Sentiments*

George-Marios Angeletos
MIT and NBER

Jennifer La'O
Chicago Booth and NBER

November 15, 2012

Abstract

This paper develops a new theory of fluctuations—one that helps accommodate the notions of "animal spirits" and "market sentiment" in unique-equilibrium, rational-expectations, macroeconomic models. To this goal, we limit the communication that is embedded in a neoclassical economy by allowing trading to be random and decentralized. We then show that the business cycle may be driven by a certain type of extrinsic shocks which we call *sentiments*. These shocks formalize shifts in expectations of economic activity without shifts in the underlying preferences and technologies; they are akin to sunspots, but operate in unique-equilibrium models. We further show how communication may help propagate these shocks in a way that resembles the spread of fads and rumors and that gives rise to boom-and-bust phenomena. We finally illustrate the quantitative potential of our insights within a variant of the RBC model.

*This paper supersedes earlier working papers that circulated in 2008 and 2011 under the titles "Sentiments" and "Decentralization, Communication and the Origins of Fluctuation". For stimulating discussions and feedback, we are grateful to colleagues and to seminar participants at Bern, Chicago, Harvard, Michigan, MIT, Northwestern, NYU, PennState, Princeton, Ohio State, the Atlanta and New York FRBs, the ECB, and numerous conferences. We owe special thanks to Markus Brunnermeier, Tarek Hassan, John Hassler, Per Krusell, and Venky Venkateswaran for conference discussions of our paper, and to Rochelle Edge, Refet Gürkaynak and Burçin Kisacikoğlu for sharing their data with us. Last but certainly not least, we are grateful to the editor, Stephen Morris, and to four anonymous referees, for constructive criticisms and detailed suggestions that greatly improved the paper.

1 Introduction

Fluctuations in macroeconomic activity and asset markets are tied to aggregate shifts in market expectations. Consider, for example, the recent crisis. The earlier boom in housing markets has been attributed to exuberant beliefs about future prices; the subsequent bust came with a fast reversal in these beliefs; and the ongoing recovery is said to hinge on how quickly firms and households regain their confidence in the economy.

These observations are commonplace; they merely pinpoint to the apparent co-movement of market expectations and market outcomes. The challenge for the macroeconomist is to formalize, and then quantify, the deeper forces that might be driving this co-movement.

In the standard paradigm, these forces are modeled as exogenous random shocks to preferences and technology, the stock of capital, or other payoff-relevant fundamentals.[1] To many economists, this is unsatisfactory: shifts in market sentiment and aggregate demand often appear to obtain without obvious innovations in people's tastes and abilities, firms' know-how, and the like. Motivated by this conviction, a long tradition in macroeconomics has therefore sought to rationalize the observed fluctuations as the product of "animal spirits" in models that feature multiple equilibria, while another approach has opted to explain the same phenomena as departures from rationality.[2]

In this paper, we are motivated by the same theme but follow a different methodological route, shifting the focus on the communication and the coordination that is facilitated by the market mechanism. In particular, we show that as long as frictions in communication prevent agents from reaching exactly the same expectations about economic activity, aggregate fluctuations in these expectations may be driven by a certain type of extrinsic shocks which we call *sentiments*. These shocks are akin to sunspots, but operate in unique-equilibirium economies.

The main body of the paper develops the key ideas within a stylized Walrasian economy, while an extension illustrates the quantitative potential within a richer, RBC-like model. Moving beyond these particular models, the broader contribution is to show how extrinsic variation in market expectations and forces akin to animal spirits can be accommodated in the modern DSGE paradigm without abandoning the discipline of either rational expectations or equilibrium uniqueness.

Model. We consider a convex neoclassical economy in which agents are rational, markets are competitive, the equilibrium is unique, and there is no room for randomization devices. To sharpen our results, we also rule out aggregate shocks to preferences, technologies, or any other payoff-relevant fundamentals. More crucially, we depart from the standard paradigm by introducing a trading friction. This serves precisely two roles in our model: it introduces idiosyncratic trading uncertainty, and it limits the communication that takes place through markets or other means.

[1] To avoid confusion, let us fix some terminology. By "standard paradigm" we refer to the class of micro-founded, unique-equilibrium, rational-expectations, general-equilibrium models that have dominated academic research since the RBC revolution. By "fundamentals" we refer to any payoff-relevant variable, such as preferences, endowments, technologies, and government policies, or news thereof. Finally, by "extrinsic shocks" we refer to any residual, payoff-irrelevant, random variable.

[2] For the first approach, see, e.g, Azariadis (1981), Benhabib and Farmer (1994), Cass and Shell (1983), Diamond (1982), Cooper and John (1988), and Guesnerie and Woodford (1992); for the second, Akerlof and Shiller (2009).

The economy is thus split into multiple islands (Lucas, 1972), which are heterogeneous in terms of TFP, information, and trading opportunities. Each island specializes in production of a certain good but wishes to consume also the good of at least one other island; this gives rise to trade. Importantly, this trade is decentralized and takes place through random matching: in each period, each island meets and trades with only one other, randomly selected, island. Furthermore, certain employment and production choices are made in anticipation of these trading opportunities but before the observation of the actual terms of trade. Finally, communication is impeded in the sense that the islands may be unable to talk to one another or otherwise reach the same expectations about relevant economic outcomes, such as the terms of their trade, prior to their physical meeting.

These modeling choices seek to capture a simple but important fact. When a firm makes her employment and production decisions, she does not have the option to communicate with all the potential consumers whom she may meet and trade with later on, and whose decisions will ultimately determine the firm's own profitability. Similarly, consumers face uncertainty about the beliefs and intentions of other agents whose choices will ultimately determine their own employment and income. Of course, some communication does take place through markets, social networks, the media, and other means. However, this communication is far from perfect, leaving agents with diverse beliefs about current and future economic conditions. What is essential for our results is the imperfection of this kind of communication, not the precise details of how we model it.

Results. As with any other rational-expectations framework, the equilibrium of our economy is defined as the fixed point between market outcomes (actual allocations and prices) and market expectations (expectations of allocations and prices). Furthermore, any variation in these endogenous variables must ultimately be driven by some sort of exogenous shock. The question of interest for us, as for the literature on coordination failures and sunspot fluctuations, is whether the equilibrium variation in market expectations is spanned by the variation in exogenous payoff-relevant variables and beliefs thereof, or whether there is also some residual, extrinsic variation.

Theorem 1 establishes that the aforementioned fixed point exists and is unique, ruling out the usual formalization of self-fulfilling fluctuations. Theorem 2 establishes that extrinsic variation in market expectations is nevertheless possible as long as these expectations remain imperfectly aligned across different agents—which, in turn, can be true as long as communication is imperfect.

To understand this result, take any two islands i and j that are about to meet and trade. Next, note that the output of each island is pinned down by the local preferences and technologies, and the local belief about the upcoming terms of trade: other things equal, an island produces more if it expects its terms of trade to improve. Finally, consider the following question: can there exist states of Nature in which both islands expect their terms of trade to improve?

Clearly, this cannot be the case if communication is perfect: if island i expects its terms to improve, and if both islands share the same beliefs about market outcomes, then island j must expect its own terms to deteriorate. As we show in Theorem 1, this logic guarantees that, whenever equilibrium expectations are homogeneous across agents, actual macroeconomic outcomes are pinned down by the underlying fundamentals, even if the latter are not per se known.

Now consider the case where communication is imperfect, so that the two islands are holding heterogeneous beliefs about the terms of their trade. This means that there can exist states of Nature in which they both expect their terms to improve, as well as states of Nature in which they both expect their terms to deteriorate. What is more, these events can be correlated in the cross-section of the economy, giving rise to aggregate fluctuations.

During a boom, each island produces more because it expects its trading partner to produce more and hence the demand for its own product to increase. During a recession, each island expects its demand to be low and acts in a way that drives down the demand for other islands. These fluctuations therefore have the same flavor, and the same empirical content, as the self-fulfilling fluctuations that obtain in models with multiple equilibria.

What drives these fluctuations is a particular kind of aggregate shocks, which we call "sentiment shocks". These shocks impact the information that is available to each island, without however affecting the latter's beliefs either about the aggregate fundamentals (which are fixed) or about the idiosyncratic fundamentals of its trading partner (which are random). In this sense, these shocks are extrinsic. These shocks nevertheless impact *equilibrium* expectations, because they effectively alter the equilibrium belief that each island forms about the choices of other islands. One can thus think of, say, a positive sentiment shock as a shock that rationalizes the optimism of one island by making that island receive news (signals) that other islands are themselves optimistic.

These shocks can thus also be understood as shocks to higher-order beliefs. By imposing that the *aggregate* fundamentals are fixed and common knowledge, we rule out the particular type of higher-order uncertainty that has been the focus of previous work (e.g., Morris and Shin, 2002, 2003, Woodford, 2003). Nevertheless, by introducing trading frictions and imperfect communication, we open the door to higher-order uncertainty at the *micro* level: when two islands are matched together, they are uncertain, not only about each other's productivities, but also about each other's beliefs of their productivities, each other's beliefs of their beliefs of their productivities, and so on. The fluctuations we document reflect correlated variation in this kind of higher-order beliefs.

That being said, we prefer to interpret our sentiment shocks as shocks to *first-order* beliefs of *endogenous* economic outcomes. In the theory, agents never *need* to form higher-order beliefs. Rather, they need only to form first-order beliefs of the relevant equilibrium allocations and prices. Furthermore, surveys contain evidence merely on this kind of first-order beliefs. Finally, there are multiple specifications of the belief hierarchy that are consistent with the same joint distribution for the model's equilibrium outcomes, which means that the former cannot be uniquely identified by data on the latter. By contrast, what can be identified is the extrinsic variation in first-order beliefs of economic activity—this is what we are after in this paper.

Complementing this perspective, we argue that correlation in the relevant expectations may emerge endogenously as agents learn from realized market outcomes or otherwise exchange their beliefs. Furthermore, we show that such communication may serve as a powerful propagation mechanism for the type of fluctuations we formalize—leading to contagion effects akin to the spread of fads and rumors, and giving rise to boom-and-bust cycles like those experienced in recent years.

Finally, to illustrate the broader applicability and the quantitative potential of our theory, we embed a tractable variant of our sentiment shocks in the RBC framework. We then show that our theory appears to have no serious difficulty in matching key business-cycle facts such as the co-movement of employment, output, consumption, and investment, or the cyclicality of measured labor wedges and output gaps.

Layout. The rest of the paper is organized as follows. Section 2 introduces our framework. Section 3 characterizes the equilibrium. Section 4 contains our main results regarding the possibility of extrinsic fluctuations. Section 6 shows how communication helps generate fad dynamics and boom-and-bust cycles. Section 7 explores the quantitative potential. Section 8 concludes. Appendices A and B contain the proofs and a detailed analysis of the model of Section 7.

2 Model

The economy consists of a continuum of islands, indexed by $i \in \mathcal{I} = [0,1]$. Each island is populated by a representative household and a representative, locally-owned firm. All agents are price-takers. Each island produces a single good, which can either be consumed at "home" or be traded for a good produced "abroad" (by some other island). Production exhibits constant returns to scale with respect to local labor, which is supplied elastically by the local household, and local land, which is in fixed supply. Time is discrete, indexed by $t \in \{0, 1, ...\}$, and each period contains two stages. Employment and production are set in stage 1, while trading and consumption occur in stage 2. Finally, and importantly, trading takes place through random pair-wise matching.

Firms and technologies. Consider the firm of island i. Its technology is given by

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

where y_{it} is the quantity produced, A_i is the local total factor productivity (TFP), n_{it} is the labor input, k_{it} is the land input, and $\theta \in (0,1)$ parameterizes the income share of labor. The profit of this firm is $\pi_{it} = p_{it} y_{it} - w_{it} n_{it} - r_{it} k_{it}$, where p_{it} denotes the local price of the local good, w_{it} denotes the local wage, and r_{it} the local rental rate of land.

TFP varies across islands but not over time, thus ruling out both aggregate and idiosyncratic shocks. The cross-sectional distribution of TFP is described by a p.d.f. $\mathcal{F}_A : \mathcal{A} \to (0,1)$, where $\mathcal{A}$ is a compact subset of $\mathbb{R}_+$. This distribution is invariant over time and common knowledge—and so is the exact mapping from the identity i of a particular island to its idiosyncratic productivity A_i.

Households and preferences. Preferences on island i are given by

$$\mathcal{U}_i = \sum_{t=0}^{\infty} \beta^t \left[U\left(c_{it}, c_{it}^*\right) - V\left(n_{it}\right) \right]$$

where $\beta \in (0,1)$ is the discount factor, $c_{it} \in \mathbb{R}_+$ and $c_{it}^* \in \mathbb{R}_+$ are the consumptions of, respectively, the "home" and the "foreign" good,[3] $U(c_{it}, c_{it}^*)$ is the utility flow from these two forms of consumption,

[3] To have well-defined preferences over the entire commodity space, we can think of the home agents as either being indifferent among the goods of all other islands, or as liking only the good of their current random match.

$n_{it} \in \mathbb{R}_+$ is labor supply, and $V(n_{it})$ is the implied disutility. U and V are given by

$$U(c, c^*) = \left(\frac{c}{1-\eta}\right)^{1-\eta} \left(\frac{c^*}{\eta}\right)^{\eta} \qquad \text{and} \qquad V(n) = \frac{n^{\epsilon}}{\epsilon},$$

where $\eta \in (0,1)$ parameterizes the extent to which there is specialization and trade (the fraction of "home" expenditure that is spent on the "foreign" good), while $\epsilon > 1$ parameterizes the Frisch elasticity of labor supply. Finally, the period-t budget constraint is given by

$$p_{it} c_{it} + p^*_{it} c^*_{it} \leq w_{it} n_{it} + r_{it} K + \pi_{it} \tag{2}$$

where p_{it} and p^*_{it} denote the local prices of, respectively, the home and the foreign good, and K is the fixed endowment of land. In equilibrium, $k_{it} = K$.

