Table 15: Posterior Odds of Model A vs Model B

	Model A →	flex prices without	sticky prices with	without
Model B ↓				
flex prices, with confidence		0.00	0.98	0.02
flex prices, without confidence		–	1.00	0.90
sticky prices, with confidence		0.00	–	0.00

E. Log-Linear Solution

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

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

Thus consider the FOCs of the local planning problem we studied in Section 5 and log-linearize them around the deterministic steady state. Let a tilde (\sim) over a variable denote the log-deviation of this variable from its steady-state value. The log-linearized equilibrium conditions take the following form:

$$\tilde{y}_{it} = \tilde{a}_t + \alpha \tilde{k}_{it} + (1 - \alpha) \tilde{n}_{it} \quad (20)$$

$$(1 + \nu) \tilde{n}_{it} = \mathbb{E}_{it} \tilde{Y}_t - \mathbb{E}_{it} \tilde{c}_{it} \quad (21)$$

$$\mathbb{E}'_{it} \tilde{c}_{i,t+1} - \tilde{c}_{it} = (1 - \beta(1 - \delta)) \mathbb{E}'_{it} [\tilde{Y}_{t+1} - \tilde{k}_{i,t+1}] \quad (22)$$

$$\tilde{Y}_t = (1 - s) \tilde{c}_{it} + \tilde{s}_{it} \quad (23)$$

$$\tilde{k}_{i,t+1} = \delta \tilde{i}_{it} + (1 - \delta) \tilde{k}_{it} \quad (24)$$

where s is the investment share.

The interpretation of these equations should be familiar: (20) is the local production function; (21) is the optimality condition for labor; (22) is the Euler condition; (23) is the local resource constraint, with s denoting the saving rate in steady state; and finally (24) is the law of motion for local capital. The only peculiarity in the above system is the presence of two distinct expectation operators \mathbb{E}_{it} and \mathbb{E}'_{it} , which denote the local expectations, respectively, stage 1 and stage 2 of period t . The difference between these two expectation operators derives from the fact that islands form beliefs about one another's signals and thereby about Y_t in stage 1 on the basis of their misspecified priors, but observe the true state of nature and the true realized Y_t in stage 2. Along with the timing convention we have adopted, this explains why the first expectation shows up in the optimality condition for labor, while the second shows up in the optimality condition for consumption/saving.

At this state, it is important to keep in mind the following. The aggregate-level variables are, of course, obtained from aggregating individual-level variables. Since all islands are identical (recall that we are focusing on the limit with $\sigma \rightarrow 0$), the equilibrium values of the aggregate variables coincide with the equilibrium values of the corresponding individual variables. E.g., in equilibrium, it is ultimately the case that $y_{it} = Y_t$ for all i, t and all states of nature. This is because all islands receive the same signals and the same fundamentals. However, this does not mean that one can just replace the island-specific variables in the above conditions with the aggregate ones, or vice versa. This is for two reasons. First, when each island picks its local outcomes, it takes the aggregate

outcomes as given. Second, even though the “truth” is that all islands receive the same signals, in stage 1 of each period each island believes that the signals of other can differ from each own signal and, accordingly, reasons that y_{it} can differ from Y_t even when all other islands follow the same strategy as itself. Keeping track of this delicate matter is key to obtaining the (correct) solution to the model.

The same principle applies to the general class of DSGE models we consider below. Accordingly, the solution method we develop in this appendix deals with this delicate matter by (i) using appropriate notation to distinguish the signal received by each agent/island from either the average signal in the population or the true underlying shock to fundamentals; and (ii) choosing appropriate state spaces for both the individual policy rules and the aggregate ones.

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

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

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

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

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

where z_{it} is the signal received by agent i , ξ_t is the higher-order belief shocks, and Δ is a loading matrix. We next let \bar{z}_t denote the average signal in the economy and note that the “truth” is that

$z_{it} = \bar{z}_t = s_t$. Yet, this truth is publicly revealed only in stage 2 of period t . In stage 1, instead, each island believes, incorrectly, that

