

Decentralization, Communication, and the Origins of Fluctuations*

George-Marios Angeletos
MIT and NBER

Jennifer La'O
Chicago Booth

June 29, 2011

Abstract

We consider a class of convex, competitive, neoclassical economies in which agents are rational; the equilibrium is unique; there is no room for randomization devices; and there are no shocks to preferences, technologies, endowments, or other fundamentals. In short, we rule out *every* known source of macroeconomic volatility. And yet, we show that these economies can be ridden with large and persistent fluctuations in equilibrium allocations and prices.

These fluctuations emerge merely because decentralized trading impedes communication and, in so doing, permits self-fulfilling waves of optimism and pessimism to obtain even when the equilibrium is unique. What is more, these fluctuations require no divergence between private and social motives. They therefore need not open the door to stabilization policy, despite their apparent arbitrariness and their total disconnect from fundamentals.

*This paper supersedes earlier working papers that circulated in 2008 and 2009 under the title “Sentiments” and that contained different models but a similar overall message. For stimulating discussions and feedback, we thank Robert Barro, Roland Benabou, Olivier Blanchard, John Campbell, Arnaud Costinot, Ben Friedman, John Geanakoplos, Emmanuel Farhi, Erik Hurst, David Laibson, Patrick Kehoe, N. Gregory Mankiw, Stephen Morris, Alessandro Pavan, Richard Rogerson, Hyun Song Shin, Ivan Werning, Mirko Wiederholt, and seminar participants at Bern, Harvard, MIT, PennState, Princeton, the Atlanta and New York FRBs, the 2011 UCSB-LAEF conference on Information and Macroeconomics, the 2010 Bank of Portugal Monetary Economics Conference, the 2010 Midwest Macroeconomics Meeting at Michigan State, the 2010 CRETE conference in Tinos, the 2011 UCSB LAEF Conference on “Putting Information Into (or Taking it out of) Macroeconomics”, the 2011 CREI Conference on “Information, Beliefs and Expectations in Macroeconomics”, and the 2011 North American Meeting of the Econometric Society at Washington University in St. Louis. Finally, we owe special thanks to Venky Venkateswaran, Tarek Hassan, and Per Krusell for conference discussions of our paper.

1 Introduction

The last three decades of macroeconomic research have significantly progressed our understanding of the business-cycle implications of a variety of propagation mechanisms. Yet, no serious progress has been made in our formalizations of the *origins* of fluctuations: much like the prototypical RBC model, any state-of-the-art DSGE model ultimately attributes the bulk of fluctuations to exogenous shocks in technologies and preferences, or some mysterious “wedges”.

The notion that fluctuations are driven by shocks to fundamentals such as preferences and technologies is ubiquitous within the modern macroeconomic paradigm, but remains hardly convincing outside its realm.¹ In fact, macroeconomists themselves often appear to use these notions only as convenient short-cuts that help them match the observed fluctuations within their formal models. Once turning to policy matters, a literal interpretation of these shocks is quickly abandoned. Instead, these shocks are customarily re-interpreted as shifts in “aggregate demand” and “consumer confidence”, or as “cost-push shocks”, “labor-market shocks”, and “asset-market shocks”—underscoring that these shocks are only metaphors for forces we have failed to comprehend and model more properly.²

These short-cuts are problematic. By conveniently recasting the residuals of one model as the structural shocks of another, they obscure the boundaries between the successes and the failures of our understanding. Perhaps more importantly, they put the policy lessons of the last three decades in serious jeopardy, in so far these lessons appear to depend on literal interpretations of the aforementioned shocks. For instance, what is the precise meaning of the optimality of flexible-price allocations in prototypical RBC and New-Keynesian models if the preference and technology shocks that populate these models are proxies for, say, “animal spirits” and “irrational exuberance”?

Motivated by these considerations, this paper develops a novel formalization of the origins of fluctuations—one that builds upon the neoclassical framework but dispenses entirely with the ubiquitous notions of aggregate shocks to fundamentals such as preferences and technologies.

More specifically, we consider a class of convex, competitive, rational-expectations economies in which the equilibrium is unique and there is no room for either randomization devices (sunspots and lotteries) or deterministic cycles and chaotic dynamics. We further impose that there are no shocks to preferences, technologies, government policies, or any other kind of fundamentals. In short, we stay comfortably within the boundaries of the neoclassical framework while also ruling out *all* known sources of macroeconomic volatility. And yet, we show that these economies may still exhibit rich and persistent fluctuations in equilibrium allocations and prices.

¹This notion has been criticized by multiple generations of economists over the years. Some sought to deconstruct this notion on the basis of empirical evidence, while others attacked it on the basis of sheer plausibility. Most recently, Krugman (2009) calls this notion “*silly*” and invites a return to the old Keynesian paradigm. In a similar vein, Shiller (2009) argues that “*animal spirits offer an explanation for why we get into recessions in the first place—for why the economy fluctuates as it does.*”

²Consistent with this view, Kocherlakota (2010) writes: “*The sources of disturbances in macroeconomic models are (to my taste) patently unrealistic... Macroeconomists use these notions [of collective shocks to preferences and technologies] only as convenient shortcuts to generate the requisite levels of volatility in endogenous variables.*”

Basic insight. The apparent paradox is explained by a simple, yet powerful insight—one that regards the role that the market mechanism plays in facilitating communication and thereby in facilitating coordination.

This role has been a cornerstone of neoclassical thinking at least since Hayek (1945): markets are supposed to help the economy achieve an efficient utilization of its resources without the need for a “center” either to coordinate their actions or to aggregate the information that is dispersed in society regarding people’s needs, tastes, and abilities. Yet, this idea has been pushed to an absurd extreme within the Arrow-Debreu framework: the notion of a central planner has been replaced with the notion of centralized markets, leaving no space for decentralization to have a meaningful bite on the extent of communication and coordination.

Indeed, in the Arrow-Debreu framework, it is *as if* all agents in the economy gather in a single “conference room” and talk to one another till they reach common knowledge, not only about their tastes, abilities, and other fundamentals, but also—and most crucially—about their intended courses of action. But, if agents can reach common knowledge about their intended courses of action and if, in addition, their incentives are aligned with social objectives, it should be no surprise that they are also able to coordinate on first-best outcomes.³

The presumption of this kind of flawless, instantaneous communication is ubiquitous in modern macroeconomics—but it is only a convenient abstraction. In reality, no agent trades with, or talks to, all other agents at the same time. Rather, most market and other social interactions are segmented and people are unable to reach common knowledge about either their opinions or their likely courses of action. This kind of imperfect communication is manifested in the anxiety with which financial markets, the media, and the general public monitor macroeconomic statistics for clues about aggregate economic activity. If communication had been perfect, what could market participants possibly learn from such signals of one another’s choices?

Our contribution is to show that this fact alone explains why the “invisible hand” may fail to coordinate the economy on first-best outcomes, making room for self-fulfilling beliefs and “animal spirits” to emerge even when the equilibrium is unique and private incentives are perfectly aligned with social ones. Imperfect communication *means* imperfect coordination.

Formalization and preview. We embed these ideas in an elementary DSGE model as follows. The economy consists of multiple “islands” (Lucas, 1972; Lucas and Prescott, 1974), each populated by a large number of firms and households. Islands are heterogeneous, and each one produces a specialized good, but wishes to consume also the good of other islands. This structure thus captures three basic ingredients of any realistic economy: heterogeneity, specialization, and trade.

Next, building on the search-theoretic tradition, we capture the decentralization of market interactions by introducing random matching: certain trades take place through random pair-wise

³The preceding discussion is not intended to trivialize the welfare theorems: their value rests precisely on showing how the “invisible hand” helps align private incentives with social objectives. Also, it is well understood that deviations from first best can obtain from misalignment between private and social incentives, whether such misalignment originates from externalities, market power, or financial frictions. By abstracting from this kind of market failures, we seek to isolate the positive and normative implications of imperfect communication.

matches (Diamond, 1982; Kiyotaki and Wright, 1998; Lagos and Wright, 2005). We nevertheless maintain perfect competition because this matching involves pairs of islands (Lucas, 1972) instead of pairs of individuals. We furthermore rule out thick-market externalities, as well as any shocks to the matching technology. Random matching thus serves precisely, and only, two purposes in our framework: it introduces idiosyncratic trading uncertainty; and it impedes communication.

Once agents from different islands “meet” and trade, they also communicate: they share any information they may have previously acquired either by past trades with other islands or by any other means. In effect, this means that the islands within any given match reach the same beliefs about all relevant economic outcomes once they trade. Nonetheless, these islands had to take some of their employment and production decisions *prior* to their meeting. These decisions could thus have depended on differential beliefs about the state of the economy. As these decisions ultimately determine the actual terms of their trade, a certain form of extrinsic uncertainty emerges along the unique equilibrium: beliefs of terms of trade are no more spanned by beliefs of fundamentals.

Equilibrium allocations and prices are thereby shown to respond to a certain type of extrinsic shocks that we call “sentiments”. These shocks are akin to sunspots, except that they operate in a unique-equilibrium economy and, importantly, are not common knowledge—if people could reach common knowledge about these shocks, they would also reach common knowledge about their courses of action, in which case the equilibrium impact of these shocks would vanish and the economy would attain first-best outcomes. These shocks thus formalize the uncertainty economic agents face about aggregate economic activity because, and only because, of imperfect communication.