Matching, timing, and information. To simplify, the matching is assumed to be uniform and i.i.d. over time: each island has an equal probability of being matched with any other island. Nature draws all the matches at the beginning of time, but does not reveal who is matched with whom and when. Thus fix a period t and a pair of islands that have been matched together in that period. In stage 2, the two islands meet, figure out that they were in the same match, and trade. The two islands, however, choose their employment and production levels in stage 1, *before* observing either their identities or the terms of their trade. Key economic decisions are thus made in anticipation of future trading opportunities, and with incomplete information about these opportunities.

Our results do not depend on the precise details of how we model the information structure. To be concrete, however, we will assume (i) that exogenous information arrives only in stage 1 of each period and (ii) that every island shares its information with its trading partner once the two meet in stage 2. The flow of information and the timing of choices are thus as in the following figure.

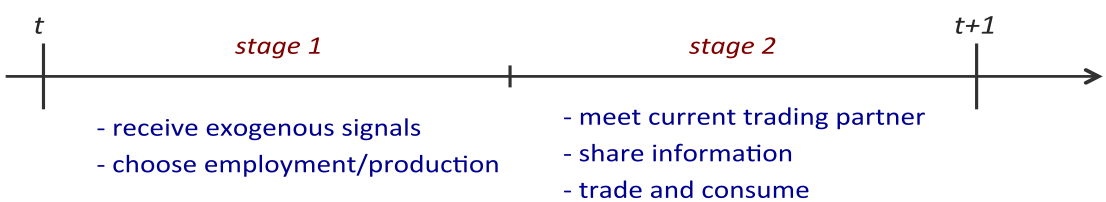

More formally, for each t, we fix a compact set $\mathcal{X}_t \subset \mathbb{R}^n$ and let x_{it} be a random variable drawn from $\mathcal{X}_t$. This variable represents the signal(s) that island i receives in stage 1 of that period and can be quite arbitrary. For instance, it may contain information, not only about the TFP of i's trading partner, but also about the information that the latter has acquired either by Nature or by past trades. We will consider specific examples in due course. For now, we only impose a certain form of symmetry: the signal received by a particular island does not depend per se on either its own "name" or the precise identities of its trading partners. It follows that all the relevant information that is available to an island in stages 1 and 2 of period t can be summarized in, respectively, the variables $\omega_{it} \in \Omega_t$ and $z_{it} \in \mathcal{Z}_t$, which are defined recursively as follows: for all $t \geq 0$, $\omega_{i,t} = (z_{i,t-1}, x_{i,t})$ and $z_{it} = (\omega_{it}, \omega_{m_t(i),t})$, where $m_t(i)$ henceforth denotes i's match in period t and where $z_{i,-1} \equiv A_i$.

That is, information sets (or "types", or "local states") are updated either by the arrival of exogenous signals in stage 1 or by the endogenous information exchange during stage 2.[4]

Sentiment shocks. The joint distribution of the signals x_{it} in the population of islands is allowed to depend on an exogenous random variable ξ_t drawn from a compact set $\Xi \subset \mathbb{R}^n$. This variable is akin to a sunspot in the sense that it affects information sets without affecting either the true aggregate fundamentals or any agent's beliefs about these fundamentals (for the latter are fixed and common knowledge). As will become clear in due course, we can further refine the notion that this variable is extrinsic by imposing that variation in ξ_t does not cause variation in any island's belief about the TFP level of either its own current and future trading partners, or of any other match in the economy. This variable will thus permit us to introduce aggregate variation in beliefs of equilibrium outcomes without any variation in beliefs of fundamentals. To fix language, we refer to ξ_t as a "sentiment shock". The history of this shock is denoted by $\xi^t \equiv (\xi_1, ...\xi_t)$.

Asset markets. When two islands meet, they trade their specialized goods, but are not allowed to trade any financial assets. Given the specification of matching we have assumed, this is without any loss of generality: since the probability that these islands will meet again the future is zero, such trading would not take place even if it were allowed. But even if we were to modify the model so that the aforementioned probability is non-zero, the essence of our results would not change: such trades would facilitate risk-sharing, but would not eliminate the communication friction.[5]

Equilibrium definition. The underlying probability space of our model is quite rich, as it involves the realizations of all matches and signals in the population. For our purposes, however, it suffices to focus on the joint distribution of the history ξ^t of the sentiment shock and of the pair of information sets $(\omega_{it}, \omega_{jt})$ of an arbitrary match (i, j). We assume that this distribution is represented by a continuous probability density function, which we henceforth denote by $\mathcal{P}_t(\omega_{it}, \omega_{jt}, \xi^t)$. Next, note that any allocation and price system can be represented with a collection of functions $\{n_t, k_t, y_t, w_t, r_t, p_t, p_t^*, c_t, c_t^*\}_{t=0}^{\infty}$ such that, for all islands, dates, and possible states, $n_{it} = n_t(\omega_{it})$, $k_{it} = k_t(\omega_{it})$, $y_{it} = y_t(\omega_{it})$, $w_{it} = w_t(\omega_{it})$, $r_{it} = r_t(\omega_{it})$, $p_{it} = p_t(z_{it})$, $p_{it}^* = p_t^*(z_{it})$, $c_{it} = c_t(z_{it})$, and $c_{it}^* = c_t^*(z_{it})$, with $z_{it} = (\omega_{it}, \omega_{jt})$ and $j = m_t(i)$.[6] We require that these functions be continuous; this guarantees that all relevant expectations are well defined and permits us to apply the contraction mapping theorem to prove existence and uniqueness of the equilibrium. Modulo these qualifications, a competitive equilibrium is defined in an otherwise conventional manner.

Definition 1. *An equilibrium is a collection of continuous allocation and price functions such that (i) given current prices and expectations of future prices, the allocations are optimal for households and firms; (ii) prices clear all markets; and (iii) expectations are rational.*

[4] Accordingly, the sets Ω_t and $\mathcal{Z}_t$ are compact and constructed recursively by letting $\mathcal{Z}_{-1} = \mathcal{A}$ and $\Omega_t = \mathcal{Z}_{t-1} \times \mathcal{X}_t$ and $\mathcal{Z}_t = \Omega_t \times \Omega_t$ for any $t \geq 0$.

[5] What would, of course, eliminate the friction is the introduction of complete and *centralized* markets, for then all relevant information would get perfectly aggregated (Grossman, 1981). Our notion of extrinsic fluctuations hinges on departing from this unrealistic extreme, but not on the precise details of how this departure takes place.

[6] Note that the price functions p_t and p_t^* must satisfy $p_t(\omega, \omega')/p_t^*(\omega, \omega') = p_t^*(\omega', \omega)/p_t(\omega', \omega)$ for all $\omega, \omega' \in \Omega_t$. This simply means that any two islands that trade face, of course, the same the terms of trade.

3 Equilibrium characterization

We now characterize the equilibrium.[7] Consider first the consumption decisions of the household of island i during stage 2 of period t. Let λ_{it} denote the Lagrange multiplier on its budget and normalize the local nominal prices so that $\lambda_{it} = 1$. Optimal consumption choices satisfy

$$U_c\left(c_{it}, c_{it}^*\right) = p_{it} \qquad \text{and} \qquad U_{c^*}\left(c_{it}, c_{it}^*\right) = p_{it}^*. \tag{3}$$

By trade balance, $p_{it}^* c_{it}^* = p_{it}(y_{it} - c_{it})$. By market clearing, $c_{it} + c_{jt}^* = y_{it}$. Combining these conditions with the corresponding ones for i's trading partner (denoted here by j), and using the Cobb-Douglas specification of U, we obtain the following:

$$c_{it} = (1-\eta)y_{it}, \qquad c_{it}^* = \eta y_{jt}, \qquad \text{and} \qquad p_{it} = y_{it}^{-\eta} y_{jt}^{\eta}. \tag{4}$$

The interpretation of these results should be familiar from international trade theory: a fraction $1-\eta$ of the good of each island is consumed at "home", while the rest is "exported"; and the terms of trade increase with the "foreign" supply relative to the "home" one.[8]

Consider now the labor-supply and labor-demand decisions that the local household and the local firm take during stage 1 of period t. These are given by the following first-order conditions:

$$V'(n_{it}) = w_{it} \qquad \text{and} \qquad w_{it} = \mathbb{E}_{it}\left[p_{it}\right] \theta \frac{y_{it}}{n_{it}}, \tag{5}$$

where $\mathbb{E}_{it}[\cdot]$ is a short-cut for the rational expectation conditional on ω_{it}. In words, workers equate the wage with the expected marginal disutility of effort, while firms equate the wage with the expected marginal revenue product of labor. It follows that the local marginal disutility of labor is equated with expected local marginal revenue product of labor.

This last finding means that we can understand the *local* equilibrium of any given island as the solution to the problem of a (benevolent) local planner that takes as given the local beliefs of terms of trade. The *general* equilibrium is then pinned down by requiring that these beliefs are consistent with the local equilibrium behavior of each island, that is, by requiring p_{it} to satisfy (4).

Proposition 1. *The equilibrium production levels and the equilibrium terms of trade solve the following fixed-point problem:*

$$y_t(\omega) = \left(\theta^{\vartheta} A_t(\omega) K^{1-\theta}\right)^{\frac{1}{1-\vartheta}} \left(\int_{\Omega_t} p_t(\omega,\omega')\mathcal{P}_t(\omega'|\omega)d\omega'\right)^{\frac{\vartheta}{1-\vartheta}} \tag{6}$$

$$p_t(\omega,\omega') = y_t(\omega)^{-\eta} y_t(\omega')^{\eta} \tag{7}$$

where $\vartheta \equiv \frac{\theta}{\epsilon} \in (0,1)$, $A_t(\omega)$ identifies the productivity of an island of type $\omega \in \Omega_t$, and $\mathcal{P}_t(\omega'|\omega)$ is the probability that this island attaches to meeting an island of type $\omega' \in \Omega_t$.

[7] The derivations that follow presume that an equilibrium allocation exists and is interior. Interiority follows for the Inada conditions on V and on the technology; existence is verified along with uniqueness in Theorem 1.

[8] Island i's terms of trade are given by $R_{it} \equiv \frac{p_{it}}{p_{it}^*}$ (the ratio of "export" to "import" prices). Note then that $R_{it} = \frac{y_{jt}}{y_{it}} = p_{it}^{1/\eta}$. Since this is an increasing function of p_{it}. we henceforth interpret p_{it} also as the terms of trade.

Proposition 1 is an example of the fixed-point relation between equilibrium outcomes and equilibrium expectations that is endemic to any rational-expectations economy. This fixed point is particularly simple here, and is essentially static because there is no savings technology or other intertemporal payoff linkages. However, as illustrated by the dynamic variant that we study in Section 6, our insights apply more generally.

Interestingly, this fixed point can also be understood as the perfect Bayesian equilibrium of a *fictitious* game among the islands. To see this, substitute (7) into (6) to get the following:

$$\log y_{it} = (1-\alpha) f_i + \alpha\, \mathbf{E}_{it}[\log y_{jt}], \tag{8}$$

where $f_i \equiv \frac{1}{1-\vartheta} \log\left(\theta^{\vartheta} A_i K^{1-\theta}\right)$ summarizes i's fundamentals, $\alpha \equiv \frac{\eta}{\eta+(1-\vartheta)/\vartheta} \in (0,1)$ is a scalar that is pinned down by preference and technology parameters, and $\mathbf{E}_{it}$ is an adjusted expectation operator defined by $\mathbf{E}_{it}[X] \equiv H^{-1}\left(\mathbb{E}_{it}\left[H(X)\right]\right)$, with $H(X) \equiv \exp(\eta X)$, for any random variable X. It follows that we can represent our economy as a game in which the players are the islands (or their local planners), their choices are their output levels, their best responses are described by (8), and the coefficient α is, in effect, the degree of strategic complementarity.

This game-theoretic interpretation reveals an important connection between our micro-founded business-cycle economy and the class of more abstract coordination games studied by Morris and Shin (2002) and Angeletos and Pavan (2007): it is *as if* the islands are trying to coordinate their production choices. We will revisit this connection in Sections 4.3 and 4.4. For now, we note that conventional general-equlibrium effects are the sole origin of what looks like strategic interaction: our model is a Walrasian economy, not a game; the actual agents (firms and households) are infinitesimal price-takers, not strategic players; and the interdependence of incentives across islands is a by-product of competitive market interactions, not a symptom of production externalities, market failures, and the like.[9] Indeed, the kind of strategic interdependence that is stylized by (8) is endemic to the market mechanism: the choices of any given firm or consumer hinge of her expectations of future market conditions, which in turn hinge on the choices of other firms and consumers.

Putting aside these interpretations, we can show that condition (8) defines a contraction mapping over the space of continuous functions that map the local state of an island, $\omega_{it} \in \Omega_t$, to its equilibrium output, $y_{it} \in \mathbb{R}_+$. The following is then immediate.

Theorem 1. *The equilibrium exists and is unique.*

The proof of this result rests on the assumption that Ω_t is compact. Without this, we cannot generally guarantee existence. Yet, whenever an equilibrium exists, it has to be unique by the fact that $\alpha \in (0,1)$. In the closed-form examples we consider in the sequel, Ω_t is not compact, but the unique equilibrium is obtained by guessing and verifying.

[9] The property that output levels are strategic complements ($\alpha > 0$) rather than strategic substitutes ($\alpha < 0$) hinges on the partial-equilibrium property that an island's employment and output responds positively to its terms of trade. This property is hard-wired in our model by the particular preference specification we have assumed, but seems to be the empirically relevant case even if we move outside our model. Furthermore, even if α had been negative, the possibility of extrinsic fluctuations would remain. The only difference is that in this case pessimistic beliefs would stimulate labor supply and economic activity, which seems counterfactual.

4 Extrinsic Fluctuations

We now proceed to study whether the equilibrium can exhibit extrinsic fluctuations, that is, whether economic outcomes can vary with the sunspot-like shock ξ_t.

As we show below, answering this question does not require one to know the precise details of the information structure. Rather, it suffices to inspect the *endogenous* expectations that the different islands end up forming about one another's level of economic activity. With this in mind, we define perfect communication as follows.