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

Note next that the stage-1 variables, y_{it} , can depend on the local signal z_{it} , along with the commonly-observed belief shock ξ_t and the backward-looking (predetermined) state variables x_{it}^b and X_t , but cannot depend on either the aggregate signal \bar{z}_t or the underlying fundamental s_t , because these variables are not known in stage 1. By contrast, the stage-2 decisions depend on the entire triplet (z_{it}, \bar{z}_t, s_t) . As already mentioned, the truth is that these three variable coincide. Nevertheless, the islands believe in stage 1 that the average signal can differ from either their own signal or the actual fundamental. Accordingly, it is important to write stage-2 strategies as functions of the three conceptually distinct objects in (z_{it}, \bar{z}_t, s_t) in order to do specify the appropriate equilibrium beliefs in stage-1. (Note that this is equivalent to expressing the stage-2 strategies as functions of the realized values of the stage-1 variables y and Y , which is the approach we took in the characterization of the recursive equilibrium in Section 5.) In what follows, we show how this belief structure facilitates a tractable solution of the aforementioned general DSGE model.

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

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

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

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

$$X_t = \Theta_X X_t^b + \Theta_s s_t + \Theta_\xi \xi_t,$$

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

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

$$\begin{aligned} y_{it} &= M_x(x_{it}^b - X_t) + M_X X_t + M_{EY} \mathbb{E}_{it} Y_t + M_f \mathbb{E}_{it}(x_{it}^f - X_t^f) + M_F \mathbb{E}_{it} X_t^f + M_s z_{it} \quad (25) \\ x_{it+1}^b &= N_x(x_{it}^b - X_t) + N_X X_t + N_y(y_{it} - Y_t) + N_Y Y_t + N_f(x_{it}^f - X_t^f) + N_F X_t^f + N_s s_t \quad (26) \end{aligned}$$

$$\begin{aligned} P_{f0} \mathbb{E}'_{it} x_{it+1}^f &= P_{f1}(x_{it}^f - X_t^f) + P_{F0} \mathbb{E}'_{it} X_{t+1}^f + P_{F1} X_t^f + P_x(x_{it}^b - X_t) + P_X X_t + \\ &+ P_{y0}(\mathbb{E}'_{it} y_{it+1} - \mathbb{E}'_{it} Y_{t+1}) + P_{Y0} \mathbb{E}'_{it} Y_{t+1} + P_{y1}(y_{it} - Y_t) + P_{Y1} Y_t + P_s s_t \quad (27) \end{aligned}$$

where

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

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

$$y_{it} = \Lambda_x^y(x_{it}^b - X_t) + \Lambda_X^y X_t + \Lambda_z^y z_{it} + \Lambda_\xi^y \xi_t \quad (28)$$

$$x_{it}^f = \Gamma_x^f(x_{it}^b - X_t) + \Gamma_X^f X_t + \Gamma_z^f z_{it} + \Gamma_{\bar{z}}^f \bar{z}_t + \Gamma_s^f s_t + \Gamma_\xi^f \xi_t \quad (29)$$

where the Λ 's and Γ 's are coefficients (matrices), whose equilibrium values are to be obtained in the sequel. Following our earlier discussion, note that the stage-2 policy rules are allowed to depend on the triplet (z_{it}, \bar{z}_t, s_t) , while the stage-1 policy rules are restricted to depend only on the local signal z_{it} . It is also useful to note that we would obtain the same solution if we were to represent the stage-2 policy rules as functions of y_{it} and Y_t in place of, respectively, z_{it} and \bar{z}_t : the latter two variables enter the equilibrium conditions that determine the stage-2 decisions, namely conditions (26) and (27), *only* through the realized values of the stage-1 outcomes y_{it} and Y_t .

Obtaining the solution. We obtain the solution in three steps. In step 1, we start by characterizing the equilibrium determination of the stage-1 policy rules, taking as given the stage-2 rules. Formally, we fix an arbitrary rule in (29); we assume that all islands believe that the stage-2 variables are determined according to this rule; and we then look for the particular rule in (28) that solves the fixed-point relation between y_{it} and Y_t described in (25) under this assumption. This step, which we can think of as the “static” component of the equilibrium, gives as a mapping from

Γ matrices to the Λ matrices. In step 2, we obtain a converse mapping by characterize the policy rules for the forward-looking variables that solve conditions (26) and (27) under the assumption that that the stage-1 outcomes are determined according to an arbitrary rule in (29). We can think of this step as solving for the “dynamic” component of the equilibrium. In step 3, we use the fixed-point between these two mappings to obtain the overall solution to the model.