As different agents meet and trade, communication does take place. Information about the aforementioned sentiment shocks thus diffuses in the economy in a manner akin to the “percolation of information” in Duffie and Manso (2007) and Duffie, Giroux, and Manso (2009). This communication, however, does not regard the aggregate fundamentals: the latter are common knowledge, hence there is nothing to be learned about them. Rather, the role played by communication is to help people share their beliefs about one another’s choices and, in so doing, to facilitate coordination.

The resulting fluctuations can feature rich, hump-shaped dynamics, capturing the “waves of optimism and pessimism” that many observers associate with business-cycle and asset-price phenomena. These waves may initially hit only a small fraction of the population, but may build force later on as they spread throughout the population in a manner akin to “fads”. Sooner or later, however, these waves die off. A boom looks like a period during which people get increasingly “exuberant”, only to “come back to their senses” after a while; a recession looks like a period during which people lose their “confidence” in the economy, only to regain it later on. In line with Keynesian tradition, these fluctuations can be interpreted as “animal spirits”, or as shifts in “aggregate demand”.

From a game-theoretic perspective, these fluctuations are rationalized by movements in second- and higher-order beliefs that are correlated in the cross-section of the economy and that obtain despite the absence of correlated shocks to either first-order beliefs or the true payoffs. In other words, booms are periods where people become more optimistic about other people’s beliefs and choices, without however being any more optimistic about the underlying fundamentals.

This reveals a connection between our paper and the pertinent literature on informational frictions and higher-order uncertainty, most notably Townsend (1985), Morris and Shin (2002, 2003), and Woodford (2003). We elaborate on this connection in due course. For now, we point out that this literature fails to address the key challenge that motivates our paper: in this literature, pretty much as in the standard paradigm, aggregate fluctuations obtain *only* because of aggregate uncertainty in fundamentals.⁴ By contrast, our paper dispenses entirely with this kind of aggregate uncertainty. Instead, it shows how *aggregate* fluctuations can obtain from purely *idiosyncratic* trading risk (herein formalized as random matching) once communication is imperfect.

In so doing, our paper complements Diamond (1982), who also sought to obtain aggregate fluctuations from trading frictions but did so only by introducing non-convexities in the matching technology and multiple equilibria. By contrast, our contribution highlights how trading frictions open the door to aggregate fluctuations merely by impeding communication.

More broadly, our paper complements a vast earlier literature that sought to formalize the Keynesian notions of “coordination failure” and “animal spirits” in models with multiple equilibria.⁵ While we are sympathetic to this class of models, our contribution here is to show how self-fulfilling phenomena are systemic features of *any* decentralized market economy—including convex, unique-equilibrium economies that have so far been considered to be inconsistent with such phenomena.

Last but not least, the self-fulfilling phenomena we document need not reflect any divergence between private and social incentives. Rather, they are consistent with a notion of constrained efficiency that embeds the bite that decentralization has on communication—market outcomes can be improved upon *only* by facilitating more communication. By implication, there is no room for fiscal, monetary, or regulatory policies that seek to manipulate people’s incentives and the allocation of resources while taking the extent of decentralization and communication as given. Our theory therefore helps rehabilitate the Keynesian notions of “coordination failure” and “animal spirits” at the heart of the neoclassical paradigm, while also rebuking the conventional wisdom that these forces are prima-facia evidence of the need for active stabilization policy.

Layout. The rest of the paper is organized as follows. Section 2 use a pedagogical example to illustrate the key ideas behind our theory. Section 3 introduces the baseline model. Section 4 characterizes the equilibrium and presents our core positive results. Section 5 discusses complementary interpretations of our results. Section 6 studies efficiency. Section 7 concludes. The Appendix includes any proofs omitted in the main text.

⁴This uncertainty takes different facades depending on the application of interest: it may regard shocks to monetary policy (Lucas, 1972, Woodford, 2003, Mackowiak and Wiederholt, 2009, Amador and Weill, 2010); shocks to aggregate productivity (Townsend, 1985, Angeletos and La’O, 2009, Lorenzoni, 2010); or shocks to abstract payoffs (Morris and Shin, 2002). One way or another, however, fluctuations are driven either by innovations in the true fundamentals, or by noisy news about them, pretty much as in conventional, symmetric-information models that introduce “news shocks” (Beaudry and Portier, 2006; Jaimovich and Rebelo, 2009).

⁵See, inter alia, Azariadis (1981), Benhabib and Farmer (1994, 1999), Cass and Shell (1983), Cooper and John (1988), Diamond and Fudenberg (1989), Guesnerie and Woodford (1992), Howitt and McAfee (1992), Matsuyama (1991), Shell (1977), and Woodford (1991).

2 An example

To understand the basic insight behind our contribution, let us momentarily consider the following pedagogical example. The economy is populated by two groups of agents, “bakers” and “shoemakers”. Bakers produce only bread but like wearing shoes; shoemakers produce only shoes but like eating food. In every period, a baker and a shoemaker are randomly matched together. In the “morning”, they decide how hard to work and how much of their respective goods to produce; in the “afternoon”, they meet and engage in competitive barter trade.

This is similar to Diamond (1982), except that matching and trade take place with probability one: there is no risk of unexplored trading opportunities. Furthermore, there is no aggregate uncertainty in fundamentals: the profile of tastes and abilities among bakers and shoemakers is fixed and common knowledge. Nevertheless, each agent faces idiosyncratic risk about his match: a given baker, for example, could be matched with either a low-ability or a high-ability shoemaker.

Clearly, effort and production choices during the morning depend on the terms of trade are expected to obtain in the afternoon: the baker works harder whenever he expects a higher relative price for bread, the shoemaker works harder whenever he expects a higher relative price for shoes.⁶ In turn, terms of trade are pinned down by effort and production choices: the relative price of bread increases with the supply of shoes, the relative price of shoes increases with the supply of bread.

The question that we then wish to raise is the following. Take a baker-shoemaker pair and fix their tastes and abilities, as well as their beliefs about one another’s tastes and abilities. Can there be events in which both agents become “exuberant” in the sense that they both expect a favorable shift in their terms of trade and, in so doing, end up working harder and producing more?

The answer to this question is negative within the canonical neoclassical paradigm: whenever the baker expects the price of bread to improve, the shoemaker *must* expect the price of shoes to deteriorate. However, this is not merely due to rationality. Rather, it is also because these agents are, in effect, presumed to share common knowledge about their effort levels and the terms of their trade at *all* points of time—even *before* they have the chance to meet and trade.

But now suppose that the baker and the shoemaker are unable to communicate *prior* to their trade. As long as this is the case, they may face non-trivial uncertainty about each other’s effort and production choices, and hence they may reach diverging beliefs about the likely terms of their upcoming trade. As a result, there can exist events in which *both* agents form optimistic beliefs about the terms of their trade, in which case each agent ends up working harder and producing more only because he expects, in effect, his trading partner to do so. Symmetrically, there can exist events in which both agents cut down on their effort and output only because they turn mutually pessimistic about the terms of their trade, without any change in their tastes and abilities.

This example illustrates how imperfect communication opens the door to self-fulfilling beliefs. The rest of our paper embeds this simple but powerful insight within an otherwise conventional macroeconomic model and studies the type of fluctuations that can thus obtain.

⁶This argument presumes that substitution effects dominate income effects, which seems the most plausible scenario; similar arguments, however, can be made in the alternative case.

3 The baseline model

The economy consists of a continuum of islands, indexed by $i \in \mathcal{I} \equiv [0, 1]$. Each island is populated by a representative household and a representative locally-owned firm. Each island produces a single good, which can either be consumed at “home” (by the island that produces it) or be traded for a good produced “abroad” (by some other island). Production exhibits constant returns to scale with respect to two inputs: local labor, which is supplied elastically by the local household, and local land, which is in fixed supply. Time is discrete, indexed by $t \in \{0, 1, \dots\}$, and each period contains two stages. Employment and production choices take place in stage 1, while trading and consumption take place in stage 2. All agents are rational and price-takers, and all markets (whether for labor or for commodities) clear at competitive prices.

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

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

where y_{it} is the quantity produced, n_{it} is the labor input, k_{it} is the land input,⁷ A_i is the local productivity (TFP), and $\vartheta \in (0, 1)$ 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. Note that productivities are allowed to vary in the cross-section of islands, thus allowing for some heterogeneity, but not over time, thus ruling out both aggregate and idiosyncratic productivity shocks.

Households and preferences. Preferences on island i are given by

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

where $\beta \in (0, 1)$ is the discount factor, c_{it} is consumption of the “domestic” good, c_{it}^* is consumption of the “foreign” good, $U(c_{it}, c_{it}^*)$ is the utility flow from these two forms of consumption, n_{it} is labor effort, and $V(n_{it})$ is the implied disutility. U is increasing and concave, while V is increasing and convex. For simplicity, the following functional forms are assumed:

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

where $\eta \in (0, 1)$ parameterizes the extent which there is specialization and trade (the fraction of “domestic” 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

$$p_{it}c_{it} + p_{it}^*c_{it}^* \leq w_{it}n_{it} + r_{it}K + \pi_{it} \quad (2)$$

The left-hand side is total expenditure, with p_{it}^* denoting the local price of the good that the island “imports” from its trading partner during stage 2; the right-hand side is total income, with K denoting the fixed supply of land.

⁷As the notation suggests, we can think of land also as capital, provided we keep it in fixed supply.

Heterogeneity and random matching. Although our approach dispenses with *aggregate* uncertainty in fundamentals, it requires *idiosyncratic* uncertainty in trading opportunities—think, for example, of the idiosyncratic uncertainty firms face about their customers, or of the one that traders face in over-the-counter asset markets. To achieve this, we assume that islands are heterogeneous and that trading opportunities arrive randomly, as the product of random pair-wise matching.