Definition 2. *The economy exhibits perfect communication if and only if the following property holds along the unique equilibrium: for any period t, any state of nature, and any given match, the two islands within that match share the same belief about each other's output levels.*

Accordingly, we find it useful to introduce the following aggregate measure of the relevant expectations. Let $\mathbb{E}_{it}[\log y_{jt}]$ measure an island's forecast of the output of its trading partner and let $\log \mathcal{B}_t$ be the average of these forecasts in the cross-section of islands. We henceforth interpret $\mathcal{B}_t$ as a proxy of the average optimism or pessimism in the economy.[10] Finally, we measure aggregate output, Y_t, by the logarithmic average of local output in the cross-section of islands.[11]

We can then state our key result as follows.

Theorem 2. *Along the unique equilibrium, aggregate output Y_t and the average expectation $\mathcal{B}_t$ can vary with the extrinsic shock ξ_t if and only if communication is imperfect.*

In the remainder of this section we prove this theorem in two steps, starting with the "only if" part and then proving the "if" by specific example. In the next section we then proceed to discuss the broader insight and its empirical content.

Perfect communication. Since each island knows its own output, if two islands have reached the same beliefs about their output levels, they must *know* each other's output. It follows that we can drop the expectation operator in condition (8) and solve for the equilibrium output of the two islands as a function of the local fundamentals. Aggregating across islands then gives the following.

Proposition 2. *When communication is perfect, Y_t and $\mathcal{B}_t$ are invariant to ξ_t. Furthermore,*

$$\log Y_t \;=\; \log \mathcal{B}_t \;=\; \kappa + \tfrac{1}{1-\vartheta}\bar{a} \tag{9}$$

where $\bar{a} \equiv \int \log A\, \mathcal{F}_A(A) dA$ measures aggregate TFP and $\kappa \equiv \frac{1}{1-\vartheta}\log(\theta^{\vartheta} K^{1-\theta})$.

To reach this result, we have effectively imposed that the islands reach common knowledge of each other's output levels, which in turn implies that they also reach common knowledge of each other's TFP levels. This property can be relaxed without affecting the essence of the result. In

[10] Another plausible proxy is the average forecast of *aggregate* output. The results we present in the sequel do not hinge on which of the two proxies one uses for empirical purposes; see footnote 15.

[11] Whenever we refer to the cross-sectional average of some island-specific variable, we mean the expectation of that variable conditional on the aggregate state. For example, $\log Y_t \equiv \int \log y_{it} di \equiv \int y_t(\omega) \mathcal{P}(\omega|\xi^t) d\omega$.

particular, to get the above result it suffices to impose that $\log \mathbb{E}_{it} p_{it} = -\log \mathbb{E}_{jt} p_{jt}$. Intuitively, this means that whenever an island expects its terms of trade to improve, its partner expects the exact opposite. As long as this is true, the joint output of the two islands continues to be pinned down by their fundamentals, even though the islands may have not reached common knowledge of them. Finally, the above result extends directly to the case of aggregate TFP shocks, irrespectively of the information that the islands might have about these shocks.[12] These facts underscore that the key issue is the beliefs that agents form about the relevant equilibrium outcomes (output levels or terms of trade), not the information they may, or may not, have about one another's fundamentals.

Imperfect communication. We now show that allowing for heterogeneity in the aforementioned kind of beliefs opens the door to extrinsic fluctuations. The example we use for this purpose is intentionally hard-wired; it also rules out any persistence in equilibrium beliefs by imposing that the sentiment shock is i.i.d. over time.[13] A discussion of the broader insight follows in Section 5; an example with richer belief dynamics is developed in Section 6.

The land endowment is normalized to $K = 1$. The cross-sectional distribution of TFP is log-normal: $\log A_i \sim \mathcal{N}(0, \sigma_A^2)$, $\sigma_A > 0$. The extrinsic shock is i.i.d Normal over time: $\xi_t \sim \mathcal{N}(0, \sigma_\xi)$, $\sigma_\xi > 0$. Finally, the exogenous signal received by i is given by the pair $x_{it} = (x_{it}^1, x_{it}^2)$, where

$$x_{it}^1 = \log A_j + u_{it}^1 \qquad \text{and} \qquad x_{it}^2 = x_{jt}^1 + \xi_t + u_{it}^2,$$

where $j = m(i,t)$ is i's trading partner, and where $u_{it}^1 \sim \mathcal{N}(0, \sigma_{u1}^2)$ and $u_{it}^2 \sim \mathcal{N}(0, \sigma_{u2}^2)$ are idiosyncratic noises, with $\sigma_{u1}, \sigma_{u2} > 0$. Note that x_{it}^1 represents a private signal that i receives about j's TFP, while x_{it}^2 represents a private signal that i receives about j's information about its own TFP. The shock ξ_t then introduces an aggregate noise component in the second type of signals.

Notice here that the posterior belief of island i about the TFP of its trading partner is pinned down by the signal x_{it}^1 alone, which is itself invariant to the sentiment shock ξ_t. It follows that ξ_t does not affect beliefs of either aggregate or idiosyncratic fundamentals. Yet, as we verify below, ξ_t triggers aggregate fluctuations, in both aggregate output and forecasts of economic activity.

Proposition 3. *Consider the equilibrium of the economy described above.*

(i) $\log Y_t$ *and* $\log \mathcal{B}_t$ *are increasing linear functions of* ξ_t.

(ii) There exist scalars $\phi_0, \psi_0 \in \mathbb{R}$ *and* $\phi_a, \phi_1, \phi_2, \psi_a, \psi_1, \psi_2 \in \mathbb{R}_+$ *such that, for all* (i, t, ω_{it}),

$$\begin{aligned} \log y_{it} &= \phi_0 + \phi_a \log A_i + \phi_1 x_{it}^1 + \phi_2 x_{it}^2 \\ \mathbb{E}_{it} \log p_{it} &= \psi_0 - \psi_a \log A_i + \psi_1 x_{it}^1 + \psi_2 x_{it}^2 \end{aligned}$$

Part (i) characterizes the aggregate behavior of the economy: variation in ξ_t triggers positive co-movement in aggregate economic activity, as measured by Y_t, and in the average sentiment, as measured by $\mathcal{B}_t$. Part (ii) reveals the micro-level behavior that rests beneath these fluctuations: an

[12] In particular, if we let $\mathcal{F}$ be time-varying and relax the assumption that the latter is common knowledge, Proposition 2 continues to hold as soon as we replace the constant $\bar{a}$ with the corresponding time-varying $\bar{a}_t$.

[13] This explains why the equilibrium aggregates we characterize in Proposition 3 below are functions of only the current shock ξ_t as opposed to its entire history.

increase in either x^1_{it} or x^2_{it} leads island i to expect an improvement in its terms of trade, which explains why y_{it} increases with either of these signals, and thereby also with ξ_t.

To build intuition for this result, suppose for a moment that the output of island i depended only on local TFP. It would then be optimal for island j to condition its own output, not only on its own TFP, but also on x^1_{jt}: a higher x^1_{jt} signals that i's output is likely to be higher and hence that the demand for j's product is also likely to be higher (equivalently, that its terms of trade will improve). But then it would become optimal for island i to raise its own production when it observes either a higher x^1_{it} or a higher x^2_{it}, for either observation would now signal that island j is likely to produce more and hence that the demand of i's product is likely to be higher. This explains why an island's expected terms of trade and its output increase with either signal.

The above intuition is based on recursive reasoning—equivalently, on iterating the contraction mapping behind Propositon 1. While illuminating, this is not strictly needed. A simpler intuition emerges once one focuses directly on the fixed point. In equilibrium, *either* of the two signals serves as a signal of the likely level of demand. The fact that one signal is intrinsic while the other is extrinsic is irrelevant to the decisions of firms and households. Rather, all that matters for them is simply that either signal contains "news" about the level of economic activity in other islands, and hence about the likely level of demand for the local product. Whenever a positive innovation occurs in ξ_t, *all* islands receive "good news" of the extrinsic type. For firms, this means an increase in expected marginal returns, which motivates them to expand their production and raise their demand for labor and land. In equilibrium, this stimulates employment and output, while also raising the wage and the rental rate (and thereby land prices). All in all, the economy ends up experiencing a boom that may appear self-fulfilling in the eyes of an outside observer.[14]

The insight that emerges is more general than the example. Since the aggregate fundamentals are fixed, any boom or recession in our model *necessarily* reflects a random shift in one island's optimism or pessimism about another island's employment and production. The particular information structure we have assumed in the preceding example in order to engineer such random shifts in beliefs of economic activity is not meant to be realistic; it is only a simple illustration of the broader insight. We elaborate on this point in the sequel.

5 Discussion and broader insights

In this section we elaborate on the generality, applicability, and empirical content of our insights. First, we explain how our extrinsic fluctuations can be understood as the symptom of aggregate variation in *higher-order* beliefs of *exogenous* fundamentals. Next, we discuss the theoretical and empirical reasons that motivate us to side-step this game-theoretic representation and, instead, interpret our sentiment shocks as shocks to *first-order* beliefs of *endogenous* economic outcomes. Finally, we discuss the modeling role that trading frictions play in our environment.

[14] Note here how our theory formalizes news of economic activity in terms of extrinsic forces rather than news about fundamentals. We revisit this point, which connects to the recent work on "news shocks", in Section 7.

Higher-order beliefs vs sentiments. Consider the following generalization of the example studied in the previous section. Fix a finite $H > 1$ and suppose that the signal x_{it} that is received by island i in period t is given by $x_{it} = (x_{it}^1, x_{it}^2, ..., x_{it}^H)$, where

$$x_{it}^1 = \log A_j + \varepsilon_{it}^1 \qquad \text{and} \qquad x_{it}^h = x_{jt}^{h-1} + \varepsilon_{it}^h \;\; \forall h \in \{2, ..., H\}$$

That is, islands receive signals of the signals... of the signals of others. Suppose further that the error terms ε_{it}^h have both idiosyncratic and aggregate components: $\varepsilon_{it}^h = \xi_t^h + u_{it}^h$, where ξ_t^h is the aggregate component and u_{it}^h is the idiosyncratic one. These components are uncorrelated across h and t, as well as with one another, and are drawn from Normal distributions with zero means and variances $(\sigma_\xi^h)^2$ and $(\sigma_u^h)^2$, respectively. Finally, to contrast our sentiment shocks to conventional technology shocks, let us also introduce aggregate TFP shocks: $\log A_{it} = a_i + \bar{a}_t$, where a_i is a Normally distributed island-specific fixed effect and $\bar{a}_t$ is a commonly-known aggregate shock.

Consider now the implied hierarchy of beliefs within a particular match (i.e., i's belief of A_j, i's belief of j's belief of A_i, and so on). Variation in ξ_t^h causes variation in beliefs of order h and above, but not in beliefs of order lower than h. Each of these shocks therefore has a *distinct* effect on the hierarchy of beliefs about the fundamentals. Nevertheless, as shown in the next proposition, these shocks are completely indistinguishable when it comes to equilibrium behavior: the entire aggregate variation in macroeconomic outcomes and in expectations of economic activity is spanned by a single composite of all these shocks, which we denote below by $\bar{\xi}_t$. It is then only this composite shock that we wish to think of as the proper measure of what a "sentiment shock" is.

Proposition 4. *Consider the equilibrium of the economy described above. There exist scalars $\Phi > 0$ and $\Psi > 0$ such that*

$$\log Y_t = \Phi\, \bar{a}_t + \Psi\, \bar{\xi}_t \qquad \textit{and} \qquad \log \mathcal{B}_t = \Phi\, \bar{a}_t + \bar{\xi}_t \tag{10}$$

where $\bar{\xi}_t$ is a linear combination of $(\xi_t^1, ..., \xi_t^h)$.

To understand this result, recall from condition (8) that the equilibrium output of each island depends *only* on its *first-order* beliefs of the level of output in other islands—not on the details of the information structure upon which these beliefs are formed. It follows that the entire extrinsic variation in aggregate outcomes can be captured in a single random variable $\bar{\xi}_t$, which summarizes the combined impact of all the exogenous shocks $\{\xi_t^h\}_{h=0}^H$ on either aggregate output, Y_t, or the average expectation, $\mathcal{B}_t$. One can thus think of $\bar{\xi}_t$ as *the* sentiment shock.

This result clarifies two points. First, there are multiple ways to shock the information structure so as to obtain the extrinsic fluctuations we are interested in: *any* of the ξ_t^h shocks serves our goals. And second, what matters for the observables of the theory is only the equilibrium variation in the *first-order* beliefs of economic activity. The precise details of how the variation in this kind of expectations is engineered by certain signals, or by shocks to, say, tenth-order beliefs of one another's TFP, is neither of particular interest to us, nor of any empirical relevance.

Indeed, suppose that an "econometrician" views the available data on aggregate employment and output, perhaps along with surveys of economic forecasts, through the lens of our model. This data may well permit the econometrician to identify separately the composite extrinsic shock $\bar{\xi}_t$ from the technology shock $\bar{a}_t$. For example, the technology shock $\bar{a}_t$ can first be identified by data on the Solow residual, and the composite sentiment shock $\bar{\xi}_t$ can then be identified by the residual variation in observed output. Survey evidence on market expectations could the be used either as alternative source of identification, or for testing the model.[15] By contrast, the information structure and the hierarchy of beliefs about the underlying fundamentals are not uniquely identified: there are multiple specifications of these objects that give rise to *exactly* the same joint distribution for equilibrium expectations and equilibrium outcomes.[16]

The insight that emerges here is quite general: the empirical content of *any* rational-expectations model rests exclusively on the joint distribution of equilibrium expectations and equilibrium outcomes. Proposition 4 illustrates this insight in a particularly stark way: a single number, the composite shock $\bar{\xi}_t$, happens to summarize the entire aggregate extrinsic variation in equilibrium expectations. Clearly, this stark property hinges on a number of simplifying assumptions that are embedded in the micro-foundations of our model.[17] In richer models, one often needs a larger state space in order to keep track of the equilibrium dynamics. Yet, the broader insight survives: for applied purposes, the key issue is the equilibrium variation in first-order beliefs of economic outcomes, not the underlying belief hierarchies.