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

Thus suppose that all islands follow a policy rule as in (29) and consider the beliefs that a given island i forms, under this assumption, about the stage-2 variables x_{it}^f and X_t^f . From (29), we have

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

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

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

which also implies that

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

Plugging the above in (25), the equilibrium equation for y_{it} , we get

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

Equivalently,

$$\begin{aligned} y_{it} &= (M_x + M_f \Gamma_x^f)(x_{it}^b - X_t) + (M_X + M_F \Gamma_X^f) X_t + M_{EY} \mathbb{E}_{it} Y_t \\ &\quad + (M_s + M_F(\Gamma_z^f + \Gamma_{\bar{z}}^f + \Gamma_s^f)) z_{it} + \left(M_F \Gamma_\xi^f + M_F \Gamma_{\bar{z}}^f \Delta + (M_F - M_f) \Gamma_z^f \Delta \right) \xi_t \end{aligned} \tag{30}$$

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

Thus suppose that this solution takes the form of a policy rule as in (28). If all other island follow this rule, then at the aggregate we have

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

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

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

Plugging this into (30), we obtain the following best response for island i :

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

For this to be consistent with our guess in (28), we must have

$$\Lambda_x^y = M_x + M_f \Gamma_x^f \tag{31}$$

$$\Lambda_X^y = (I - M_{EY})^{-1}(M_X + M_F \Gamma_X^f) \tag{32}$$

$$\Lambda_z^y = (I - M_{EY})^{-1} \left[M_s + M_F(\Gamma_z^f + \Gamma_{\bar{z}}^f + \Gamma_s^f) \right] \tag{33}$$

$$\Lambda_\xi^y = (I - M_{EY})^{-1} \left\{ M_F(\Gamma_\xi^f + \Gamma_{\bar{z}}^f \Delta) + (M_F - M_f)\Gamma_z^f \Delta + M_{EY}\Lambda_z^y \Delta \right\} \tag{34}$$

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

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

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

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

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

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

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

It follows that

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

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

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

For the aggregate next-period forward variables we have

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

and therefore

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

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

$$\begin{aligned} y_{it} - Y_t &= \Lambda_x^y(x_{it}^b - X_t) + \Lambda_z^y(z_{it} - \bar{z}_t) \\ x_{it}^f - X_t^f &= \Gamma_x^f(x_{it}^b - X_t) + \Gamma_z^f(z_{it} - \bar{z}_t) \\ Y_t &= \Lambda_X^y X_t + \Lambda_z^y \bar{z}_t + \Lambda_\xi^y \xi_t \\ X_t^f &= \Gamma_X^f X_t + (\Gamma_z^f + \Gamma_{\bar{z}}^f) \bar{z}_t + \Gamma_s^f s_t + \Gamma_\xi^f \xi_t \end{aligned}$$

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

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

Equivalently,

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

and hence

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

where

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

It follows that

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

or equivalently

$$\mathbb{E}'_{it}x^f_{it+1} = \Phi_x(x^b_{it} - X_t) + \Phi_X X_t + \Phi_z z_{it} + \Phi_{\bar{z}} \bar{z}_t + \Phi_s s_t + \Phi_\xi \xi_t \quad (35)$$

where

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

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

$$\mathbb{E}'_{it}X^f_{t+1} = \Phi_X X_t + \Phi_z z_{it} + \Phi_{\bar{z}} \bar{z}_t + \Phi_s s_t + (\Phi_\xi + \Gamma_z^f \Delta Q)\xi_t \quad (36)$$

With the above steps, we have calculated all the objects that enter the Euler condition (27). We can thus proceed to characterize the fixed-point relation that pins down the solution for the stage-2 policy rule.

To ease the exposition, let us repeat the Euler condition (27) below:

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

Use now (35) to write the left-hand-side of the Euler condition as

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

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

$$\begin{aligned}P_{f0}\mathbb{E}'_{it}x^f_{it+1} &= P_{f1} \left\{ \Gamma_x^f(x^b_{it} - X_t) + \Gamma_z^f(z_{it} - \bar{z}_t) \right\} + \\ &\quad + P_{F0} \left\{ \Phi_X X_t + (\Phi_z + \Phi_{\bar{z}})\bar{z}_t + \Phi_s s_t + (\Phi_\xi + \Gamma_z^f \Delta Q)\xi_t \right\} \\ &\quad + P_{F1} \left\{ \Gamma_X^f X_t + (\Gamma_z^f + \Gamma_{\bar{z}}^f)\bar{z}_t + \Gamma_s^f s_t + \Gamma_\xi^f \xi_t \right\} + P_x \left\{ x^b_{it} - X_t \right\} \\ &\quad + P_X X_t + P_{y0} \left\{ \Lambda_x^y \left(\Omega_x(x^b_{it} - X_t) + \Omega_z(z_{it} - \bar{z}_t) \right) - \Lambda_z^y \Delta Q \xi_t \right\} \\ &\quad + P_{Y0} \left\{ \Lambda_X^y (\Omega_X X_t + (\Omega_z + \Omega_{\bar{z}})\bar{z}_t + \Omega_s s_t + \Omega_\xi \xi_t) + \Lambda_z^y R_s t + (\Lambda_\xi^y + \Lambda_z^y \Delta)Q \xi_t \right\} \\ &\quad + P_{y1} \left\{ \Lambda_x^y(x^b_{it} - X_t) + \Lambda_z^y(z_{it} - \bar{z}_t) \right\} + P_{Y1} \left\{ \Lambda_X^y X_t + \Lambda_z^y \bar{z}_t + \Lambda_\xi^y \xi_t \right\} + P_s s_t\end{aligned}$$

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

$$P_{f0}\Phi_x = P_x + P_{f1}\Gamma_x^f + P_{y0}\Lambda_x^y\Omega_x + P_{y1}\Lambda_x^y \quad (37)$$

$$(P_{f0} - P_{F0})\Phi_X = P_{F1}\Gamma_X^f + P_X + P_{Y0}\Lambda_X^y\Omega_X + P_{Y1}\Lambda_X^y \quad (38)$$

$$P_{f0}\Phi_z = P_{f1}\Gamma_z^f + P_{y0}\Lambda_x^y\Omega_z + P_{y1}\Lambda_z^y \quad (39)$$

$$(P_{f0} - P_{F0})\Phi_{\bar{z}} = P_{F0}\Phi_z + (P_{F1} - P_{f1})\Gamma_z^f + P_{F1}\Gamma_{\bar{z}}^f + P_{Y0}\Lambda_X^y(\Omega_z + \Omega_{\bar{z}}) - P_{y0}\Lambda_x^y\Omega_z + (P_{Y1} - P_{y1})\Lambda_z^y \quad (40)$$

$$(P_{f0} - P_{F0})\Phi_s = P_{F1}\Gamma_s^f + P_{Y0}(\Lambda_X^y\Omega_s + \Lambda_z^y R) + P_s \quad (41)$$

$$(P_{f0} - P_{F0})\Phi_\xi = P_{F0}\Gamma_z^f\Delta Q + P_{F1}\Gamma_\xi^f + P_{Y0}\left\{\Lambda_X^y\Omega_\xi + \Lambda_\xi^y Q\right\} + (P_{Y0} - P_{y0})\Lambda_z^y\Delta Q + P_{Y1}\Lambda_\xi^y \quad (42)$$

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

Step 3. Steps 1 and 2 resulted in two systems of equations in the Λ and Γ matrices, namely system (31)-(34) and system (37)-(42). We now look at the joint solution of these two systems, which completes our guess-and-verify strategy and gives the sought-after equilibrium policy rules.

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

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

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

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

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

With the sum $\bar{\Gamma}_s^f \equiv \Gamma_z^f + \Gamma_{\bar{z}}^f + \Gamma_s^f$ determined as above, we can next obtain each of its three components as follows. First, Γ_s^f can be obtained from (41):

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

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

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

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

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

Plugging the definition of Φ_z and Ω_z in the above, we have

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

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

Consider now the matrices Λ_x^y and Γ_x^f . These are readily obtained from (31) and (37) once we replace the already-obtained results. It is also straightforward to check that these matrices correspond to the solution of the version of the model that shuts down all kinds of uncertainty but allows for heterogeneity in the backward-looking state variables (“wealth”).