Heterogeneity is herein assumed to be in productivities, but it could as well be in tastes, endowments, or any other primitive characteristic. The cross-sectional distribution of productivities is given by a probability function $\mathcal{F}_A : \mathcal{S} \rightarrow (0, 1)$, with support $\mathcal{S}_A \subset \mathbb{R}_+$ that is finite and has at least two elements. 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 .

The “matching technology” (that is, the probability that any two islands get matched together) is also invariant and common knowledge. For expositional simplicity, we further assume that the matching is uniform: in any date and state, each island has an equal probability to be matched with any other island.

These assumptions rule out aggregate shocks to any type of fundamentals, including the matching technology, as a source of aggregate volatility. They also distinguish our paper from the pertinent literature on informational frictions, which invariably presumes aggregate shocks to fundamentals and removes common knowledge about them.

Timing, information, and communication. As mentioned before, each period has two stages. In stage 1, each island operates in isolation, not having yet the opportunity to meet nor otherwise communicate with its realized match. At this point, the firms and workers of each island make their employment and production choices while facing uncertainty about the likely type of their island’s trading partner—and hence about the terms of trade they will face in stage 2. Once stage 2 arrives, the two islands meet. At this point, they observe each other’s identities, productivities, and production levels, and they share any information they may have accumulated in the past. They then trade their goods within a competitive pair-wise market and they finally consume.

Before any two islands meet, they may have some private information about their likely match. One part of this information could be the product of the communication that takes place through decentralized market interactions; this is an integral part of our analysis. Another part could be the product of choices that we treat as exogenous to our model, such as market search and other forms of information collection.

We formalize the latter kind of information as follows. Pick an arbitrary pair of islands (i, j) that get matched during period t and let (A_i, A_j) be their productivities. In stage 1 of that period (i.e., prior to their actual meeting), the two islands receive private signals about their match. The private signal of island i is denoted by $\tilde{\omega}_{it}$, the one of j by $\tilde{\omega}_{jt}$. These signals may be correlated with one another, as well as with the two islands’ productivities, thus permitting us to capture the beliefs that agents may hold about the abilities, tastes, or other intrinsic characteristics of their likely trading partners. At the same time, these signals may depend on an extrinsic aggregate random variable ξ_t , which is itself uncorrelated with the productivities of any match.

Specifically, we first let ξ_t be drawn from a finite set \mathcal{S}_ξ according to a Markov transition probability function $\mathcal{F}_\xi : \mathcal{S}_\xi^2 \rightarrow [0, 1]$, with $\mathcal{F}_\xi(\xi_t|\xi_{t-1})$ denoting the probability of ξ_t conditional on ξ_{t-1} . We then let $\mathcal{S}_{\tilde{\omega}}$ denote the finite support of the aforementioned match-specific signals and $\mathcal{G} : \mathcal{S}_{\tilde{\omega}}^2 \times \mathcal{S}_A^2 \times \mathcal{S}_\xi \rightarrow [0, 1]$ the distribution, conditional on ξ_t and the productivities of a particular match, that Nature uses to draw the aforementioned signals. That is, $\mathcal{G}(\tilde{\omega}, \tilde{\omega}'|A, A', \xi)$ identifies the probability that any trading pair with productivity levels A and A' will receive signals, respectively, $\tilde{\omega}$ and $\tilde{\omega}'$, given any particular value ξ of the aforementioned extrinsic variable.⁸

Sentiment shocks. By construction, the shock ξ_t affects neither the cross-sectional profile of productivities (which is constant and common knowledge), nor the matching technology. To strengthen the notion that this shock is utterly disconnected from either the true fundamentals or people’s beliefs about the fundamentals, we further impose that the signal structure \mathcal{G} is such that the posterior belief of any island about the productivity of its trading partner is invariant to ξ_t .⁹ Nonetheless, we will soon show that, along the unique equilibrium of the economy, this shock can trigger aggregate variation in people’s beliefs about allocations and prices—and, in so doing, in actual allocations and prices as well—as long as people lack common knowledge about this shock. This shock thus formalizes the extrinsic risk that emerges because, and only because, of imperfect communication. To fix language, we henceforth refer to ξ_t as the “sentiment shock”.

Information sets, contingencies, and equilibrium. We denote with $\xi^t \equiv \{\xi_0, \dots, \xi_t\}$ the history of the sentiment shock up to period t , and with ω_{it}^s the information set (or “private history”, or “local state”, or “type”) of island i in stage s of period t , for any $i \in [0, 1]$, $t \in \{0, 1, \dots\}$, and $s \in \{1, 2\}$. Following the preceding description of the economy, these information sets can be defined recursively as follows. First, $\omega_{i0}^1 = (A_i, \tilde{\omega}_{i0})$; that is, the information set in the beginning of time is merely the local productivity and the local signal about the current match. Next, for all $t \geq 0$, $\omega_{it}^2 = (\omega_{it}^1, \omega_{jt}^1)$, where $j = m(i, t)$ henceforth identifies the realized match (trading partner) of island i in period t ; that is, the stage-2 information set of an island is the combination of its own stage-1 information set and the one of its trading partner. Finally, for all $t \geq 1$, $\omega_{it}^1 = (\omega_{i,t-1}^2, \tilde{\omega}_{it})$; that is, the stage-1 information set in each period after the first one is the stage-2 information set of the previous period plus the local signal about the current match.¹⁰

To simplify notation, we henceforth let $\omega_{it} \equiv \omega_{it}^1$, that is, we drop the $s = 1$ index for the stage-1 information sets. We then denote the support of these sets with \mathcal{S}_t (the support of ω_{it}^2 is then simply the square of \mathcal{S}_t) and denote with $\mathcal{P}_t(\omega_{it}, \omega_{jt}, \xi^t)$ the joint probability of the stage-1 information sets of a trading pair along with the history of the extrinsic shock. The corresponding marginal and conditional probabilities are denoted by $\mathcal{P}_t(\omega_{it})$, $\mathcal{P}_t(\xi^t)$, $\mathcal{P}_t(\omega_{it}|\xi^t)$, $\mathcal{P}_t(\xi^t|\omega_{it})$, and so on.

⁸A law of large numbers is assumed to apply so that $\mathcal{G}(\tilde{\omega}, \tilde{\omega}'|A, A', \xi)$ is also the fraction of such pairs of island that receive signals $(\tilde{\omega}, \tilde{\omega}')$.

⁹This boils down to a restriction on $\mathcal{G}(\tilde{\omega}, \tilde{\omega}'|A, A', \xi)$ so that the implied distribution for A' conditional on $(A, \tilde{\omega})$ is the same as the one conditional on both $(A, \tilde{\omega})$ and ξ . That is, ξ can affect the beliefs that each island forms about the information of other islands, but not about their fundamentals.

¹⁰In principle, the private history of each island also includes the identities of its past trading partners, the identities of the latter’s own past partners, and so on. These elements, however, are irrelevant in equilibrium.

Any allocation and price system can thus be represented with a collection of functions $(n, y, w, r, q, p, p^*, c, c^*)$ such that, for all islands, all dates, and all local histories, $n_{it} = n_t(\omega_{it})$, $y_{it} = y_t(\omega_{it})$, $w_{it} = w_t(\omega_{it}^2)$, $r_{it} = r_t(\omega_{it}^2)$, $q_{it} = q_t(\omega_{it}^2)$, $p_{it} = p_t(\omega_{it}^2)$, $p_{it}^* = p_t^*(\omega_{it}^2)$, $c_{it} = c_t(\omega_{it}^2)$, and $c_{it}^* = c_t^*(\omega_{it}^2)$, where, recall, $\omega_{it}^2 = (\omega_{it}, \omega_{jt})$, with $j = m(i, t)$ denoting i 's trading partner in period t . That is, an island's employment and output depend only on its own beginning-of-period information set, while the commodity prices it ends up facing in stage 2, the realized real wage and rental rates, its consumption bundle, and the local asset prices depend also on the information the island obtains in stage 2 from its trading partner.¹¹ Notwithstanding these qualifications, we define a competitive equilibrium in an otherwise conventional manner.

Definition 1. *An equilibrium is a collection of state-contingent allocation and price functions, $(n, y, w, r, q, p, p^*, c, c^*)$, such that (i) the associated allocations are optimal for households and firms, and (ii) the associated prices clear all markets, for all islands, all dates, and all histories.*

Remark. The information structure considered above should not be taken too literally. What we ultimately seek to capture is the beliefs that people form about economic outcomes such as labor-market conditions, consumer demand, asset returns, and so on. The exact details of how these beliefs are formed in reality are beyond our complete understanding, just as are the complex ways through which people communicate through markets, the media, the blogosphere, private relationships, and so on. The exogenous signals $\tilde{\omega}_{it}$ and the sentiment shock ξ_t are ultimately modeling devices that help us introduce exogenous variation in this kind of beliefs.

Also note that we have ruled out trading in financial assets, or other risk-sharing arrangements. This is only for expositional simplicity. Introducing such trading *within* islands would permit us to price certain assets (e.g., the local land), but would not affect equilibrium allocations, simply because agents are identical within each island. Allowing such trading *across* islands may affect allocations by permitting islands to diversify their idiosyncratic trading risks. Yet, as we discuss in Section 6, this issue is orthogonal to our contribution: the self-fulfilling phenomena we document survive even if markets are complete in the sense that, whenever two islands meet, they are free to trade any good, asset, or insurance contract they might possibly wish. The only essential friction is that communication remains imperfect so that any two islands cannot reach common knowledge about their private histories and courses of action *prior* to their meeting.