These points explain why we insist to interpret our "sentiment shocks" as extrinsic movements in (first-order) expectations of economic activity rather than as shocks to higher-order beliefs of fundamentals. To reinforce this interpretation, it is useful to study a variant model that departs from rationality in order to introduce *exogenous* shocks to this kind of expectations.[18] Thus suppose that an island's forecast of its trading partner's output obeys the following ad hoc law of motion:

$$\hat{\mathbb{E}}_{it} \log y_{jt} = \log y^*_{jt} + \zeta_t, \tag{11}$$

where $\log y^*_{jt}$ denote the equilibrium output levels that obtain in the perfect-information, rational-expectations benchmark[19] and ζ_t is an exogenous shock that perturbs the agents' expectations away from this benchmark. One may thus think of ζ_t as *irrational* shifts to "market psychology."

[15]The most natural empirical counterpart for the type of expectations that matter in the theory seems to be the forecasts that firms make about their own sales (or demand). In the absence of such data, one could also proxy $\mathcal{B}_t$ by the average forecast of *aggregate* output. Indeed, as shown in the proof of Proposition 4, these two distinct kinds of forecasts are positively correlated within our model, which justifies this point.

[16]Here, and throughout the paper, we use the term "equilibrium expectations" as synonymous to "first-order beliefs of endogenous economic outcomes". This is consistent with the spirit of the entire rational-expectations literature.

[17]First, there is no capital and matching is pairwise, guaranteeing that the only first-order beliefs that matter are those that each island forms about the contemporaneous output of a single other island. Second, preferences are homothetic and technologies are CRS, implying the corresponding reduced-form game features linear best responses. Third, the information structure is Gaussian, so that any belief (probability distribution) can be captured by its mean and its variance, and all variances are held constant, so that ultimately only means vary. Finally, all the exogenous shocks are i.i.d. over time, ruling out complicated learning dynamics as in, say, Townsend (1985).

[18]Expectations of economic outcomes are themselves endogenous in any rational-expectations context. If one wishes to introduce exogenous shifts to this kind of expectations, some departure from rationality is inevitable.

[19]For the purposes of the present exercise, y^*_{jt} is merely a particular function of the exogenous fundamentals (A_{it}, A_{jt}), not an endogenous element of the equilibrium.

Proposition 5. *Consider the equilibrium of the non-rational-expectations variant described above.*[20] *There exist scalars $\Phi > 0$ and $\Psi > 0$ such that*

$$\log Y_t = \Phi\,\bar{a}_t + \Psi\,\zeta_t \qquad \text{and} \qquad \log \mathcal{B}_t = \Phi\,\bar{a}_t + \zeta_t$$

Furthermore, the scalars $\Phi > 0$ and $\Psi > 0$ are the same as those in Proposition 4.

The present variant is therefore observationally equivalent to the example we studied in the beginning of this section. But whereas the present variant violates rational expectations, our earlier example did not. In this sense, our notion of sentiment shocks provides an exact rationalization of random, and seemingly irrational, shifts in expectations of economic activity.

Trading frictions and extrinsic volatility. By impeding the ability of economic agents to coordinate their beliefs and actions on those that would have obtained in the frictionless Arrow-Debreu benchmark, trading frictions permit us to accommodate the Keynesian notion that recessions are the product of some kind of "coordination failure". This is similar in spirit to Diamond (1982). But whereas Diamond formalized the aforementioned notion by tying trading frictions to thick-market externalities and multiple equilibria, we achieve the same objective merely by letting trading frictions impede communication along the unique equilibrium of the economy.

Trading frictions also permit us to sustain aggregate volatility in equilibrium outcomes without *any* aggregate shocks to fundamentals such as preferences and technologies. While the latter property is not strictly needed,[21] it helps sharpen our theoretical contribution relative to the pertinent literature; it also adds a degree of flexibility for quantitive purposes.

To elaborate on this last point, consider the class of games studied in Morris and Shin (2002) and Angeletos and Pavan (2007); the observations we make below extend more generally to the pertinent macroeconomics literature on informational frictions. In this class of games, the equilibrium can be represented as the fixed point to the following relation:

$$y_i = (1-\alpha)\mathbb{E}_i[A] + \alpha\mathbb{E}_i[Y] \tag{12}$$

where y_i is the action of agent i, Y is the corresponding aggregate, A is the aggregate payoff-relevant fundamental, and $\alpha \in (0,1)$ is a scalar parameterizing the degree of strategic complementarity.[22] The literature then proceeds by specifying a particular information structure—sometimes exogenous, sometimes endogenous—and characterizing the resulting equilibrium. But even if we do not spell out the details of the information structure, we can obtain a tight upper bound on the equilibrium level of aggregate volatility as follows. First, aggregate condition (12) across i and iterate over the

[20]Since beliefs are hereby treated as exogenous, the definition of the equilibrium must be adjusted accordingly: an equilibrium is now a collection of allocations that are optimal for the households and the firms taking as given the aforementioned, exogenously-specified, beliefs, along with the wages and prices that clear the various markets.

[21]In fact, if we add unobservable aggregate shocks to fundamentals, we can engineer additional extrinsic volatility from higher-order uncertainty about these shocks.

[22]The precise interpretation of these variables varies from application to application. For example, the relevant fundamental is an aggregate monetary shock in Woodford (2003), Mankiw and Reis (2002), and Mackowiak and Wiederholt (2009), whereas it is an aggregate TFP shock in Angeletos and La'O (2009) and Lorenzoni (2009).

expectation operator to obtain the the equilibrium aggregate action as a weighted average of the hierarchy of beliefs about the aggregate shock:

$$Y \;=\; (1-\alpha)\sum_{h=1}^{\infty}\alpha^{h-1}\overline{\mathbb{E}}^{h}[A] \tag{13}$$

where $\overline{\mathbb{E}}^{h}[.]$ denotes the h-th order average forecast. Next, recall that the variance of the forecast of any random variable is necessarily no larger than the variance of the variable itself. It follows that $Var(\overline{\mathbb{E}}^{h}[A]) \geq Var(\overline{\mathbb{E}}^{h+1}[A])$ for all h. Using this fact along with (13), we infer that

$$Var(Y) \leq Var(A).$$

That is, the aggregate volatility in economic activity is bounded from above by the aggregate volatility in fundamentals. As the latter vanishes, the former also vanishes.

These facts are true no matter whether the macroeconomic volatility is driven by actual changes in the fundamentals (say, TFP shocks) or by noise in signals about them ("noise shocks"). Furthermore, the upper bound is attained, namely $Var(Y) = Var(A)$, when the fundamentals are perfectly known. In practical terms, this means that introducing incomplete information about aggregate shocks to technology or other fundamentals is likely to be counter-productive on its own right if the ultimate goal is to explain the observed business cycle with smaller such shocks.

Our theory offers a simple resolution to this conundrum. Contrast condition (8) in our model with condition (12) above. The key formal difference is that an agent's best response hinges on his forecast of the idiosyncratic action of a random trading partner rather than his forecast of the aggregate action.[23] As a result, even if there are no aggregate shocks to fundamentals, we can always sustain an arbitrarily high level of aggregate volatility by (i) assuming sufficiently large idiosyncratic risk and (ii) engineering enough correlation in the agents' beliefs of their idiosyncratic economic outcomes. The precise quantitative value of this added flexibility remains to be explored.

6 Communication, Contagion, and Boom-and-Bust Cycles

Our fluctuations hinge on the existence of correlated movements in agents' beliefs of economic activity. In our preceding analysis, this correlation was hard-wired in exogenous signals. More naturally, such correlation may emerge as the by-product of how agents communicate through, say, the markets, social networks, or the media—communication *means* correlation.

We illustrate this idea in this section by considering an example in which an exogenous sentiment shock hits only a few agents in the beginning, but then spreads endogenously in the rest of the economy as these agents trade and communicate with other agents. The resulting belief dynamics resemble the spread of fads and rumors and give rise to phenomena akin to boom-and-bust cycles.[24]

[23] Another difference between (8) and (12) is that the idiosyncratic fundamental A_i, or f_i, shows up in (8) instead of the aggregate fundamental A in (12). This difference, however, is not crucial on its own right.

[24] To be clear, the particular aspect of boom-and-bust phenomena that we seek to accommodate in this section is the underlying waves in market expectations, not the interplay between financial markets and the real economy. Adding such an interplay—as, e.g., in La'O (2010)—could enrich the propagation of the belief waves we document.

Setup. Consider the following variant of our model. At $t = 0$, the islands are split into two equally-sized groups. TFP is the same within a group but differs across groups. Think of these groups as "North" and "South", let A_N and A_S be the respective TFP levels, and assume that these are i.i.d. draws from a log-Normal distribution.

Each of these two groups is then split into two subgroups. Islands in the first subgroup observe nothing more than their own productivities; we refer to them as "uninformed". Islands in the second subgroup, which we refer to as "partially informed", get to see two additional signals. Similarly as in Section 4, these signals are given by $x_N^1 = \log A_S + \varepsilon_N$ and $x_N^2 = x_S^1 + \xi$ for the North, and $x_S^1 = \log A_N + \varepsilon_S$ and $x_S^2 = x_N^1 + \xi$ for the South, where $\varepsilon_N, \varepsilon_S$ and ξ are all Normally distributed, independent of one another, and independent of the TFP draws. The initial fraction of partially informed islands is given by $\chi \in (0, 1/2)$; the rest are uninformed.

The exogenous aggregate state is summarized in $\tilde{s} = (A_N, A_S, \varepsilon_N, \varepsilon_S, \xi)$. Once Nature draws $\tilde{s}$ at $t = 0$, no other aggregate shock ever hits the economy, and no further exogenous information ever arrives—islands learn only in an endogenous manner, as they meet and "talk" to one another. The entire dynamics we document below are thus the sole product of this kind of communication.

To obtain a closed-form solution of the equilibrium, the random matching is assumed to take the following form. First, an uninformed island can meet either a similarly uninformed island from its own productivity group, in which case it learns nothing, or a partially informed one from it own productivity group, in which case it learns the latter's information and hence turns into a partially informed island next period. Second, a partially informed island can meet either an uninformed one from its own productivity group, in which case it learns nothing itself, or a partially informed one from the *other* productivity group, in which case they both learn the entire state $\tilde{s}$ and turn into a third category, which we call "fully informed". Third, a fully informed island can only meet with a fully informed from its own productivity group. And finally, each island knows beforehand (in stage 1) whether it is matched with an island of the same or different information category.

This structure defines a three-step "information ladder". The uninformed islands are at the bottom, the partially informed in the middle, and the fully informed at the top. In each period, an island ascends at most one step in this ladder, depending on the information of its match. Eventually, all islands reach the top, but this takes time. The dynamics we document below are a manifestation of how the population ascends this ladder.

Results. It is easy to check that the only islands whose employment and production choices are sensitive to the initial sentiment shock, ξ, are partially informed islands that expect to be matched with other partially informed islands. These islands behave in essentially the same way as in the example of Section 4. But, whereas in this earlier example all the islands behaved in that fashion, here only a fraction does. Furthermore, this fraction evolves over time, due to the communication that takes place as islands meet and trade.

To fix ideas, we henceforth focus on positive realizations for ξ, which translates to a wave of optimism. We accordingly refer to the partially-informed islands as "exuberant" and let μ_t be the fraction of such islands in the population.

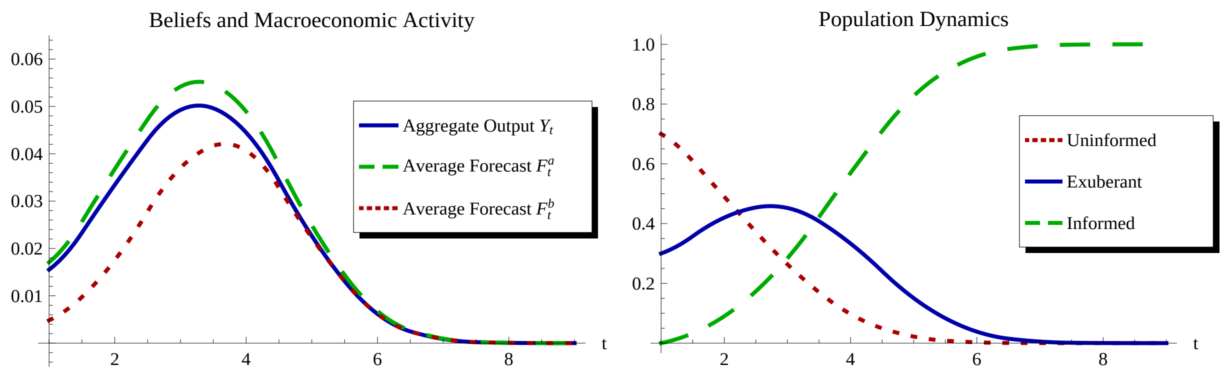

Figure 1: Contagion and boom-and-bust cycle. The left panel illustrates the response of aggregate output to a positive sentiment shock (solid line), along with that of the average of the beliefs that each island holds about either the output of its trading partner (dashed line) or aggregate output (doted line). The right panel illustrates the underlying population dynamics.

Proposition 6. *(i) The economy experiences a "fad": the fraction of "exuberant" islands, μ_t, initially increases, but later on falls and eventually converges to zero.*

(ii) There exists a scalar $\Phi > 0$ such that the dynamic response of aggregate output to the initial sentiment shock is given by

$$\frac{\partial \log Y_t}{\partial \xi} = \Phi \mu_t, \ \forall t.$$

These results are illustrated in Figure 1. The left panel documents the dynamic response of aggregate output, and of the islands' forecasts, to the sentiment shock. The right panel reveals the underlying population dynamics (i.e., the evolution of the distribution of islands along the aforementioned information ladder). It is evident that the dynamics of actual and expected output track the dynamics of the fraction of "exuberant" islands, which is first increasing and then decreasing. A similar result holds for wages and employment, as well as for asset (land) prices, which are, in effect, forecasts of future economic activity. The economy thus experiences a "wave of optimism" that builds up force for a while, only to fade away later on—there is a boom followed by a bust.