To complete our solution, what remains is to determine the matrices Γ_ξ^f and Λ_ξ^y . These matrices solve conditions (34) and (42), which we repeat below:

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

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

$$\begin{aligned} \Omega_\xi &= N_Y\Lambda_\xi^y + N_F\Gamma_\xi^f \\ \Phi_\xi &= \Gamma_X^f(N_Y\Lambda_\xi^y + N_F\Gamma_\xi^f) + (\Gamma_\xi^f + \Gamma_z^f\Delta)Q \end{aligned}$$

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

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

where

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

Note that A , B , and C are determined by primitives, plus some of the coefficients that we have also characterized. The above equation therefore gives us the unique solution for the matrix Γ_ξ^f as a function of the primitives of the model. Λ_ξ^y is then readily obtained from (34).

This completes the solution.

References

- Alvarez, F., F. Lippi, and L. Paciello (2011): “Optimal Price Setting with Observation and Menu Costs,” *Quarterly Journal of Economics*, 126, 1909–1960.
- Angeletos, G.M., and J. La’O (2014): “Sentiments,” *Econometrica*, 81, 739–779.
- Bai, Y., V. Ríos-Rull, and K. Storesletten (2012): “Demand Shocks As Productivity Shocks,” mimeo.
- Barsky, R.B., S. Basu, and K. Lee (2014): “Whither News Shocks?,” NBER Macroeconomics Annual 2014.
- Barsky, R.B., and E.R. Sims (2011): “News Shocks and Business Cycles,” *Journal of Monetary Economics* 58, 273–289.
- Barsky, R.B., and E.R. Sims (2012): “Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence,” *American Economic Review*, 102, 1343–1377.
- Beaudry, P., and F. Portier (2006): “Stock Prices, News, and Economic Fluctuations,” *American Economic Review*, 96, 1293–1307.
- Beaudry, P., and F. Portier (2013): “Understanding Noninflationary Demand-Driven Business Cycles,” NBER Macroeconomics Annual 2013.
- Benati, L. (2014): “Do TFP and the Relative Price of Investment Share a Common I(1) Component?,” mimeo.
- Benhabib, J., and R. Farmer (1994): “Indeterminacy and Increasing Returns,” *Journal of Economic Theory*, 63, 19–41.
- Benhabib, J., P. Wang, and Y. Wen (2014): “Sentiments and Aggregate Fluctuations,” *Econometrica*, forthcoming.
- Bergmann, D., T. Heumann, and S. Morris (2014): “Information and Volatility,” mimeo.
- Bergmann, D., and S. Morris (2013): “Robust Predictions in Games of Incomplete Information,” *Econometrica*, 81, 1251–1308.
- Blanchard, O., J. L’Huillier, and G. Lorenzoni (2013): “News, Noise, and Fluctuations: An Empirical Exploration,” *American Economic Review*, 103, 3045–3070.
- Blanchard, O., and D. Quah (1989): “The Dynamic Effects of Aggregate Demand and Supply Disturbances,” *American Economic Review*, 79, 655–73.
- Bloom, N. (2009): “The Impact of Uncertainty Shocks,” *Econometrica*, 77, 623–685.
- Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta, and S. Terry (2012): “Really Uncertain Business Cycles,” NBER Working Paper 18245.

- Buera, F., and B. Moll (2012): “Aggregate Implications of a Credit Crunch,” NBER Working Paper 17775.
- Burnside, C., M. Eichenbaum, and S. Rebelo (2011): “Understanding Booms and Busts in Housing Markets,” mimeo.
- Caroll, C., (2001): “Macroeconomic Expectations of Households and Professional Forecasters,” *Quarterly Journal of Economics*, 118, 169–298.
- Cass, D., and K. Shell (1983): “Do Sunspots Matter?” *Journal of Political Economy*, 91, 193–227.
- Chari, V., P. Kehoe, and E. McGrattan (2007): “Business Cycle Accounting,” *Econometrica*, 75, 781–836.
- Chari, V., P. Kehoe, and E. McGrattan (2008): “Are Structural VARs with Long-Run Restrictions Useful in Developing Business Cycle Theory?,” *Journal of Monetary Economics*, 55, 1337–1352.
- Christiano, L., M. Eichenbaum, and C. Evans (2005): “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 113, 1–45.
- Christiano, L., M. Eichenbaum, and M. Trabandt (2014): “Understanding the Great Recession,” NBER Working Paper No. 20040.
- Christiano, L., R. Vigfusson (2005): “Maximum Likelihood in the Frequency Domain: the Importance of Time-to-Plan,” *Journal of Monetary Economics*, 50, 789–815.
- Christiano, L., C. Ilut, R. Motto, and M. Rostagno (2008): “Monetary Policy and Stock Market Boom-Bust Cycles,” ECB Working Paper No. 955.
- David, J., H. Hopenhayn, and V. Venkateswaran (2014): “Information, Misallocation and Aggregate Productivity,” mimeo.
- Eusepi, S., and B. Preston (2011): “Expectations, Learning and Business Cycle Fluctuations,” *American Economic Review*, 101, 2844–2872.
- Evans, G., and S. Honkapohja (2001): “*Learning and Expectations in Macroeconomics*,” Princeton University Press, Princeton, NJ.
- Farmer, R. (2012): “Confidence, Crashes and Animal Spirits,” *Economic Journal*, 122, 1–16.
- Fernald, J. (2014): “A Quarterly, Utilization-adjusted Series on Total Factor Productivity,” Federal Reserve Bank of San Francisco Working Paper 2012–19.
- Guesnerie, R. (1992): “An Exploration of the Eductive Justifications of the Rational Expectations Hypothesis,” *American Economic Review*, 82, 1254–1278.
- Huo, Z., and Takayama, N. (2014) “Rational Expectations Models with Higher Order Beliefs,” mimeo.