4 Equilibrium and self-fulfilling fluctuations

In this section we first characterize the general equilibrium of our economy, we next establish our key result regarding the possibility of self-fulfilling fluctuations, and we finally present a series of examples that illustrate the rich fluctuations that can be accommodated by our theory.

¹¹Note that that the wage and the rental rate are allowed to be contingent on information that will be revealed after trade has taken place. This convention is inconsequential for our results. Also, 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 $\omega, \omega' \in \mathcal{S}_t$. This simply means that any two islands that trade face, of course, the same the terms of trade.

4.1 Characterization

Consider the behavior of the household of island i during stage 2 of period t . Let $\lambda_{it} = \lambda_t(\omega_{it}^2)$ denote the Lagrange multiplier on its budget. Its optimal consumption choices satisfy

$$U_c(c_{it}, c_{it}^*) = \lambda_{it} p_{it} \quad \text{and} \quad U_{c^*}(c_{it}, c_{it}^*) = \lambda_{it} p_{it}^*. \quad (3)$$

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), imposing market clearing, and using the Cobb-Douglas specification of U , we obtain the following:

$$c_{it} = (1 - \eta)y_{it}, \quad c_{it}^* = \eta y_{jt}, \quad \text{and} \quad \lambda_{it} p_{it} = U_c(c_{it}, c_{it}^*) = y_{it}^{-\eta} y_{jt}^{\eta} \quad (4)$$

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 local price of the ‘‘home’’ good increases with ‘‘foreign’’ supply.

Price effects like the above are central to our theory, for they encapsulate the interdependence of economic decisions that trade introduces in *any* market economy. Indeed, this interdependence is present even within the Arrow-Debreu framework, simply because the market prices faced by any particular agent are ultimately pinned down by the behavior of other agents.

Turning back to the characterization of the equilibrium, consider the optimal labor-supply and labor-demand decisions of, respectively, the local household and the local firm during stage 1. These are given by, respectively,

$$V'(n_{it}) = \mathbb{E}_{it}[\lambda_{it} w_{it}] \quad \text{and} \quad \mathbb{E}_{it}[\lambda_{it} w_{it}] = \mathbb{E}_{it}[\lambda_{it} p_{it}] \vartheta \frac{y_{it}}{n_{it}}, \quad (5)$$

where $\mathbb{E}_{it}[\cdot]$ is a short-cut for the rational expectation conditional on ω_{it} . These conditions, too, have a familiar interpretation: workers equate the real wage with the marginal disutility of their effort, firms equate the real wage with the marginal revenue product of labor. The only difference from the standard paradigm is that firms and workers alike face uncertainty about the commodity prices (equivalently, their island's terms of trade) that will be determined in stage 2.

Combining conditions (4) and (5), we reach the following condition, which equates the marginal cost of labor in each island with the local expectation of its marginal revenue product:

$$V'(n_{it}) = \vartheta \mathbb{E}_{it}[y_{jt}^{\eta} y_{it}^{1-\eta} n_{it}^{-1}]. \quad (6)$$

Along with the production function (1) and the fact that land is in fixed supply, this condition pins down the equilibrium levels of employment and output in each island as functions of the local productivity, the local land, and the local expectations of the level of output in the island's trading partner. We thus reach the following proposition, which reduces the general equilibrium to a sequence of tractable fixed-point problems, one for each period.

Proposition 1. For each t , let \mathcal{Y}_t be the set of real, positive-valued functions with domain \mathcal{S}_t , and define the operator $\mathcal{T}_t : \mathcal{Y}_t \rightarrow \mathcal{Y}_t$ as follows: for any $f \in \mathcal{Y}_t$ and any $\omega \in \mathcal{S}_t$,

$$\mathcal{T}_t f(\omega) = (1 - \hat{\alpha}) \left\{ \frac{\log A_t(\omega) + \hat{\vartheta} \log K}{1 - \hat{\vartheta}} \right\} + \hat{\alpha} \left\{ H^{-1} \left(\sum_{\omega' \in \mathcal{S}_t} H(f(\omega')) \mathcal{P}_t(\omega'|\omega) \right) \right\} \quad (7)$$

where $\hat{\vartheta} \equiv \frac{\vartheta}{\epsilon} \in (0, 1)$, $\hat{\alpha} \equiv \frac{\eta}{\eta + (1 - \hat{\vartheta})/\hat{\vartheta}} \in (0, 1)$, $H(x) \equiv \eta \exp(x)$, $A_t(\omega)$ identifies the productivity of an island of type $\omega \in \mathcal{S}_t$, and $\mathcal{P}_t(\omega'|\omega)$ is the probability that this island meets an island of type $\omega' \in \mathcal{S}_t$.¹² 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 .

To interpret the above fixed-point relation, it helps to re-write it as follows:

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

where $z_i \equiv \frac{1}{1 - \hat{\vartheta}} (\log A_i + \hat{\vartheta} \log K)$ summarizes the local fundamentals and \mathbf{E}_{it} is a risk-adjusted expectation operator, defined by $\mathbf{E}_{it}[x] \equiv H^{-1}(\mathbb{E}_{it}[H(x)])$. The equilibrium level of output of an island is thus given as a weighted average of the local fundamentals and the local expectations of the equilibrium level of output in the island's trading partner. In this regard, the competitive equilibrium of our economy can be re-cast as the perfect Bayesian equilibrium of a game in which the players are the islands and their best responses are the fixed-point relation in (8): it is *as if* each island seeks to align its own output with the output of its trading partner.

This game-theoretic interpretation reveals a certain resemblance between our economy and the more abstract incomplete-information games studied, *inter alia*, by Morris and Shin (2002) and Angeletos and Pavan (2007). We revisit this point in Section 5. For now, we note that competitive trade is the sole origin of what looks like strategic interaction in our economy.¹³ This observation explains why the strength of the interdependence of economic outcomes across islands, as measured by the coefficient $\hat{\alpha}$, is increasing in the extent of specialization and trade, as parameterized by η . More generally, Proposition 1 encapsulates the interdependence of economic choices that is endemic to *any* multi-good, multi-agent economy, whether trading takes place in competitive or non-competitive fashion.¹⁴

¹²The 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 $\omega \in \mathcal{S}_t$, returns the first element of ω , while \mathcal{P}_t follows from the exogenous stochastic structure of the economy.

¹³Indeed, our model is a Walrasian economy, not a game; the actual agents (firms, workers, consumers) are infinitesimal price-takers, not strategic players; and the interdependence across islands reflects merely the endogeneity of their competitive terms of trade, not any kind of externality in preferences, technology and the like.

¹⁴For instance, a similar interdependence is present in the New-Keynesian framework in terms of pricing decisions: the target price of each firm depends on the aggregate price level (the prices set by other other firms). Alternatively, a similar result is proved in Angeletos and La'O (2009), in terms of quantities this time, for an RBC economy that features centralized markets for the consumption goods, a Dixit-Stiglitz preference structure, monopoly power, and flexible nominal prices. Finally, La'O (2010) obtains a similar interdependence from collateral constraints. A safe conjecture is thus that our insights can be embedded in richer models of the macroeconomy, which might help bring our theory closer to the data. In this paper we opt to stay away from such extensions only for pedagogical reasons: we seek to show that imperfect communication *alone* opens the door to self-fulfilling fluctuations and "animal spirits" even at the very heart of the neoclassical paradigm.

Putting aside these points, Proposition 1 helps identify the set of equilibria of our economy with the set of fixed points of the operators \mathcal{T}_t . Using Blackwell’s sufficiency conditions, it is then easy to check that, for all t , \mathcal{T}_t is a contraction mapping with modulus equal to $\hat{\alpha} \in (0, 1)$. It follows that \mathcal{T}_t admits a unique fixed point, which in turn proves the following.

Proposition 2. *The equilibrium exists and is unique.*

4.2 Imperfect communication and self-fulfilling fluctuations

To fix language, we say that the economy exhibits “self-fulfilling fluctuations” if and only if aggregate uncertainty in equilibrium allocations and prices obtains without aggregate shocks either to the true fundamentals—namely, preferences, technologies, and matching—or to any agent’s beliefs of these fundamentals. This definition is akin to the conventional definition of sunspot fluctuations (Cass and Shell, 1983), except that here we are concerned with an economy in which the equilibrium is unique and there is no room for randomization devices (sunspots, lotteries, etc.).

We next define (im)perfect communication as follows.

Definition 2. *The economy exhibits “perfect communication” if any two islands that have been matched together in a given period are able to reach common knowledge about each other’s production levels and/or the terms of their upcoming trade as early as in stage 1 of that period.*

This clarifies the notion of imperfect communication that is at the heart of our theory: communication regards people’s ability to share the same beliefs about equilibrium allocations and prices, which means in effect reaching common knowledge about one another’s courses of action.

With these definitions at hand, we state our core positive result as follows.

Theorem 1. *The economy can feature self-fulfilling fluctuations along its unique equilibrium if and only if communication is imperfect.*

The “only if” part of this theorem epitomizes the current state of the art: perfect communication rules out self-fulfilling fluctuations in our setting as in any other conventional, unique-equilibrium, macroeconomic model. In particular, perfect communication guarantees that allocations and prices within any particular match of islands are pinned down by local preferences and technologies, which in turn implies that aggregate outcomes are pinned down by the cross-sectional profile of these fundamentals and the matching technology. By itself, this rules out self-fulfilling fluctuations. Since we have also ruled out aggregate shocks to either this profile or the matching technology, it follows that our economy features no aggregate volatility whatsoever as long as communication is perfect.