During the boom phase, more and more islands receive "good news" about the level of economic activity in other islands, and hence about their terms of trade. For those islands that were born exuberant at $t = 0$, this news arrives exogenously, from Nature. For those islands that become exuberant in any subsequent period, these news arrive endogenously, as these islands meet islands that were already exuberant. Finally, as time passes, more and more islands become fully informed. The bust phase is thus associated with a "correction" in previously exuberant beliefs. Communication causes the fraction of exuberant islands first to increase and then to fall.

The contagion effects behind these population dynamics are reminiscent of those discussed, inter alia, in Shiller (2005) and Akerlof and Shiller (2009): "irrational exuberance" is said to spread in the economy as one agent hears "stories" from other agents. In fact, our dynamics are very similar to those found in Burnside, Eichenbaum, and Rebelo (2011), in a study of the recent boom-and-bust cycle in housing prices. But whereas these authors model the contagion between different agents as the product of behavioral (irrational) heuristics, here we show that it may be merely the symptom of the (imperfect) communication that takes place via the market mechanism and other social interactions. Exuberance then spreads *because* of rationality.

Putting aside any interpretation, three additional remarks are worth making regarding the mechanics of our theory, as illustrated in the above example. First, although our theory—like any other theory—requires an exogenous initial trigger for fluctuations in endogenous economic outcomes to kick off, this trigger may rest in a small fraction of the population and nevertheless give rise to a pervasive wave of optimism or pessimism in the entire economy. Second, as long as communication is imperfect, more communication may actually amplify our fluctuations: markets, macroeconomic statistics, the media, and the blogosphere may serve as channels of contagion. Finally, to the extent that communication gets finer and finer with time, equilibrium beliefs must eventually converge, which guarantees that the impact of any given extrinsic shock eventually vanishes. The fluctuations we formalize in this paper therefore embed, not only a natural propagation mechanism, but also a natural mean-reverting mechanism: booms must be followed by busts, recessions by recoveries.

7 A Quantitative Exploration

Although our contribution is primarily methodological, we also wish to illustrate its quantitative potential. Towards this goal, we consider a variant of the RBC model that replaces the conventional notion of technology shocks with our notion of sentiment shocks.

Setup. To accommodate capital accumulation, we reinterpret the specialized goods that are traded across the islands as intermediate inputs into the production of a local final good, which in turn is used either for consumption or for investment. Trade takes place in terms of these inputs.

The local resource constraint of island i (equivalently, the market-clearing condition for the final good) is given by $c_{it} + i_{it} = y_{it}$, where c_{it} is consumption, i_{it} is gross investment, and y_{it} is the output of the local final-good sector. Capital accumulates according to the following law of motion:

$$k_{i,t+1} = \left(1 - \Delta\left(e_{it}\right)\right) k_{it} + i_{it},$$

where k_{it} denotes the local capital stock, $\Delta\left(e_{it}\right)$ is the rate of capital depreciation, and e_{it} is the rate of capital utilization. As in King and Rebelo (2000), capital depreciation is an increasing convex function of capital utilization: $\Delta(e) = \frac{\delta}{\mu} e^{\mu}$, with $\delta > 0$ and $\mu > 1$.[25]

The intermediate-good sector makes its input and production choices prior to observing the relevant terms of trade; this introduces essentially the same type of terms-of-trade uncertainty as in our baseline model. The production of the local intermediate input is given by

$$q_{it} = A_i (e_{it} k_{it})^{1-\theta} (n_{it})^{\theta},$$

[25] By introducing variable capital utilization, we are able to generate procyclical labor productivity despite the absence of aggregate technology shocks. If we remove this feature, labor productivity becomes counter-cyclical in response to sentiment shocks, but the rest of the cyclical properties of the model are not seriously affected.

where n_{it} denotes local employment. Profits are given by $\pi^q_{it} = p_{it}q_{it} - w_{it}n_{it} - (r_{it} + \Delta(e_{it}))k_{it}$, where p_{it} is the price of the local intermediate good, w_{it} is the local wage, and r_{it} is the local rental rate of capital net of depreciation costs.

The production of the final good is given by

$$y_{it} = \tfrac{1}{\zeta} \left(h_{it}\right)^{1-\eta} \left(h^*_{it}\right)^{\eta}$$

where h_{it} and h^*_{it} are the "home" and "foreign" intermediate inputs and $\zeta \equiv (1-\eta)^{(1-\eta)}\eta^{\eta}$ is a constant. Profits are given by $\pi^y_{it} = y_{it} - p_{it}h_{it} - p^*_{it}h^*_{it}$; in equilibrium, profits are zero because the technology is constant returns to scale and the final producers adjust production after observing all prices. Market clearing for intermediate inputs imposes $q_{it} = h_{it} + h^*_{jt}$, where j stands for i's trading partner during period t.

Finally, consider the representative household of island i. Its preferences are standard:

$$\mathcal{U}_i = \sum_{t=0}^{\infty} \beta^t \left[U(c_{it}) - V(n_{it})\right].$$

where $U(c) = \frac{1}{1-\gamma}c^{1-\gamma}$, $V(n) = \frac{1}{\epsilon}n^{\epsilon}$, $\gamma > 0$, and $\epsilon > 1$. Its budget constraint is given by

$$c_{it} + i_{it} = \pi^y_{it} + \pi^q_{it} + w_{it}n_{it} + r_{it}k_{it}.$$

Characterization. As in our baseline model, a partial characterization of the equilibrium can be obtain without spelling out the details of the information structure.

Proposition 7. *Any equilibrium allocation solves the following system:*

$$\begin{aligned}
V'(n_{it}) &= \theta\zeta\mathbb{E}_{it}\left[U'(c_{it})\frac{y_{it}}{n_{it}}\middle|\,\omega_{it}\right] \\
\Delta'\left(e_{it}\right)e_{it} &= \left(1-\theta\right)\zeta\mathbb{E}_{it}\left[U'(c_{it})\frac{y_{it}}{k_{it}}\middle|\,\omega_{it}\right] \\
U'(c_{it}) &= \beta\mathbb{E}_{it}\left[U'(c_{i,t+1})\left(1+\left(1-\theta\right)\zeta\tfrac{\mu}{1+\mu}\frac{y_{it+1}}{k_{it+1}}\right)\middle|\,z_{it}\right] \\
c_{it} + k_{it+1} &= y_{it} + (1-\Delta\left(e_{it}\right))k_{it} \\
y_{it} &= q_{it}^{1-\eta}q_{jt}^{\eta} \\
q_{it} &= A_{it}(e_{it}k_{it})^{1-\theta}(n_{it})^{\theta}
\end{aligned}$$

The top four conditions should be familiar: they are the optimality conditions for labor and capital utilization, the Euler condition, and the resource constraint. The remaining two conditions specify the production levels of the various goods. Compared to the standard RBC model, the only essential novelties are (i) that the income of each island depends on the production choices of another island, through the relevant terms-of-trade effect; and (ii) that expectations are heterogeneous.

The key mechanism thus remains the same as in our baseline model: booms and recessions are driven by extrinsic shocks to beliefs about "demand" (about the output of other islands). Interestingly, however, these fluctuations now manifest, not only in employment, but also in investment and capital utilization. What is more, as all these decisions are infinitely forward-looking, economic activity in one period may respond to extrinsic belief shifts about economic activity far in the future—it is *as if* the islands are playing a dynamic game in which an island's optimal employment, consumption and investment choices during one period depend on the expected output of its likely trading partner, not only in the current period, but also in all future periods.

Priors and sentiments. While the aforementioned characterization of the equilibrium is conceptually straightforward, a numerical solution remains challenging in the case of persistent sentiment shocks, because of an infinite-regress problem similar to that in Townsend (1983). To bypass this challenge, we introduce a heterogeneous-prior variant of the information structure that permits us to embed persistent belief shocks while ruling out complicated learning dynamics.[26]

Under this variant, each island receives a single signal about its trading partner, given by

$$x_{it} = \log A_{jt} + \varepsilon_{it}, \tag{14}$$

where ε_{it} is an error term. The islands continue to share a common prior about the underlying fundamentals, but have heterogeneous priors about these error terms. In particular, each island believes (i) that its own error is unbiased, drawn from a Normal distribution with zero mean and variance $\sigma_\varepsilon^2 > 0$, and (ii) that the errors of all other islands are biased, drawn from a Normal distribution with the same variance but a mean equal to ξ_t, where ξ_t follows a Markov process.

This variable plays the same modeling role as in the preceding analysis: a positive innovation in ξ_t raises the islands' higher-order beliefs of one another's fundamentals without affecting the corresponding first-order beliefs. The key change with this variant is therefore the computational gain: if we assume that ξ_t is commonly known to all agents, then (see Appendix B) the log-linearized dynamics of the economy can be summarized in a linear policy rule $\Gamma : \mathbb{R}_+^2 \to \mathbb{R}_+$ such that

$$\tilde{K}_{t+1} = \Gamma(\xi_t, \tilde{K}_t).$$

where the tilde indicates log-deviation from steady state. This is akin to the policy rule of the standard RBC model, except that the conventional TFP shock has been replaced by our sentiment shock. The numerical implementation thus becomes straightforward.

Numerical results. We work at quarterly frequency and set $\beta = .99$, $\gamma = 2$, $\theta = .65$, $\epsilon = 2$, and $\mu = 2$; these values are consistent with King and Rebelo (2000). Next, we assume that ξ_t follows an AR(1) process: $\xi_t = \rho\xi_{t-1} + \nu_t$, where $\rho \in (0, 1)$ and ν_t is i.i.d. Normal with mean zero and variance σ_ξ^2. We set $\rho = .98$, which builds strong persistence in our fluctuations. The remaining parameters (η, σ_A, σ_ε and σ_ξ) then matter for aggregate dynamics only through a single composite coefficient, which itself scales up and down all aggregate outcomes. Exploiting this property, we fix

[26] This variant is in tension with the strong version of rational expectations, which requires that the prior of each agent coincides with the objective truth. Before opting for this variant, we thus experimented with few alternatives that maintained the common-prior assumption. One was to guess a low-dimensional state space as in Woodford (2003). Another was to let the shock become common knowledge with a lag of $T \geq 2$ periods. Unfortunately, neither of these attempts worked out. A remaining possibility which we did not explore is to restrict the state space and search for a "myopic equilibrium" as in Krusell and Smith (1998).

	The Model		U.S. Data	
	std. dev	corr(X,Y)	std. dev	corr(X,Y)
Y	1.67	1.00	1.67	1.00
N	1.41	1.00	1.43	0.87
C	1.21	0.98	1.27	0.80
I	4.14	0.96	5.48	0.82
Y/N	0.26	0.99	0.82	0.51
LW	4.95	-1.00	4.98	-0.82

Table I. This table documents the business-cycle statistics of our model along with those of the US economy. All quantities are in quarterly frequency and HP-filtered. See Appendix B for details.

$\eta = 1$ and $\sigma_a = \sigma_\varepsilon = \sigma_\xi = \sigma$, and then set $\sigma = 0.038$, which induces the variance of the HP-filtered aggregate output in our model to match the corresponding moment in the US data.[27] We then simulate the dynamics of the economy and report the model's HP-filtered business-cycle statistics in Table I, along with the corresponding statistics for the US economy.

Given that the volatility of output is matched by design, the question of interest is whether our model also matches the relative volatility and the co-movement of all the other macroeconomic variables. As evident in Table I, our model is quite successful in this respect. Sentiment shocks cause employment, consumption, investment, and labor productivity to co-move with output, as in actual business cycles. Furthermore, the quantitative effects are in the ballpark of the actual data.

Relative to the standard RBC model, we do worse in that we generate little pro-cyclicality in labor productivity. This is simply because we have ruled out technology shocks.[28] But we also do better in that we match the observed counter-cyclicality of the labor wedge, which is an important feature of the data (Chari, Kehoe and McGrattan, 2007; Shimer, 2009).[29] To understand this property, consider the stripped-down version of our model where capital and utilization are fixed. As a negative sentiment shock causes firms to turn pessimistic about their profitability, labor demand and employment fall. As this happens, the average labor productivity actually goes up. Standard business-cycle accounting will thus register the resulting recession as an increase in the implicit tax on labor. By the same token, the recession may manifest as an increase in the measured output gap—and can thus be interpreted as the symptom of "insufficient aggregate demand".

It is also worth contrasting the cyclical properties of our theory with those of the literature on "news shocks". Bound by conventional DSGE practice, this literature formalizes news of economic activity as news of future technology. In so doing, it faces a significant difficulty in generating the

[27]Our calibration strategy is consistent with standard DSGE practice, where the various shocks are estimated so that the model matches the data.

[28]If capital utilization were fixed, labor productivity would have been countercyclical, due to diminishing returns; employment, consumption and investment would, however, remain procyclical. Also, note that the procyclicality of labor productivity in the data has actually vanished during the last two decades (Gali and van Rens, 2010). If we focus on this period, the quantitative performance of our theory improves.

[29]The wedge is defined, in logs, by $LW_t \equiv \log\left[\frac{U'(C_t)}{V'(N_t)}\theta\frac{Y_t}{N_t}\right]$, so that a positive wedge maps to a tax on labor.

observed joint procyclicality in employment, consumption, and investment (Beaudry and Portier, 2006). The usual fixes involve exotic preferences (Jaimovich and Rebelo, 2009), a suboptimal monetary policy that responds to noise by stimulating aggregate demand (Lorenzoni, 2010), or some combination of the two (Christiano, Ilut, Motto, Rostagno, 2008). By contrast, our theory permits one to formalize news of economic activity as extrinsic forces that induce optimistic beliefs about both current and future firm profitability and consumer income. This helps stimulate employment, consumption and investment in a similar fashion as a conventional technology shock, which explains why our theory has no serious difficulty in capturing the observed co-movement in these variables.

8 Concluding remarks

Are business cycles and fluctuations in asset markets driven by changes in preferences and technologies? Or are they driven by "animal spirits", "market psychology", and self-reinforcing waves of optimism and pessimism?

This question is not just an empirical matter. To address it, one must first propose a precise theory that formalizes the popular but vague notions of "animal spirits" and the like; to paraphrase Lucas (2001), one needs equations that explain what these words mean.