- Hansen, L., and T. Sargent (2007): “*Introduction to Robustness*,” Princeton University Press. Princeton: Princeton University Press.
- Ilut, C., and M. Schneider (2014): “Ambiguous Business Cycles,” *American Economic Review*, 104 (8), 2368–2399.
- Jaimovich, N., and S. Rebelo (2009): “Can News About the Future Drive the Business Cycle?,” *American Economic Review*, 99, 1097–1118.
- Justiniano, A., G. Primiceri, and A. Tambalotti (2010): “Investment Shocks and Business Cycles,” *Journal of Monetary Economics*, 57, 132–145.
- Lorenzoni, G., (2009): “A Theory of Demand Shocks,” *American Economic Review*, 99, 2050–2084.
- Mackowiak, B., and M. Wiederholt (2009): “Optimal Sticky Prices Under Rational Inattention,” *American Economic Review*, 99, 769–803.
- Mackowiak, B., and M. Wiederholt (2014): “Business Cycle Dynamics Under Rational Inattention,” mimeo.
- Mankiw, N. G., and R. Reis (2002): “Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve,” *Quarterly Journal of Economics*, 117, 1295–1328.
- Melosi, L. (2014): “Estimating Models with Dispersed Information”, *American Economic Journal: Macroeconomics*, 6, 1-31.
- Morris, S., and H. S. Shin (2002): “The Social Value of Public Information,” *American Economic Review*, 92, 1521–1534.
- Morris, S., and H. S. Shin (2001): “Rethinking Multiple Equilibria in Macroeconomics,” NBER Macroeconomics Annual 2000, 139-161, MIT Press.
- Nimark, C. (2013): “Dynamic Higher Order Expectations,” mimeo.
- Pavan, A. (2014): “Attention, Coordination, and Bounded Recall,” mimeo.
- Rondina, G., and T. Walker (2012): “Information Equilibria in Dynamic Economies with Dispersed Information,” mimeo.
- Sala, L., (2013): “DSGE Models in the Frequency Domain,” IGIER Working Paper 504.
- Sims, C., (2003): “Implications of Rational Inattention,” *Journal of Monetary Economics*, 50, 665–690.
- Scheinkmann, J., and W. Xiong (2003): “Overconfidence, Short-Sale Constraints and Bubbles,” *Journal of Political Economy* 111, 1183–1219.
- Smets, F., and R. Wouters (2007): “Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach,” *American Economic Review*, 97, 586–606.
- Uhlig, H. (2003): “What Drives GNP?,” mimeo.

- Woodford, M., (2002): “Imperfect Common Knowledge and the Effects of Monetary Policy,” in *Knowledge, Information, and Expectations in Modern Macroeconomics: in Honor of Edmund S. Phelps*, ed. by P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford. Princeton: Princeton University Press.
- Woodford, M., (2013): “Macroeconomic Analysis Without the Rational Expectations Hypothesis,” *Annual Review of Economics*, 5, 303–346.
- Weinstein, J., and M. Yildiz (2007): “A Structure Theorem for Rationalizability with Applications to Robust Predictions of Refinements,” *Econometrica*, 75, 365–400.