The “if” part of the theorem contains the key theoretical contribution of our paper: imperfect communication opens the door to self-fulfilling fluctuations. As anticipated in the Introduction, the basic intuition for this result is quite simple. Take any period t and any pair of islands (i, j) that have been matched together in this particular period. Once these two islands meet and trade in stage 2 of that period, they a fortiori reach symmetric information. Nonetheless, in stage 1, these islands may have private information, either because they have received different signals from Nature, or

because they have met with, and “talk to”, different islands in the past. As a result of this, the beliefs that these islands form, in equilibrium, regarding the (endogenous) terms of their upcoming trade may well not be spanned by their beliefs about the (exogenous) fundamentals. In short, extrinsic uncertainty emerges despite the uniqueness of the equilibrium. Aggregate fluctuations then emerge as the symptom of aggregate (correlated) movements in beliefs of economic outcomes, without any aggregate shocks either to the fundamentals themselves or to peoples’ beliefs of the fundamentals.

From a game-theoretic perspective, these correlated belief movements must ultimately be rationalized by correlated shocks to higher-order beliefs that obtain without correlated shocks either to fundamentals or to first-order beliefs. Yet, this interpretation is not strictly needed. In a competitive, large-scale economy like ours, people do not need to engage in higher-order reasoning; they only need to form rational expectations about the statistic model that governs equilibrium prices and aggregate economic outcomes. The self-fulfilling fluctuations we document are thus best understood as form of extrinsic uncertainty in aggregate economic outcomes.

In the remainder of this section, we use a series of examples to illustrate the richness and the empirical plausibility of the fluctuations that can be accommodated by our theory. We also elaborate on how the aforementioned correlation in beliefs could be, at least in part, the product of imperfect, decentralized communication.

4.3 A tractable Gaussian example

Consider the following specification of our model. K is normalized to 1. The cross-sectional distribution of productivities is log-normal: $\log A_i \sim \mathcal{N}(0, \sigma_A^2)$, for some $\sigma_A > 0$. For any two islands i and j that are matched together in period t , the exogenous signal that i receives about j is given by the pair $\tilde{\omega}_{it} = (x_{it}, s_{it})$, where

$$x_{it} = \log A_j + \varepsilon_{it} \quad \text{and} \quad s_{it} = x_{jt} + \xi_t + u_{it},$$

and where $\varepsilon_{it} \sim \mathcal{N}(0, \sigma_x^2)$ and $u_{it} \sim \mathcal{N}(0, \sigma_s^2)$ are idiosyncratic noises, with $\sigma_x, \sigma_s > 0$. That is, each island receives some information about the productivity of its trading partner, as well as about the latter’s information. Finally, the sentiment shock is i.i.d. over time and Normally distributed: $\xi_t \sim \mathcal{N}(0, \sigma_\xi^2)$, for some $\sigma_\xi > 0$.¹⁵

Let $\log Y_t$ and $\log N_t$ denote aggregate output and employment, defined as the cross-sectional averages of local output and employment, $\log y_{it}$ and $\log n_{it}$. Next, let $\bar{E} \log Y_t$ and $\bar{E} \log N_t$ denote the corresponding average beliefs, defined as the cross-sectional average of the corresponding local beliefs. Finally, note that the beliefs island i holds about the productivity of its trading partner are pinned down by the signal x_{it} alone, which is itself invariant in ξ_t . It follows that variation in ξ_t does not cause variation in any agent’s beliefs about the fundamentals of its trading partner, nor of any other agent in the economy. Yet, as the next proposition verifies, innovations in ξ_t trigger self-fulfilling fluctuations.

¹⁵This assumption implies that the fluctuations we document below are uncorrelated over time; we will add persistence in the following subsection.

Proposition 3. Consider the unique equilibrium of the Gaussian economy described above.

(i) There exist vectors $\Phi_0 \in \mathbb{R}^2$ and $\Phi_\xi \in \mathbb{R}_+^2$ such that, for all dates and states of nature,

$$\begin{bmatrix} \log Y_t \\ \log N_t \end{bmatrix} = \Phi_0 + \Phi_\xi \xi_t \quad \text{and} \quad \begin{bmatrix} \bar{E} \log Y_t \\ \bar{E} \log N_t \end{bmatrix} = \Phi_0 + \frac{\rho_\xi}{\rho_\varepsilon + \rho_\xi} \Phi_\xi \xi_t$$

(ii) There exist coefficients $\phi_0, \psi_0 \in \mathbb{R}$ and $\phi_a, \phi_x, \phi_s, \psi_a, \psi_x, \psi_s \in \mathbb{R}_+$ such that, for all islands, dates, and states of nature,

$$\begin{aligned} \log y_{it} &= \phi_0 + \phi_a \log A_i + \phi_x x_{it} + \phi_s s_{it} \\ \mathbb{E}_{it} \log p_{it} &= \psi_0 - \psi_a \log A_i + \psi_x x_{it} + \psi_s s_{it} \end{aligned}$$

Part (i) establishes that variation in the extrinsic variable ξ_t triggers positive co-movement in aggregate economic activity and in people’s beliefs about it. As a result, booms and recessions appear to be “fueled” by forces akin to “animal spirits”, that is, by movements in market beliefs that are utterly disconnected from either beliefs about the fundamentals or the fundamentals themselves. Yet, these “animal spirits” reflect neither multiple equilibria nor deviations from rationality.

Part (ii) reveals the micro-level behavior that rests beneath these aggregate self-fulfilling fluctuations. Naturally, i ’s output increases with its own productivity. In turn, this explains why j ’s output increases with x_{jt} : a higher x_{jt} informs j that i ’s productivity and output will be higher, which is “good news” about j ’s terms of trade. But as long as j behaves in this way, a higher s_{it} signals that j is likely to produce more, which is now “good news” about i ’s terms of trade. It follows that i ’s output increases with s_{jt} , even though the latter does not affect i ’s beliefs about either its own or its partner’s fundamentals.

To see now how part (ii) explains part (i), consider a positive innovation in ξ_t . If this innovation were common knowledge, no island would react to it. But as this raises s_{it} for all i , each island forms optimistic beliefs about its terms of trade.¹⁶ For local firms, this means an increase in their expectations of the relative price of their product, which in turn motivates them to raise their production and their demand for all inputs (labor and land). For local workers, this translates to an increase in their expected real wage, which motivates them to work harder. For local land owners, it means an increase in the rental rate, which can boost local asset prices.¹⁷ As a result, the economy ends up experiencing a self-fulfilling boom—with the optimism of one agent (or one island) justifying that of another. Conversely, a negative innovation in ξ_t causes a self-fulfilling “loss of confidence” in the sense of triggering pessimistic beliefs about aggregate economic activity, which in turn manifest themselves in a recession.

¹⁶Note that the response of y_{it} to ξ_t is symmetric across all i . As a result, the *realized* terms of trade do not move with ξ_t . Yet, variation in ξ_t causes variation in *expected* terms of trade as of stage 1, because ξ_t is not common knowledge in that stage. Variant examples where ξ_t moves also some realized terms of trade can also be constructed.

¹⁷This in turn suggests that in variants of our model that introduce capital, the boom would manifest not only in employment, output, consumption, but also in investment.

4.4 Persistence

The preceding example featured no persistence: aggregate fluctuations were uncorrelated over time. This, however, was only for expositional simplicity: persistent self-fulfilling fluctuations obtain as long as the sentiment shock is persistent and does not become common knowledge.

To illustrate this, we let the sentiment shock be driven by the following process:

$$\Delta\xi_t = \mu\Delta\xi_{t-1} + v_t,$$

where Δ denotes the first-difference operator, $\mu \in (0, 1)$ parameterizes the persistence of the sentiment process, and $v_t \sim \mathcal{N}(0, \sigma_\xi^2)$ is the innovation term. Next, to maintain the analysis tractable, we assume that ξ_t becomes common knowledge after $T > 1$ periods and let $\Xi_t \equiv \{\xi_t, \xi_{t-1}, \dots, \xi_{t-T}\}$ denote the history of the sentiment shock over the last T periods. It follows that the communication that takes place at the end of each period is akin to the observation of a series of noisy signals on the vector Ξ_t (see Appendix). We can thus show that the equilibrium level of aggregate output is given by a log-linear function of Ξ_t :

$$\log Y_t = \Phi' \Xi_t \tag{9}$$

where Φ is a vector in \mathbb{R}^{T+1} (and Φ' denotes its transpose).

Figure 1 illustrates this result for the case of a positive innovation in the sentiment shock.¹⁸ In the left panel of this figure, the solid line represents the impulse response of aggregate output, while the other two lines reveal the underlying belief dynamics, as summarized in the cross-sectional average of the belief that each island holds either about aggregate output (dotted line) or about the output of its trading partner (dashed line). Evidently, the boom is fueled by a “wave of optimism” that obtains without any change in preferences, technologies, or other fundamentals. This wave keeps building force for a while, before eventually fading away. As a result, aggregate output, and along with it employment, consumption, and asset prices, follows a hump shape—there is a boom, and then a bust.

Hump-shaped dynamics like the above are commonly obtained in Structural VAR exercises. Accommodating them in standard macroeconomic models typically requires the introduction of adjustment costs, habit formation, and other sources of inertia in either preferences or technologies. Here, these hump-shaped dynamics originate merely in self-fulfilling beliefs.