This paper makes a contribution in precisely this direction: we develop a novel formalization of extrinsic movements in market expectations, one that requires neither a departure from rationality nor the introduction of multiple equilibria.

To achieve this, we relax the conventional assumption that all agents share the same beliefs about the state of the economy. We then show that, once this is true, economic outcomes and market expectations may co-move in response to a certain type of extrinsic shocks which we call "sentiments". These shocks are akin to sunspots, but operate in unique-equilibrium models. They rationalize random, and seemingly inexplicable, shifts in the optimism or pessimism that economic agents may hold about one another's choices and thereby about future market conditions.

To outside observers, the resulting fluctuations might look as "self-fulfilling", or as the product of mysterious "demand shocks" that are disconnected from preferences and technologies. In this respect, they have a genuinely Keynesian flavor. They are nevertheless consistent with the neoclassical paradigm, resting merely on the heterogeneity of people's expectations and the consequent misalignment of their choices.

The combination of these points underscores what, in our view, is the relative strength of our theory. Not only is our theory capable of matching key business-cycle facts, as illustrated in the previous section; it also helps accommodate within the dominant macroeconomic paradigm a set of popular notions about the "real" workings of the economy that have so far been considered inconsistent with this paradigm.

Introducing our notion of shocks in richer DSGE models, and estimating their contribution to observed business cycles, is a natural direction for future research. Translating our insights in the context of asset markets and studying their policy implications are two other possible directions.

Appendix A: proofs

Proof of Proposition 1. From the the optimality condition for labor (5), we get

$$n_{it} = (\mathbb{E}_{it}[p_{jt}]\theta y_{it})^{\frac{1}{\epsilon}}$$

Substituting the above into the production function yields

$$y_{it} = A_i K^{1-\theta} \left(\mathbb{E}_{it}[p_{jt}]\theta y_{it}\right)^{\frac{\theta}{\epsilon}}$$

Finally, solving the above for y_{it}, and letting $\vartheta \equiv \theta/\epsilon$, we obtain

$$y_{it} = \left(\theta^{\vartheta} A_i K^{1-\theta}\right)^{\frac{1}{1-\vartheta}} \left(\mathbb{E}_{it}\left[p_{it}\right]\right)^{\frac{\vartheta}{1-\vartheta}}$$

which gives the first condition in the proposition. The second condition follows directly from condition (4).

Proof of Theorem 1. Substituting (7) into (6) and rearranging, we get

$$y_t(\omega)^{1+\eta\frac{\vartheta}{1-\vartheta}} = \left(\theta^{\vartheta} A_t(\omega) K^{1-\theta}\right)^{\frac{1}{1-\vartheta}} \left(\int_{\Omega_t} y_t(\omega')^{\eta} \mathcal{P}_t(\omega'|\omega) d\omega'\right)^{\frac{\vartheta}{1-\vartheta}}$$

Taking logs, we reach the following condition

$$\log y_t(\omega) = \frac{\frac{1}{1-\vartheta}}{1+\eta\frac{\vartheta}{1-\vartheta}} \log\left(\theta^{\vartheta} A_t(\omega) K^{1-\theta}\right) + \frac{\frac{\vartheta}{1-\vartheta}}{1+\eta\frac{\vartheta}{1-\vartheta}} \log\left(\int_{\Omega_t} y_t(\omega')^{\eta} \mathcal{P}_t(\omega'|\omega) d\omega'\right)$$

This reduces to condition (8) in the main text once we let $\alpha \equiv \frac{\eta}{\eta+(1-\vartheta)/\vartheta} \in (0,1)$ and $H(x) \equiv \eta \exp(x)$. It also means that we can recast the equilibrium allocations in period t as the solution to the above fixed point problem.

In particular, for each t, let $\mathcal{Y}_t$ be the set of real, bounded, and continuous functions with domain Ω_t, and endow this set with the sup-norm to obtain a complete metric space. Next, define the operator $\mathcal{T}_t : \mathcal{Y}_t \to \mathcal{Y}_t$ as follows: for any $f \in \mathcal{Y}_t$ and any $\omega \in \Omega_t$,

$$\mathcal{T}_t f(\omega) = (1-\alpha)\left\{\kappa + \tfrac{1}{1-\vartheta}\log A_t(\omega)\right\} + \alpha\left\{H^{-1}\left(\int_{\Omega_t} H\left(f(\omega')\right) \mathcal{P}_t(\omega'|\omega) d\omega'\right)\right\} \qquad (15)$$

where $\kappa \equiv \log(\theta^{\vartheta} K^{1-\theta})$ is a constant, $A_t(\omega)$ identifies the productivity of an island of type $\omega \in \Omega_t$, and $\mathcal{P}_t(\omega'|\omega)$ is the probability density with which this island meets an island of type $\omega' \in \Omega_t$.[30] Now, take any equilibrium and let $y_t \in \mathcal{Y}_t$ be the equilibrium output function in period t, for any t. Then, and only then, $\log y_t$ is a fixed point of $\mathcal{T}_t$.

Existence and uniqueness of the equilibrium then follows from the fact that, for all t, the operator $\mathcal{T}_t$ is a contraction with modulus equal to $\alpha \in (0,1)$. We verify this fact below by showing that $\mathcal{T}_t$ satisfies Blackwell's sufficiency conditions.

[30] The functions A_t and $\mathcal{P}_t$ are pinned down by the primitives of the economy: A_t is simply the function that, for any $\omega \in \Omega_t$, returns the first element of ω, while $\mathcal{P}_t$ follows from the exogenous stochastic structure of the economy.

(i) Monotonicity. Suppose $f, g \in \mathcal{Y}_t$ and $f(\omega) \geq g(\omega)$ for all $\omega \in \Omega_t$. First, note that

$$\mathcal{T}_t f(\omega) - \mathcal{T}_t g(\omega) = \alpha \left\{ H^{-1} \left(\int_{\Omega_t} H\left(f(\omega')\right) \mathcal{P}_t(\omega'|\omega) d\omega' \right) - H^{-1} \left(\int_{\Omega_t} H\left(g(\omega')\right) \mathcal{P}_t(\omega'|\omega) \right) d\omega' \right\}$$

Note that $\alpha > 0$ and that $H^{-1}(x) = \log(x/\eta)$, which is a monotonically increasing function. We infer that $\mathcal{T}_t f(\omega) - \mathcal{T}_t g(\omega) \geq 0$ if and only if

$$\int_{\Omega_t} \eta \exp\left(f(\omega')\right) \mathcal{P}_t(\omega'|\omega)\omega' \geq \int_{\Omega_t} \eta \exp\left(g(\omega')\right) \mathcal{P}_t(\omega'|\omega)\omega'. \tag{16}$$

Now, note that $f(\omega) \geq g(\omega)$ for all $\omega \in \Omega$ implies that $\eta \exp(f(\omega')) \geq \eta \exp(g(\omega'))$ for all $\omega \in \Omega_t$. This immediately implies that condition (16) is always satisfied. Therefore, $f \geq g$ implies $\mathcal{T}_t f \geq \mathcal{T}_t g$, which proves that $\mathcal{T}_t$ is monotonic.

(ii) Discounting. Let $a \geq 0$ be a constant. Then, using the fact that H is an exponential function, we have:

$$\begin{aligned}
\mathcal{T}_t\left[f(\omega) + a\right] &= (1-\alpha)\left\{\tfrac{1}{1-\vartheta}\log A_t(\omega)\right\} + \alpha \left\{ H^{-1} \left(\int_{\Omega_t} H\left(f(\omega') + a\right) \mathcal{P}_t(\omega'|\omega) d\omega' \right) \right\} \\
&= (1-\alpha)\left\{\tfrac{1}{1-\vartheta}\log A_t(\omega)\right\} + \alpha \left\{ H^{-1} \left(\int_{\Omega_t} H\left(f(\omega')\right) \mathcal{P}_t(\omega'|\omega) d\omega' \right) \right\} + \alpha a
\end{aligned}$$

Therefore, $\mathcal{T}_t\left[f(\omega) + a\right] = \mathcal{T}_t f(\omega) + \alpha a$, where $\alpha \in (0,1)$, which proves that $\mathcal{T}_t$ satisfies discounting.

As $\mathcal{T}_t$ satisfies both monotonicity and discounting, Blackwell's theorem applies, guaranteeing that the operator $\mathcal{T}_t$ is a contraction and completing the proof.

Proof of Theorem 2. That perfect communication rules out extrinsic fluctuations is proven in Proposition 2. The converse follows either from the simple example in Proposition 3 or from the generalized example in Proposition 4.

Proof of Proposition 2. As mentioned in the main text, perfect communication guarantees that the two islands within any given match know each other's output levels. Using this fact in condition (8) for island i and in the corresponding condition for island j, we obtain the total output of the two islands as

$$\log y_{it} + \log y_{jt} = f_i + f_j,$$

where, recall, $f_i \equiv \frac{1}{1-\vartheta}\log\left(\theta^\vartheta A_i K^{1-\theta}\right)$ identifies the local fundamentals of island i (and similarly for f_j). The result then follows from aggregating the above finding across all matches.

Proof of Proposition 3. In the proposed equilibrium, the period-t output of island j is log-normally distributed conditional on the information of island i, for any i, j, and t. Furthermore, the conditional variance $Var(\log y_{jt}|\omega_{it})$ is invariant to ω_{it}:

$$Var(\log y_{jt}|\omega_{it}) = \sigma_y^2 \equiv \phi_a^2\sigma_a^2 + \phi_1^2\sigma_{u1}^2 + \phi_2^2(\sigma_{u2}^2 + \sigma_\xi^2)$$

It follows that $\mathbf{E}_{it} \log y_{jt} = \mathbb{E}_{it} \log y_{jt} + \frac{1}{2}\eta^2\sigma_y^2$. The fixed-point condition (8) thus reduces to

$$\log y_{it} = const + (1-\alpha)\frac{1}{1-\vartheta}\, a_i + \alpha\, \mathbb{E}_{it}\left[\log y_{jt}\right] \tag{17}$$

where $a_i \equiv \log A_i$ and where $const$ is a scalar that is invariant with ω_{it} and that we henceforth ignore without any loss of generality.

We guess and verify a log-linear equilibrium under the log-normal specification for the shock and information structure. Suppose the equilibrium production strategy of the island of type ω_{jt} takes a log-linear form given by

$$\log y_{jt} = const + \phi_a a_j + \phi_1 x_{jt}^1 + \phi_2 x_{jt}^2, \tag{18}$$

for some coefficients (ϕ_a, ϕ_1, ϕ_2). It follows that i's posterior about $\log y_{jt}$ is log-normal, with

$$\mathbb{E}_{it}[\log y_{jt}] = const + \phi_a \mathbb{E}_{it}[a_{jt}] + \phi_1\left(a_i + \mathbb{E}_{it}[u_{jt}^1]\right) + \phi_2\left(x_{it}^1 + \mathbb{E}_{it}[\xi_t] + \mathbb{E}_{it}[u_{jt}^2]\right)$$

We henceforth ignore the constant terms ($const$) in order to simplify the exposition.

Let $\gamma_1 \equiv \sigma_{u1}/\sigma_A$, $\gamma_2 \equiv \sigma_{u2}/\sigma_A$, $\gamma_\xi \equiv \sigma_\xi/\sigma_A$ denote the relative noise ratios. Then,

$$\mathbb{E}_{it}[a_{jt}] = \frac{1}{1+\gamma_1^2} x_{it}^1, \quad \mathbb{E}_{it}[u_{jt}^1] = \frac{\gamma_1^2}{\gamma_1^2+\gamma_2^2+\gamma_\xi^2}\left(x_{it}^2 - a_i\right), \quad \mathbb{E}_{it}[\xi_t] = \frac{\gamma_\xi^2}{\gamma_1^2+\gamma_2^2+\gamma_\xi^2}\left(x_{it}^2 - a_i\right), \quad \mathbb{E}_{it}[u_{jt}^2] = 0$$

Substituting these expressions into (17) gives us

$$\log y_{it} = \frac{1-\alpha}{1-\vartheta}\, a_i + \alpha\left[\phi_a \frac{1}{1+\gamma_1^2} x_{it}^1 + \phi_1\left(a_i + \frac{\gamma_1^2}{\gamma_1^2+\gamma_2^2+\gamma_\xi^2}\left(x_{it}^2 - a_i\right)\right) + \phi_2\left(x_{it}^1 + \frac{\gamma_\xi^2}{\gamma_1^2+\gamma_2^2+\gamma_\xi^2}\left(x_{it}^2 - a_i\right)\right)\right]$$

By symmetry to (18), i's output must satisfy $\log y_{it} = \phi_a a_i + \phi_1 x_{it}^1 + \phi_2 x_{it}^2$. For this to coincide with the above condition for every z, it is necessary and sufficient that the coefficients (ϕ_a, ϕ_1, ϕ_2) solve the following system:

$$\phi_a = \frac{1-\alpha}{1-\vartheta} + \alpha\phi_1 - \phi_2, \qquad \phi_1 = \alpha\left(\phi_a \frac{1}{1+\gamma_1^2} + \phi_2\right), \qquad \phi_2 = \alpha\left(\phi_1 \frac{\gamma_1^2}{\gamma_1^2+\gamma_2^2+\gamma_\xi^2} + \phi_2 \frac{\gamma_\xi^2}{\gamma_1^2+\gamma_2^2+\gamma_\xi^2}\right)$$

The unique solution to this system gives us the following equilibrium coefficients.

$$\phi_a = \frac{(1-\alpha)\left(1+\gamma_1^2\right)\left(\left(1-\alpha^2\right)\gamma_1^2+\gamma_2^2+(1-\alpha)\gamma_\xi^2\right)}{(1-\vartheta)\left((1-\alpha^2)\left(\gamma_1^4+\gamma_2^2+(1-\alpha)\gamma_\xi^2\right)+\gamma_1^2\left(1-\alpha^2+\gamma_2^2+(1-\alpha)\gamma_\xi^2\right)\right)} > 0 \tag{19}$$

$$\phi_1 = \frac{(1-\alpha)\alpha\left(\gamma_1^2+\gamma_2^2+(1-\alpha)\gamma_\xi^2\right)}{(1-\vartheta)\left((1-\alpha^2)\left(\gamma_1^4+\gamma_2^2+(1-\alpha)\gamma_\xi^2\right)+\gamma_1^2\left(1-\alpha^2+\gamma_2^2+(1-\alpha)\gamma_\xi^2\right)\right)} > 0 \tag{20}$$

$$\phi_2 = \frac{(1-\alpha)\alpha^2\gamma_1^2}{(1-\vartheta)\left((1-\alpha^2)\left(\gamma_1^4+\gamma_2^2+(1-\alpha)\gamma_\xi^2\right)+\gamma_1^2\left(1-\alpha^2+\gamma_2^2+(1-\alpha)\gamma_\xi^2\right)\right)} > 0 \tag{21}$$

Furthermore, the expected equilibrium price must satisfy $\mathbb{E}_{it} \log p_{it} = \frac{1-\vartheta}{\vartheta} \log y_{it} - \frac{1}{\vartheta} a_i$. Using the above results, we have that $\mathbb{E}_{it} \log p_{it} = -\psi_a a_i + \psi_1 x_{it}^1 + \psi_2 x_{it}^2$, with

$$\psi_a = -\left(\frac{1-\vartheta}{\vartheta}\phi_a - \frac{1}{\vartheta}\right), \qquad \psi_1 = \frac{1-\vartheta}{\vartheta}\phi_1 > 0, \qquad \text{and} \qquad \psi_2 = \frac{1-\vartheta}{\vartheta}\phi_2 > 0.$$

To sign the coefficient ψ_a, it is straightforward to check the following: (i) ψ_a is strictly decreasing in γ_2, and (ii) $\lim_{\gamma_2 \to 0} \psi_a > \lim_{\gamma_2 \to \infty} \psi_a > 0$. Together, this implies that ψ_a is everywhere positive.