The mechanics of these dynamics rest on the interaction of two competing effects. On the one hand, following the initial innovation, ξ_t keeps increasing for a while; this is due to the particular stochastic process we assumed for ξ_t . On the other hand, islands keep receiving more and more information about ξ_t over time; this information is embedded in the s signals each island receives either from Nature or from meeting with other islands. The first effect helps generate the initial build up: for a while, the continuing increase in ξ_t causes all islands to get increasingly optimistic about their terms of trade, adding fuel to the boom. The second effect helps generate the eventual bust: as time pass, the impact of the initial shock to the islands’ beliefs about their terms of trade

¹⁸The solution to the above problem is explained in the Appendix. Here, we concentrate on a numerical example.




Figure 1: Persistent fluctuations. The left panel illustrates the impulse response of aggregate output to a positive sentiment shock (solid line), along with the corresponding impulse response of the cross-sectional average of the beliefs that each island holds either about aggregate output (dotted line) or about the output of its trading partner (dashed line). The right panel illustrates the comparative statics of the impulse response of aggregate output with respect to the speed of learning and the exogenous growth in the sentiment shock.

fades away, putting a stop to the self-fulfilling boom. In short, an exogenous form of “increasing exuberance” explains the initial build-up, while endogenous learning explains the eventual fading.

These two effects are further illustrated in the right panel of Figure 1. From the dotted green line, we see that the initial build-up disappears once $\rho = 0$, which removes the escalating dynamics in the exogenous sentiment process. From the dashed red line, on the other hand, it is clear that the persistence of the response of output increases when we raise σ_u , the level of noise in the s signals, simply because this slows down learning.

Translating these effects to practical terms, we invite the reader not to seek precise interpretations of either the exogenous sentiment process nor the noise variables featured in the preceding example. These are merely convenient modeling devices that help us capture the rich waves of “optimism” and “pessimism” that may be sustained by imperfect communication.

4.5 Contagion and propagation

To reinforce the last message, we now discuss how decentralized communication may serve as a powerful propagation mechanism, helping self-fulfilling beliefs spread from one agent to another in a manner akin to the spread of rumors, fads, and contagious diseases.

To build intuition, suppose the economy consists of two agents, Amy and Bob, who produce differentiated goods and expect that, with positive probability, they will have the chance to trade with each other. As in our earlier example with the bakers and the shoemakers, imperfect communication between Amy and Bob can sustain self-fulfilling variation in their beliefs about the terms of their trade. Say that both Amy and Bob have formed jointly “exuberant” beliefs about the terms of their trade—they expect each other to produce a lot and, because of these optimistic beliefs, they are themselves willing to produce more.

Now let a third agent, Cathy, enter the economy. Suppose Cathy is initially uninformed, expects to trade with Bob, but first meets Amy. Clearly, Amy has nothing to learn from this meeting. Cathy, however, inherits Amy’s optimistic beliefs about Bob. In so doing, she turns “exuberant” herself: she comes to believe that her terms of trade with Bob will be favorable and hence finds it in her best interest to raise her own production level. Finally, consider a fourth agent, David. Suppose David is also uninformed initially but gets to meet Bob and, in so doing, inherits Bob’s beliefs. David will then turn exuberant whether he expects to trade with Amy or Cathy. And so on.

We can accommodate these intuitions in our framework as follows. In the beginning of time, let the islands be split into two equally-sized groups. Productivity is the same within a group but differs across groups. Let $a_1 \equiv \log A_1$ denote the log of the productivity of group 1, and $a_2 \equiv \log A_2$ that of group 2. These are drawn independently from one another from a Normal distribution with mean zero and variance σ_a^2 .¹⁹ 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: a signal about the other groups’ productivity, and a signal about the other group’s signal about their own productivity. Similarly as in Section 4.3, these signals are given by $x_1 = a_2 + \varepsilon_1$ and $s_1 = x_2 + \xi$ for group 1, and $x_2 = a_1 + \varepsilon_2$ and $s_2 = x_1 + \xi$ for group 2, where $\varepsilon_1, \varepsilon_2$ and ξ are independent of one another as well as of the productivity draws a_1 and a_2 , with ε_1 and ε_2 drawn from Normal distributions with mean zero and variance σ_ε^2 and ξ drawn from a Normal with mean zero and variance σ_ξ^2 . The initial fraction of partially informed islands is an exogenous parameter denoted by χ . We let $\chi \in (0, 1/2)$, so that the majority of islands are initially uninformed.

Let $z = (a_1, a_2; \varepsilon_1, \varepsilon_2, \xi)$ denote the exogenous aggregate state, which is determined by Nature in the beginning of time and henceforth stays constant. Once nature draws z in the beginning of time, no further aggregate shock ever hits the economy—the only uncertainty that is realized over time is the idiosyncratic one associated with random matching. Furthermore, no island receives any further exogenous information about the underlying state. Together, these properties guarantee that all the dynamics we will document below are the sole product of the communication that takes place in the economy as different islands meet, trade, and “talk” to one another.

The matching technology is independent of z , and it 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 its 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*

¹⁹Because there are only two groups, the law of large number does not apply, leaving the economy with aggregate productivity risk. However, this can be removed by considering variants of this example that increases the number of groups with such independent productivity draws. In any event, our abstraction from aggregate fundamental risk in the preceding analysis was only for pedagogical reasons. Furthermore, the presence of such risk may complement our results, for coordination on first-best outcomes may now require perfect communication, not only about idiosyncratic fundamentals, but also about aggregate fundamentals.

productivity group, in which case they both learn the entire aggregate state z and hence turn into what we shall henceforth call “fully informed” islands next period. Finally, a fully informed island can only meet with a fully informed from its own productivity group; the match then leads to no communication, for these islands already know the entire state.

This structure defines an “information ladder”, with the uninformed islands at the bottom, the partially informed in the middle, and the fully informed at the top. The matching technology is then such that, in any given period, an island can either learn nothing from its match and hence maintain its initial position in the ladder, or can learn just enough to move exactly one step up the ladder. Eventually, all islands reach the top of the ladder (formally, fully informed is an absorbing state), but this takes time. The aggregate dynamics we document below are a manifestation of how the population of islands ascends this informational ladder.

To preserve a closed-form solution, we assume that, in stage 1 of each period, each island knows beforehand whether in stage 2 it will meet an island that is equally or differentially informed.²⁰ It is then straightforward to check that the only islands whose employment and production choices are sensitive to ξ are the partially informed islands that are matched with partially informed islands (from the other productivity group): it is only within these matches that the s signals help an island predict its trading partner’s output and, equivalently, its terms of trade. To fix language, we henceforth refer to these islands as the “exuberant” islands (with the understanding that their “exuberance” should be re-interpreted as “pessimism” in the case of a negative realization of the initial sentiment shock ξ).

These islands behave in essentially the same way as in the example of Section 4.3 (or as Amy and Bob in our preceding informal discussion). But, whereas in this earlier example all the islands were exuberant, here only a fraction are exuberant at any point of time. Furthermore, and importantly, this fraction evolves over time, as the product of the endogenous communication that takes place in the economy through matching and trading.

Proposition 4. *Let λ_t be the fraction of exuberant islands in period t .*

(i) *The economy experiences a “fad”: λ_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*

$$\frac{\partial \log Y_t}{\partial \xi} = \Phi \lambda_t \quad \forall t$$

The first part of this result underscores how decentralized communication can generate phenomena akin to the spread of fads, rumors, and contagious diseases. The second part then makes clear that the entire dynamic effects of the initial sentiment shock on aggregate economic activity originate in communication.

These properties are illustrated in Figure 3. The left panel of this figure documents the dynamic response of aggregate output, and of the average forecast of aggregate output, to the initial positive

²⁰The results can be extended to the alternative scenario, but then we would have to resort to numerical solution.

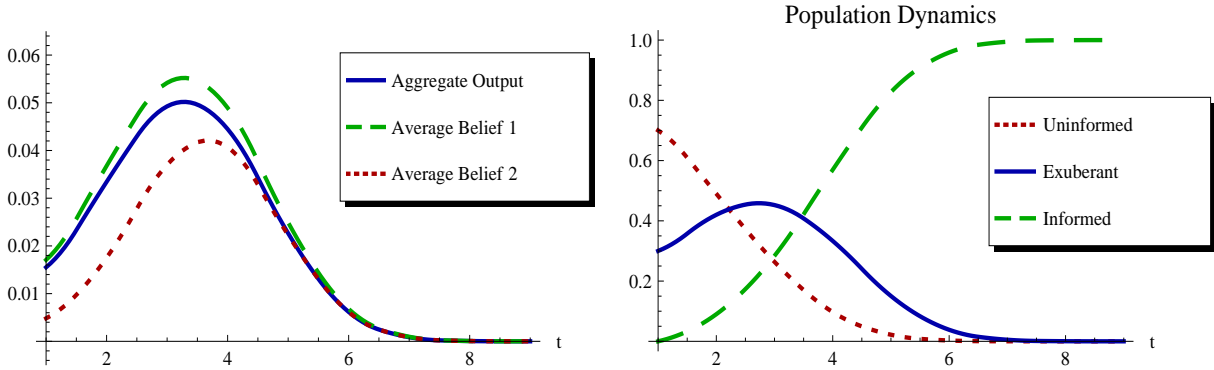


Figure 2: Contagious exuberance. The left panel illustrates the dynamics of aggregate output following the initial realization of ξ (solid line), along with the corresponding dynamics of the cross-sectional average of the beliefs that each island holds either about aggregate output (dotted line) or about the output of its trading partner (dashed line). The right panel illustrates the underlying population dynamics.

sentiment shock. The right panel documents the underlying population dynamics, that is, the evolution of the distribution of islands along the aforementioned information ladder. It is then evident that the dynamics of aggregate output, as well as those of the average forecast, track the dynamics of the fraction of “exuberant” islands. As anticipated in the proposition, this fraction is first increasing and then decreasing. The resulting boom thus takes the form of a “wave of optimism” that builds up force for a while, only to fade away after enough time.