Given the log-linear structure of equilibrium output and the log-normal specification for productivity and the noises, we find that aggregate output is given by $\log Y_t = \chi_0 + \chi_\xi \xi_t$, where $\chi_0 \equiv \phi_0 + \frac{1}{2}\left((\phi_a + \phi_1 + \phi_2)^2 + (\phi_1 + \phi_2)^2 \gamma_1\right)$ and $\chi_\xi = \phi_2$.

Next, due to the log-normal shock and information structure, we can infer that $\log \mathbb{E}_{it} y_{jt} = \mathbb{E}_{it} \log y_{jt} + const$, where $const$. Furthermore, from (17) we have that island i's belief j's log output must satisfy

$$\mathbb{E}_{it}\left[\log y_{jt}\right] = \frac{1}{\alpha}\left(\log y_{it} - (1-\alpha)\frac{1}{1-\vartheta}\, a_i\right)$$

Aggregating the above across islands and using the fact that the cross-sectional average of a_i is fixed, we get that $\log \mathcal{B}_t = const + \frac{1}{\alpha}\log Y_t$, which verifies that $\log \mathcal{B}_t$ is also a linear function of ξ_t.

Proof of Proposition 4. Let $\mathbf{x}_{it} \equiv \left(a_i, x_{it}^1, x_{it}^2, ..., x_{it}^h\right)'$ and note that

$$\mathbf{x}_{it} = \mathbf{M}\xi_t + \mathbf{m}_1 \mathbf{u}_{it} + \mathbf{m_2 u_{jt}} + \mathbf{m}_a \mathbf{a}_{ijt}$$

where $\xi_t \equiv \left(\xi_t^1, ..., \xi_t^h\right)', \mathbf{u}_{it} = \left(u_{it}^1, ..., u_{it}^h\right)', \mathbf{u}_{jt} = \left(u_{jt}^1, ..., u_{jt}^h\right)'$, and $\mathbf{a}_{ij} = (a_i, a_j)'$, and where $\mathbf{M}, \mathbf{m}_1, \mathbf{m}_2, \mathbf{m}_a$ are some fixed matrices full of zeros and ones.

We guess and verify a log-linear equilibrium under the log-normal specification for the shock and information structure. Suppose the equilibrium production strategy of the island of type ω_{jt} takes a log-linear form given by

$$\log y_{jt} = \chi \bar{a}_t + \phi \mathbf{x}_{jt}$$

for some coefficients $\chi \in \mathbb{R}$ and $\phi = (\phi_a, \phi_1, \phi_2, ..., \phi_h) \in \mathbb{R}_+^{H+1}$. It follows that i's posterior about $\log y_{jt}$ is log-normal, with

$$\mathbb{E}_{it}\left[\log y_{jt}\right] = \chi \bar{a}_t + \phi \mathbb{E}_{it}\left[\mathbf{x}_{jt}\right] \tag{22}$$

Furthermore, i's conditional expectation of $\mathbf{x}_{jt}$ is simply the projection of $\mathbf{x}_{jt}$ on $\mathbf{x}_{it}$: $\mathbb{E}_{it}\left[\mathbf{x}_{jt}\right] = \mathbf{H}\mathbf{x}_{it}$, where $\mathbf{H}$ is the relevant projection matrix. Substituting these expressions into (17) gives us

$$\log y_{it} = (1-\alpha)\,\frac{1}{1-\vartheta}\,(a_i + \bar{a}_t) + \alpha\left[\chi \bar{a}_t + \phi \mathbf{H}\mathbf{x}_{it}\right]$$

For this to coincide with $\log y_{it} = \chi \bar{a}_t + \phi \mathbf{x}_{it}$ for every ω_{it}, it is necessary and sufficient that the coefficients χ and ϕ are given the solution to the following system:

$$\chi = (1-\alpha)\,\frac{1}{1-\vartheta} + \alpha\chi \qquad \text{and} \qquad \phi = (1-\alpha)\,\frac{1}{1-\vartheta}\mathbf{e}_1 + \alpha\,(\phi\mathbf{H})', \tag{23}$$

where $\mathbf{e}_1$ is a column vector of length $h+1$ composed of zeros except for a unit in the first position. Finally, noting that $\int \mathbf{x}_{it} di = \mathbf{M}\xi_\mathbf{t}$,[31] we find that aggregate output is given by

$$\log Y_t = \chi \bar{a}_t + \phi \mathbf{M}\xi_t. \tag{24}$$

[31]Whenever we refer to the cross-sectional average $\int X_{it} di$ of some island-specific variable $X_{it} = X(\omega_{it})$, we mean the expectation of $X(\omega_{it})$ conditional on the aggregate state. That is, $\int X_{it} di \equiv \int X(\omega)\mathcal{P}(\omega|\xi^t, \bar{a}^t)d\omega$.

Furthermore, solving condition (8) for $\mathbb{E}_{it}[\log y_{jt}]$, we have that i's belief of j's output satisfies $\mathbb{E}_{it}[\log y_{jt}] = \frac{1}{\alpha}(\log y_{it} - \frac{1-\alpha}{1-\vartheta}\ (a_i + \bar{a}_t))$. The corresponding aggregate therefore satisfies

$$\log \mathcal{B}_t = \frac{1}{\alpha}\left(\log Y_t - \frac{1-\alpha}{1-\vartheta}\bar{a}_t\right) = \frac{1}{\alpha}\left(\chi - \frac{1-\alpha}{1-\vartheta}\right)\bar{a}_t + \frac{1}{\alpha}\phi\mathbf{M}\boldsymbol{\xi}_\mathbf{t} = \chi\bar{a}_t + \frac{1}{\alpha}\phi\mathbf{M}\boldsymbol{\xi}_\mathbf{t} \tag{25}$$

where we abstract from the constant (see the proof of Proposition 3) and where the last equation uses the fact that $\chi = 1/(1-\vartheta)$ from (23). The result then follows from (24) and (25) once we define the composite shock as $\bar{\xi}_t \equiv \frac{1}{\alpha}\phi\mathbf{M}\xi_t$ and let $\Phi \equiv \chi = \frac{1}{1-\vartheta}$ and $\Psi \equiv \alpha$.

Finally, let us characterize the average forecast of aggregate output, defined as

$$\log \mathcal{B}'_t \equiv \int \mathbb{E}_{it}[\log Y_t]di.$$

The goal is to show that this average forecast, which may be easier to observe in survey evidence, can be thought as a noisy empirical proxy of $\mathcal{B}_t$, which is the relevant belief aggregate in the model. By projecting ξ_t on $\mathbf{x}_{it}$, we get $\mathbb{E}\left[\boldsymbol{\xi}_\mathbf{t}|\omega_{it}\right] = \mathbf{B}\mathbf{x}_{it}$ for some matrix $\mathbf{B}$. It follows that

$$\mathbb{E}_{it}[\log Y_t] = \Phi\bar{a}_t + \phi\mathbf{M}\mathbf{B}\mathbf{x}_{it}$$

and therefore $\log \mathcal{B}'_t = \Phi\bar{a}_t + \phi\mathbf{MBM}\boldsymbol{\xi}_\mathbf{t}$. Since $\bar{\xi}_t \equiv \frac{1}{\alpha}\phi\mathbf{M}\xi_t$ and $\phi\mathbf{MBM}\boldsymbol{\xi}_\mathbf{t}$ are both functions of ξ_t, and the latter is orthogonal to $\bar{a}_t$, we can regress $\phi\mathbf{MBM}\boldsymbol{\xi}_\mathbf{t}$ on $\bar{\xi}_t$ to obtain $\phi\mathbf{MBM}\boldsymbol{\xi}_\mathbf{t} = \Lambda\bar{\xi}_t + \upsilon_t$, where $\Lambda \equiv Cov\left(\phi\mathbf{MBM}\boldsymbol{\xi}_\mathbf{t}, \frac{1}{\alpha}\phi\mathbf{M}\boldsymbol{\xi}_\mathbf{t}\right)/Var\left(\frac{1}{\alpha}\phi\mathbf{M}\boldsymbol{\xi}_\mathbf{t}\right)$ is a scalar and where υ_t is a linear function of ξ_t that is orthogonal to both $\bar{\xi}_t$ and $\bar{a}_t$. We thus get

$$\log \mathcal{B}'_t = \Phi\bar{a}_t + \Lambda\bar{\xi}_t + \upsilon_t,$$

which represents a noisy signal of $\log \mathcal{B}_t$.

Proof of Proposition 5. Once we fix the local beliefs of an island as in (11), the characterization of the local employment, wages, and output in that island follows the same steps as in our baseline model. It follows that the equilibrium output of each island is given from condition (8) after replacing the rational expectations $\mathbb{E}_{it}\log y_{jt}$ with the ad-hoc beliefs specified in (11). That is,

$$\log y_{it} = (1-\alpha)\frac{1}{1-\vartheta}\log A_{it} + \alpha\hat{\mathbb{E}}_{it}\log y_{jt},$$

where $\log A_{it} = \bar{a}_t + a_i$ is the local TFP. Aggregating this condition across all islands (and ignoring as always the constants) gives $\log Y_t = (1-\alpha)\frac{1}{1-\vartheta}\bar{a}_t + \alpha\log\mathcal{B}_t$, where $\bar{a}_t$ is the aggregate TFP shock, while aggregating (11) gives $\log\mathcal{B}_t = \log Y_t^* + \zeta_t$, where $\log Y_t^*$ is the aggregate level of output in the frictionless, rational-expectations benchmark. The result then follows for noting that $\log Y_t^* = \frac{1}{1-\vartheta}\bar{a}_t$ and letting $\Phi \equiv \frac{1}{1-\vartheta}$ and $\Psi \equiv \alpha$, which are the same coefficients as those in Proposition 4.

Proof of Proposition 6. *Part (i).* For any period and any history up to that point, the type of an island belongs to the following set:

$$\bar{\Omega} \equiv \left\{\omega_U^N, \omega_{U+}^N, \omega_P^N, \omega_{P+}^N, \omega_F^N; \omega_U^S, \omega_{U+}^S, \omega_P^S, \omega_{P+}^S, \omega_F^S\right\},$$

where, for each group $g \in \{N, S\}$, ω_U^g are uninformed islands that are matched with a uninformed island from their group, ω_{U+}^g are uninformed islands that are matched with a partially informed island, ω_P^g are partially informed islands that are matched with an uninformed island; ω_{P+}^g are partially informed islands that are matched with a partially informed island from the *other* group; and ω_F^g are fully informed that are matched with a fully informed island from their group.

The period-t cross-sectional distribution of types is thus summarized in a vector $\mathbf{m}_t \in \Delta(\bar{\Omega})$, with the n-th element of this vector giving the fraction of islands whose types is the n-th element of $\bar{\Omega}$. The dynamics of $\mathbf{m}_t$ follows directly from the presumed matching technology.

Clearly, ω_F^N and ω_F^S are absorbing states for, respectively, the North and the South. Along with the fact that $\mu_0 = \chi > 0$, this proves that μ_t must eventually decrease and must converge to zero as $t \to \infty$. Finally, the fact that μ_t must initially increase follows from the assumption $\mu_0 = \chi < 1/2$.

Part (ii). To understand the determination of equilibrium output, consider first all the matches between islands of types ω_{P+}^N and ω_{P+}^S. These matches are, in effect, identical to those featured in Section (4). The equilibrium output for these types must therefore satisfy $\log y(\omega_{P+}^N) = \phi_a a_N + \phi_1 x_N^1 + \phi_2 x_N^2$ and $\log y(\omega_{P+}^S) = \phi_a a_S + \phi_1 x_S^1 + \phi_2 x_S^2$, where the coefficients (ϕ_a, ϕ_1, ϕ_2) are given in (19)-(21). For all other matches, on the other hand, it is straightforward to check that output is given either by $\phi_a a_N$ (for the Northern islands) or $\phi_a a_S$ (for the Southern islands), where $\phi_a = \frac{1}{1-\vartheta}$. We thus infer that local output is given as follows:

$$\log y_{it} = \begin{cases} \phi_a a_N + \phi_1 x_N^1 + \phi_2 x_N^2 & \text{if } \omega_{it} = \omega_{P+}^N, \\ \phi_a a_S + \phi_1 x_S^1 + \phi_2 x_S^2 & \text{if } \omega_{it} = \omega_{P+}^S, \\ \phi_a a_i & \text{otherwise} \end{cases} \tag{26}$$

Aggregating this across all islands, we obtain

$$\log Y_t = \phi_a \bar{a} + \mu_t \left[\phi_1 \bar{\varepsilon} + \phi_2 \xi\right]$$

where $\bar{a} \equiv \frac{1}{2}(a_N + a_S)$ and $\bar{\varepsilon} \equiv \frac{1}{2}(\varepsilon_1 + \varepsilon_2)$, and where μ_t is the fraction of islands with types either ω_{P+}^N or ω_{P+}^S. The result then follows by letting $\Phi \equiv \phi_2$.