This wave is akin to the hump-shaped dynamics we obtained in Section 4.4, except that now both the initial build-up and the eventual fading emerge endogenously as the product of communication. The initial increase occurs as optimistic beliefs spread from one island to another (i.e., as uninformed islands meet partially informed islands and become “exuberant” themselves); the eventual decrease happens as more and more islands reach common knowledge about the underlying state and therefore about their production choices and their terms of trade (i.e., as partially informed islands meet with other partially informed islands and become fully informed).

These findings are reminiscent of the contagion effects 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 “fad” dynamics are similar to those found in Burnside, Eichenbaum, and Rebelo (2011): both papers feature economies in which beliefs are transmitted from some agents to others in way such that the proportion of “exuberant” agents initially rises, reaches a peak, and then eventually falls. But whereas these authors model this kind of contagion 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, not just despite of it.

To conclude, the preceding example, although quite abstract, serves three important functions. First, it underscores how communication helps propagate “animal spirits” from one agent to another

like a contagious disease—communication *means* contagion. Second, it illustrates how the correlation in beliefs that was previously entirely hard-wired into the exogenous stochastic structure of the economy can be recast, at least in part, as the by-product of communication—communication *means* correlation in beliefs. Finally, it helps explain why our theory may naturally accommodate the hump-shaped dynamics that are found in structural VAR exercises, as well as the “fad” dynamics that many associate with the recent “bubbles” in asset and housing markets.

What we have to leave outside our theory is the initial trigger: as with *any* other theory of stochastic fluctuations, our theory also needs some exogenous random impulses. But whereas the dominant approach requires *all* the exogenous impulses to obtain in fundamentals, our theory permits these impulses to be disconnected from fundamentals—shocks emerge in self-fulfilling beliefs, not just in fundamentals.

4.6 Magnitude and co-movement

A quantitative evaluation of our theory is beyond the scope of this paper. Nonetheless, looking forward to future research, we would like to give some indications regarding the likely quantitative potential of the type of fluctuations we have formalized in this paper.

Towards this goal, consider the Gaussian example studied in Section 4.3 and normalize the variances of the sentiment shock and the idiosyncratic noises to be proportional to the heterogeneity in productivities: $(\sigma_A, \sigma_\xi, \sigma_x, \sigma_s) = (1, \gamma_\xi, \gamma_x, \gamma_s) \cdot \sigma$, for some $\gamma_\xi, \gamma_x, \gamma_s, \sigma > 0$. We can thus parameterize the economy by the scalar $\sigma \in \mathbb{R}_+$ and the vector $\mathbf{e} \equiv (\beta, \vartheta, \eta, \epsilon, \gamma_\xi, \gamma_x, \gamma_s) \in \mathcal{E} \equiv (0, 1)^3 \times \mathbb{R}_+^4$. As both the intrinsic and the extrinsic components of this uncertainty are proportional to σ , the magnitude of both the aggregate and the idiosyncratic variation in equilibrium output (or employment) are also proportional to σ , which proves the following.

Proposition 5. *Consider the Gaussian example of Section 4.3. There exists a function $\Lambda : \mathcal{E} \rightarrow \mathbb{R}_+$ such the equilibrium volatility of aggregate output is given by*

$$\text{Var}(\log Y_t) = \Lambda(\mathbf{e})\sigma^2.$$

It follows that sufficient uncertainty at the micro level (sufficiently high σ) sustains arbitrarily high volatility at the macro level.

This proposition exemplifies the quantitative potential of our theory: provided the search and communication frictions faced at the micro-level are important enough (in the sense that σ is big enough), our theory can match *any* level of macroeconomic volatility, despite the uniqueness of the equilibrium and the entire absence of aggregate shocks to fundamentals.

Clearly, this is in sharp contrast to the standard paradigm. In any unique-equilibrium DSGE model, the magnitude of aggregate fluctuations is tightly connected to the magnitude of the underlying aggregate shocks to preferences, technologies, or other fundamentals. As the uncertainty in the latter vanishes, macroeconomic volatility also vanishes. A similar property holds in the pertinent lit-

erature on information frictions.²¹ By contrast, our approach permits us to decouple the magnitude of macroeconomic volatility from the magnitude of aggregate uncertainty in fundamentals.

Turning to the ability of our theory to generate the right *co-movement* in macroeconomic activity, note that employment is necessarily procyclical in our setting, simply because there is no variation either in the technology nor in the capital stock: output fluctuates in our model only because employment fluctuates. Furthermore, because there is no investment, consumption coincides with output. It follows that, once seen under the lenses of the standard neoclassical model, our employment fluctuations will register as countercyclical movements in the “labor wedge”, thereby matching an important aspect of the data (Chari, Kehoe, McGrattan, 2007; Shimer, 2009). At the same time, because the rental rate of land is procyclical, one may guess that richer variants of our framework that recast land as capital could also feature procyclical investment. Finally, to the extent that we introduce variable capital utilization and/or increasing returns, our fluctuations can also feature procyclical labor productivity and procyclical Solow residuals.²²

While quite intuitive, it is worth emphasizing that these properties of our theory are in contrast to the recent literature on “news shocks”, which was spurred by Beaudry and Portier (2006). This literature has focused on signals of future productivity, which have opposing effects on labor supply and consumption in the neoclassical framework. This literature has thus sought to obtain positive co-movement between employment and consumption by introducing exotic preferences (Jaimovich and Rebelo, 2009) or sticky prices and suboptimal monetary policy (Lorenzoni, 2010). By contrast, we abstract from productivity shocks, or news thereof, and consider an entirely different type of shock—a shock that triggers variation in labor demand as firms get more or less optimistic about the demand for their products. More precisely, variation in ξ_t causes short-run variation in each island’s expected marginal revenue product, the incentive effects of which are akin to those of a transitory productivity shock. This explains why our fluctuations feature the right co-movement between employment and consumption.

To conclude, although a serious quantitative exploration has to await future research, our theory appears to have no obvious difficulty in generating either a significant level of volatility or the right cyclical co-movement in macroeconomic activity.

5 Discussion: communication, coordination, and beliefs

Our theory helps formalize a very basic idea, one that used to be at the core of Keynesian thinking: the way economic behavior is modeled in the centralized Arrow-Debreu framework, and in modern macroeconomic models, presumes a level of coordination that is patently unrealistic.

²¹This literature accommodates two sources of volatility: innovations in aggregate fundamentals; and noise in signals about fundamentals (as in the complementary literature on “news shocks”). Nonetheless, as the aggregate uncertainty in fundamentals vanishes, *both* types of volatility also vanish.

²²A tractable variant of our model that confirms these conjectures is available upon request; this variant introduces an “investment sector” that uses labor to produce a capital good that depreciates with in one period. Whether the desired cyclical properties will obtain in more serious quantitative exercises remains open for future work.

A voluminous literature tried to capture this idea in the 80's and 90's in models with multiple equilibria.²³ Among the early seminal contributions to this literature, Diamond (1982) considered a model with search frictions and thick-market externalities. Random matching was used to capture the decentralization of trading and the possibility that some trading opportunities could go unexplored (which could then be interpreted as unemployment). Externalities in the matching technology (in the probability of meeting a valuable trading partner) were then used to generate multiple equilibria, some of which featured a lot of trade (low unemployment), while others featured little trade (high unemployment). Kiyotaki and Wright (1993) and Lagos and Wright (2005), on the other hand, use similar notions of random matching and decentralized trading in order to capture the role of money as a medium of exchange.

Building on this tradition, our paper also used random matching to capture the decentralization of market interactions. We nevertheless abstracted entirely from the issues that have dominated the pertinent literature on search, such as unemployment and money. We also ruled out any non-convexity in the matching technology—in fact, the probability of matching was exogenous in our model—or any other source of multiple equilibria. Instead, we shifted the focus to an entirely different aspect of trading frictions: the fact that they impede communication. Our key contribution was then to show, in effect, that imperfect communication *means* imperfect coordination.

In so doing, we also revisit Hayek's (1946) seminal insight regarding the role of markets in coordinating economic activity. Recall that Hayek argued, not only that markets can play an important role in communicating information, but also that they can do so more efficiently than planning systems because it seems improbable that a "center" could ever attain the extent of communication and information aggregation that is needed for implementing first-best outcomes. But then note that the Arrow-Debreu framework replaces the notion of a central planner with the notion of a centralized Walrasian auctioneer, which hardly seems to capture the essence of Hayek's argument: that the information/communication requirements of the Arrow-Debreu framework are identical to those of centralized planning mechanisms. By contrast, our approach lets decentralization have a bite on the extent of communication that can take place under *either* the market's solution concept we have considered so far, or the planner's solution concept we introduce and study in Section 6.