Proof of Proposition 7. By combining the optimality conditions for the final-good firms with market clearing (trade balance), we get

$$h_{it} = (1-\eta) q_{it}, \quad h_{it}^* = \eta q_{jt}, \quad p_{it} = q_{it}^{-\eta} q_{jt}^{\eta}, \quad \text{and} \quad p_{it}^* = q_{it}^{1-\eta} q_{jt}^{\eta-1}$$

This is similar to the baseline model; we only have to re-interpret the consumption goods in that model as the intermediate inputs in the present model.

Consider now the behavior of the intermediate-good firms. The first-order conditions with respect to labor, the capital stock, and the rate of capital utilization are, respectively, as follows:

$$\begin{aligned}
\mathbb{E}_{it}\left[\lambda_{it} w_{it}\right] &= \mathbb{E}_{it}\left[\lambda_{it} p_{it}\right] \theta \frac{q_{it}}{n_{it}} \\
\mathbb{E}_{it}\left[\lambda_{it}(r_{it} + \Delta\left(e_{it}\right))\right] &= \mathbb{E}_{it}\left[\lambda_{it} p_{it}\right] (1-\theta) \frac{q_{it}}{k_{it}} \\
\mathbb{E}_{it}\left[\lambda_{it} \Delta'\left(e_{it}\right) k_{it}\right] &= \mathbb{E}_{it}\left[\lambda_{it} p_{it}\right] (1-\theta) \frac{q_{it}}{e_{it}}
\end{aligned}$$

where λ_{it} is the marginal value of wealth on island i. These conditions simply state that the expected marginal costs of labor, capital, and capital utilization are equated with their respective expected marginal revenue products, which in turn depend on the island's expected terms of trade.

Next, on the household's side, the Envelope condition, the optimality condition for labor, and the Euler condition give the following:

$$\begin{aligned}
\lambda_{it} &= U'(c_{it}) \\
V'(n_{it}) &= \mathbb{E}_{it}\left[\lambda_{it} w_{it}\right] \\
U'(c_{it}) &= \mathbb{E}_{it}^{2}\left[\beta U'(c_{i,t+1}) \left(1 + r_{i,t+1}\right)\right]
\end{aligned}$$

where, recall, $\mathbb{E}_{it}^{2}$ denotes the expectation conditional on stage-2 information.

Combining the aforementioned conditions, using $p_{it} q_{it} = q_{it}^{1-\eta} q_{jt}^{\eta} = \zeta y_{it}$ where $\zeta = (1-\eta)^{1-\eta} \eta^{\eta}$, and adding the local resource constraint, we get the system of equations in the proposition.

Appendix B: numerical solution of the RBC variant

To simulate the equilibrium dynamics of the RBC variant of Section 7, we first log-linearize conditions in Proposition 7 to get the following linear dynamic system:

$$\epsilon \tilde{n}_{it} = \mathbb{E}_{it}^{1}\left[\tilde{y}_{it} - \gamma \tilde{c}_{it}\right] \tag{27}$$

$$(1+\mu)\, \tilde{e}_{it} = \mathbb{E}_{it}^{1}\left[\tilde{y}_{it} - \tilde{k}_{it}\right] \tag{28}$$

$$\tilde{c}_{it} = \mathbb{E}_{it}^{2}\left[\tilde{c}_{i,t+1} - \tfrac{(1-\beta)}{\gamma}\left(\tilde{y}_{it+1} - \tilde{k}_{i,t+1}\right)\right] \tag{29}$$

$$\bar{c}\tilde{c}_{it} + \bar{k}\tilde{k}_{i,t+1} = \bar{y}\tilde{y}_{it} + \left(1 - \tfrac{1-\beta}{\beta\mu}\right) \bar{k}\tilde{k}_{it} - (1+\mu)\, \tfrac{1-\beta}{\beta\mu} \bar{k}\tilde{e}_{it} \tag{30}$$

$$\tilde{y}_{it} = (1-\eta)\tilde{q}_{it} + \eta \tilde{q}_{jt} \tag{31}$$

$$\tilde{q}_{it} = a_{it} + \theta \tilde{n}_{it} + (1-\theta)\left(\tilde{e}_{it} + \tilde{k}_{it}\right) \tag{32}$$

where the bars denote steady-state values and the tildes denote log-deviations from steady state.

Let $\tilde{\rho}_{it} = \rho(\tilde{y}_{it}, \tilde{k}_{it}, \tilde{e}_{it})$ denote the right-hand side of condition (30); this identifies the overall resources that are available in stage 2, measured in terms of log-deviation from steady state. We conjecture the following island-level policy rules, along with a rule for aggregate capital:

$$\begin{aligned}
(\tilde{e}_{it}, \tilde{n}_{it}, \tilde{q}_{it}) &= f\left(a_{it}, x_{it}, \xi_t; \tilde{k}_{it}, \tilde{K}_t\right) \\
(\tilde{c}_{it}, k_{i,t+1}) &= g\left(\tilde{\rho}_{it}; \xi_t, \tilde{K}_t\right) \\
\tilde{K}_{t+1} &= \Gamma\left(\xi_t, \tilde{K}_t\right)
\end{aligned}$$

where the functions f, g, and Γ are linear. This guess is justified by the following considerations. First, an island's employment, utilization, and production choices during stage 1 depend on its own productivity, its current signal of the productivity of its trading partner, and on the perceived bias in the latter's signal for essentially the same reasons that it does in our baseline model; but now it also depends on its own capital stock, and on the aggregate capital stock, because the former enters local production while the latter is i's best forecast of the capital stock of its trading partner.[32] Second, an island consumption and investment during stage 2 are pinned down by realized resources, for the usual reasons, and by the aggregate state of the economy, for the latter determines i's beliefs of its future terms of trade, local income, and local prices. Finally, the aggregate policy rule for capital obtains from aggregating the corresponding individual policy rules and noting that the cross-sectional average of resources is ultimately pinned down by the current sentiment shock ξ_t and the current aggregate capital $\tilde{K}_t$.

We then solve the equilibrium by the method of undetermined coefficients: we write the policy rules in terms of arbitrary coefficients; we next plug these rules in the aforementioned log-linearized system (27)-(32) along with the definition of $\tilde{\rho}_{it}$ and the aggregation consistency between g and Γ; we then arrive to a system of equations in the aforementioned coefficients, which can be solved for the equilibrium. This procedure is, in effect, quite similar to the way one solves the log-linearized version of the RBC model, except for the extra complication that our log-linearized system embeds also a fixed point between island-specific and aggregate policy rules.

Once we have the policy rules, we create 1000 random time series for the sentiment shock, each of length 1250 periods. For each of these series, we compute the equilibrium time series of all the key macroeconomic variables. We next drop the first 1000 periods, in order to get rid of any dependence on initial conditions, and apply the HP-filter on the last 250 periods (which is approximately as many quarters as in our data), using the conventional weight (1600). We next compute the relevant business-cycle statistics on the HP-filtered series. We finally take averages of these statistics across all 1000 series, and report these averages in the left two columns of Table I in the main text.

Finally, to obtain the empirical counterparts of these statistics, we use the actual U.S. time series data as documented in Smets and Wouters (2007), except for two minor changes. First, we

[32] To understand why the aggregate capital stock K_t is i's best forecast of j's capital stock, recall that we have assumed that the idiosyncratic productivity of an island is i.i.d. over time and across islands. Learning about j's current productivity therefore gives no information about j's history of past productivity shocks and hence also about its capital stock. If, instead, productivity were persistent, then i's best forecasts of j would be a linear combination of the aggregate capital stock and i's signals about j's productivity. This would complicate a bit the solution, but is unlikely to affect the results.

extend the data through 2012. Second, we correct the population series for the problem identified in Edge and Gürkaynak (2010).[33]

Our data therefore covers the period 1948Q3-2012Q2, are at a quarterly frequency, and are seasonally adjusted. Output, consumption, and investment are measured by, respectively, GDP, Personal Consumption Expenditures, and Fixed Private Investment; these variables are taken from the BEA, are deflated by the BEA's GDP Price Deflator, and are normalized in per-capita terms. Employment is measured by Nonfarm Hours, as taken from the U.S. Department of Labor. Finally, the population series we use is the smoothed Civilian Noninstitutional Population aged 16 and over, provided by Edge, Gürkaynak, and Kisacikoğlu (2012) and obtained from the Federal Reserve Board. The same HP-filter is applied to this data as with the model's simulated data, and the corresponding statistics are finally reported in the right two columns of Table I.

References

[1] Akerlof, George, and Robert Shiller (2009), *Animal Spirits: How Human Psychology Drives the Economy, and Why it Matters for Global Capitalism*, Princeton: Princeton University Press.

[2] Angeletos, George-Marios, and Jennifer La'O (2009), "Noisy Business Cycles," *NBER Macroeconomics Annual 2009*, 24, 319-378.

[3] Angeletos, George-Marios, and Jennifer La'O (2011), "Decentralization, Communication, and the Origins of Fluctuations," NBER working paper 17060.

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

[5] Azariadis, Costas (1981), "Self-fulfilling Prophecies," *Journal of Economic Theory* 25 380-96.

[6] Barsky, Robert, and Eric Sims (2009), "News Shocks," NBER Working Papers 15312.

[7] Beaudry, Paul, and Franck Portier (2006), "Stock Prices, News, and Economic Fluctuations," *American Economic Review* 96, 1293-1307.

[8] Benhabib, Jess, and Roger Farmer (1994), "Indeterminacy and Increasing Returns," *Journal of Economic Theory* 63, 19-41.

[9] Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo (2011), "Understanding Booms and Busts in Housing Markets," Duke/Northwestern mimeo.

[33] Smets and Wouters (2007) uses the population series reported by the BLS. Edge and Gürkaynak (2010) notice that this series exhibits a number of extremely sharp spikes caused by the census picking up previously uncultured population growth. If we use the original series of Smets and Wooters (2007), the results in Table I remain nearly identical.

[10] Cass, David, and Karl Shell (1983), "Do Sunspots Matter?" *Journal of Political Economy* 91, 193-227.

[11] Chari, V. V., Patrick Kehoe, and Ellen McGrattan (2007), "Business Cycle Accounting," *Econometrica* 75, 781-836.

[12] Christiano, Lawrence, Cosmin Ilut, Roberto Motto, and Massimo Rostagno (2008), "Monetary policy and stock market boom-bust cycles," ECB Working Paper No. 955.

[13] Cooper, Russell, and Andrew John (1988), "Coordinating Coordination Failures in Keynesian Models," *Quarterly Journal of Economics* 103, 441-463.

[14] Diamond, Peter (1982), "Aggregate Demand Management in Search Equilibrium," *Journal of Political Economy* 90(5), 881-894.

[15] Edge, Rochelle, and Refet Gurkaynak (2010), "How Useful Are Estimated DSGE Model Forecasts for Central Bankers?," *Brookings Papers on Economic Activity*, 209-244.

[16] Gali, Jordi, and Thijs van Rens (2010), "The Vanishing Procyclicality of Labor Productivity," CREI mimeo.

[17] Grossman, Sanford J. (1981), "An Introduction to the Theory of Rational Expectations under Asymmetric Information," *Review of Economic Studies* 48(4), 541-559.

[18] Guesnerie, Roger, and Michael Woodford (1992), "Endogenous Fluctuations," in J.-J. Laffont, ed., *Advances in Economic Theory: Proceedings of the Sixth World Congress of the Econometric Society*, Cambridge: Cambrige Univ. Press.

[19] Jaimovich, Nir, and Sergio Rebelo (2009), "Can News About the Future Drive the Business Cycle?," *American Economic Review*, 99(4), 1097-1118.

[20] King, Robert, and Sergio Rebelo (2000), "Resusitating Real Business Cycles," in J. B. Taylor & M. Woodford, eds., *Handbook of Macroeconomics*, Elsevier.

[21] Kocherlakota, Narayana (2010), "Modern Macroeconomic Models as Tools for Economic Policy," *The Region* May 2010, 5-21.

[22] Krusell, Per, and Anthony Smith (1998), "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy* 106, 867-896.

[23] La'O, Jennifer (2010), "Collateral Constraints and Noisy Fluctuations," Chicago Booth mimeo.

[24] Lorenzoni, Guido (2010), "A Theory of Demand Shocks," *American Economic Review* 99(5), 2050-2084.

[25] Lucas, Robert E., Jr. (1972), "Expectations and the Neutrality of Money," *Journal of Economic Theory* 4, 103-124.

[26] Lucas, Robert E., Jr. (2001), April 2001 lecture given at Trinity University.

[27] Morris, Stephen, and Hyun Song Shin (2002), "The Social Value of Public Information," *American Economic Review* 92, 1521-1534.

[28] Morris, Stephen, and Hyun Song Shin (2003), "Global Games–Theory and Applications," in M. Dewatripont, L. Hansen, and S. Turnovsky, eds., *Advances in Economics and Econometrics*, 8th World Congress of the Econometric Society, Cambridge: Cambridge University Press.

[29] Shiller, Robert (2005), *Irrational Exuberence*, 2nd ed., Princeton: Princeton University Press.

[30] Shimer, Robert (2009), "Convergence in Macroeconomics: The Labor Wedge," *American Economic Journal: Macroeconomics* 1, 280-297.

[31] Townsend, Robert M. (1983), "Forecasting the Forecasts of Others," *Journal of Political Economy* 91, 546-588.

[32] Woodford, Michael (2003), "Imperfect Common Knowledge and the Effects of Monetary Policy," in P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, eds., *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, Princeton: Princeton University Press.