Turning to an another, complementary interpretation of our results, we remind the reader that the equilibrium of our economy can be understood as the unique rationalizable outcome of a certain game among the islands of our economy. From this perspective, our self-fulfilling fluctuations can be mapped to random variation in higher-order beliefs (the beliefs of one island about the beliefs of its likely trading partner) that is orthogonal to either first-order beliefs (the beliefs of the fundamentals) or to the fundamentals themselves. Furthermore, while the variation in fundamentals and first-order beliefs is restricted to be uncorrelated in the cross-section of matches, the aforementioned higher-order beliefs are correlated. It is this correlated higher-order uncertainty that explains our

²³See, inter alia, Azariadis (1981), Benhabib and Farmer (1994, 1999), Cass and Shell (1983), Cooper and John (1988), Diamond (1982), Diamond and Fudenberg (1989), Guesnerie and Woodford (1992), Howitt and McAfee (1992), Matsuyama (1991), Shell (1977), and Woodford (1991).

self-fulfilling fluctuations from a game-theoretic perspective.²⁴

In this respect, our contribution complements, and builds upon, Morris and Shin (2002, 2003), Woodford (2003), and a growing literature that studies the macroeconomic effects of informational frictions and higher-order uncertainty. However, as anticipated in the Introduction, there is a substantial innovation. This literature presumes aggregate shocks to fundamentals, removes common knowledge about these shocks, and studies the response of the economy to noisy signals about these shocks. In so doing, this literature enriches our understanding of propagation mechanisms and helps decompose the observed volatility between shocks to fundamentals and noisy news about fundamentals. But it does not revisit the ultimate origins of fluctuations: all fluctuations continue to hinge on uncertainty in aggregate fundamentals.²⁵

By contrast, our contribution dispenses entirely with this kind of uncertainty, shifts the focus to the trading and communication frictions agents face at the *micro* level, and shows how these frictions by themselves open the door to rich self-fulfilling fluctuations at the *macro* level. That being said, it should be clear that our results do not hinge on the absence of aggregate shocks to fundamentals.²⁶ The reason we have abstracted from such shocks in this paper is entirely pedagogical: we want to make crystal clear that what we are after is fluctuations that are driven by a certain form of “coordination failure”, and by self-fulfilling beliefs regarding endogenous economic outcomes, not by news about exogenous fundamentals.

Furthermore, whereas the aforementioned literature often shies away from the role of markets in communicating information, this communication is central to our analysis. Self-fulfilling beliefs emerge in our setting *only* because agents expect to trade and only because communication is imperfect. But trade *means* communication. Communication and self-fulfilling beliefs are thus tightly connected in our framework. This point was reinforced in Section 4.5, where communication was shown to propagate self-fulfilling beliefs from one agent to another.

Translating all these formalisms to the real world, we would like to emphasize that, at least of our purposes, signals, information sets, and higher-order beliefs are only modeling devices that help formalize how people form beliefs about the prospects of the economy. We certainly do not envision people engaging everyday in higher-order reasoning regarding one another’s beliefs and choices. Nor do we imagine “Nature” sending people exogenous signals regarding other people’s abilities,

²⁴To see this more clearly, consider our earlier Gaussian example. It is straightforward to check that the signal s_{it} moves island i ’s second and higher-order beliefs regarding its trading partner, without moving its first-order beliefs. The extrinsic shock ξ_t then induces correlation in this kind of higher-order beliefs across all the islands.

²⁵For example, Lucas (1972), Woodford (2003), Mankiw and Reis (2002), Mackowiak and Wiederholt (2009), Nimark (2008), Hellwig and Venkateswaran (2009), and Amador and Weill (2010) alike study how informational frictions impact the response of the economy to monetary shocks; Angeletos and La’O (2009), Lorenzoni (2010), Blanchard et al (2010), and Hassan and Mertens (2010) shift the focus to productivity shocks; and finally Morris and Shin (2002, 2003), Angeletos and Pavan (2007), Rondina and Walker (2010), and Amador and Weill (2011) study payoff shocks in a more abstract class of environments. One way or another, however, the ultimate origins of fluctuations in all this prior work are innovations in aggregate fundamentals (or news about them).

²⁶In fact, adding such shocks would only reinforce our results in so far information about these shocks remains dispersed, for this would introduce an additional source of lack common knowledge about equilibrium allocations and prices.

tastes, and information sets, as we have assumed in our model. In reality, firms’ employment and investment decisions are driven by their expectations of consumer demand; consumers’ spending is driven by their expectations of future employment opportunities and future income; investors’ trading positions are driven by beliefs regarding future asset returns. It is this kind of beliefs that determine actual macroeconomic outcomes and asset prices. It is movements in this kind of beliefs that market pundits and practitioners often refer to as shifts in “market psychology”, “consumer confidence”, and “investor sentiment”. And it is this kind of beliefs that are at the heart of our theoretical exploration in this paper—not the particular modeling devices we use in order to formalize these beliefs.

Finally, our insights are likely to extend well beyond the narrow boundaries of the particular model we have considered in this paper, for they rest only on how decentralization and imperfect communication impede coordination. Nonetheless, the particular micro-foundations we have favored in this paper serve two important goals. First, they highlight that our insights are relevant *even* within the neoclassical core of the modern macroeconomic paradigm. And second, they facilitate a transparent welfare analysis—an issue to which we now turn our attention.

6 Risk-sharing and Efficiency

Our results hinge on islands facing idiosyncratic risk in their terms of trade, in so far this risk opens the door to lack of common knowledge about equilibrium allocations and prices. In general, such idiosyncratic trading risk gives rise to a demand for risk-sharing arrangements. Such risk-sharing arrangements, however, were ruled out in our baseline model. In this section, we extend the analysis to a variant of our model that insulates our results from this type of considerations. By doing so, this variant also provides us with a useful benchmark for studying the welfare properties of the self-fulfilling phenomena we identify in this paper.

6.1 Risk-sharing

We now consider a variant that introduces a homogenous “numeraire” good, which enters preferences linearly and can be traded alongside the local specialized goods. In a manner akin to Lagos and Wright (2005), these features remove any value, whether private or social, for financial trades and for any other form of state-contingent transfer schemes.

Each island is endowed with a fixed amount \tilde{y} of the numeraire good in each period. The production of the specialized goods remains as before. Preferences are given by

$$\sum_t \beta^t [U(c_{it}, c_{it}^*) - V(n_{it}) + \chi \tilde{c}_{it}]$$

where \tilde{c}_{it} is the consumption of the “numeraire” good, χ is a positive scalar, U is strictly increasing

and concave, and V is strictly increasing and convex.²⁷ Finally, balanced trade requires

$$p_{it}^* c_{it}^* = p_{it}(y_{it} - c_{it}) + (\tilde{y} - \tilde{c}_{it}).$$

In words, the “imports” of the foreign specialized good can now be financed by “exports” of either the domestic specialized good or the numeraire.²⁸

The characterization of the equilibrium follows similar steps as in the baseline mode;. The prices and the consumption levels of the specialized goods are pinned down by the resource constraints and the following optimality condition:

$$U_c(c_{it}, c_{it}^*) = p_{it} = p_{jt}^* = U_{c^*}(c_{jt}, c_{jt}^*)$$

Equilibrium prices can thus be expressed as functions of the islands’ outputs:

$$p_{it} = P(y_{it}, y_{jt}),$$

for some function $P : \mathbb{R}_+^2 \rightarrow \mathbb{R}_+$. This is akin to condition (4) in the baseline model. The only difference is that P may no more admit a closed-form solution, but this is inessential.²⁹ Modulo this adjustment, the rest of the equilibrium characterization remains as in the baseline model. We thus reach the following result, which generalizes Propositions 1 and 2 from our baseline model.

Proposition 6. *The equilibrium exists and is unique. Furthermore, there exists a monotone function $G : \mathbb{R}_+^3 \rightarrow \mathbb{R}$ such that, for any $t \geq 1$, the period- t equilibrium output function solves the following fixed-point problem:*

$$\sum_{\omega' \in \mathcal{S}_t} G(y_t(\omega), A_t(\omega), y_t(\omega')) \mathcal{P}_t(\omega' | \omega) = 0 \quad (10)$$

Condition (10) is the analogue of condition (7) in the baseline model: it gives the equilibrium output of an island as an increasing function of the local productivity and the local beliefs about the likely output level of other islands. As in the baseline model, the interdependence among the islands originates merely from specialization and trade. Before, this dependence embodied certain wealth effects in the sense that terms of trade depended, not only on the aggregate supply of the two traded goods, but also on their ownership pattern. Now, all wealth effects have been absorbed by the consumption of the numeraire good. This, however, does not affect the core of our positive results: self-fulfilling fluctuations continue to emerge as long as communication is imperfect.

This is most clearly illustrated in the following special case. Let $U(c, c^*) = (c^{1/2} c^{*1/2})^\gamma$ and $V(n) = n^\epsilon$, where $0 < \gamma < 1 < \epsilon$. It is then easy to check that $P(y, y^*) = y^{-\gamma/2} y^{*\gamma/2}$ and, by

²⁷For added flexibility, we no more impose the functional forms for U and V that we used in our baseline model.

²⁸Note that prices are normalized so that the price of the numeraire good is one.

²⁹A simple expression for P obtains when preferences are symmetric between the “home” and the “foreign” good in the sense that $U(c, c^*) = U(c^*, c)$ for all c, c^* . In this case, $P(y, y^*) \equiv U_c(y/2, y^*/2)$. More generally, one can show that $P(y, y^*)$ is necessarily decreasing in y , while it is increasing (resp., decreasing) in y^* if and only if $U_{c, c^*} > 0$ (resp., < 0). The existence of self-fulfilling fluctuations then hinges on ruling out the knife-edge case where $U_{c, c^*} = 0$.

implication, condition (10) reduces to condition (7), modulo a redefinition of the scalars $\hat{\vartheta}$ and $\hat{\alpha}$. It follows that all our preceding results can immediately be recast within this special case.

Moving beyond the particular variant we have studied here, what sustains our self-fulfilling fluctuations is only that communication is imperfect. From Grossman (1981), we know that the combination of centralized and complete markets induces, in effect, perfect communication. But as long as communication remains imperfect, the details of the span of goods and assets that the agents may be able to trade are likely to be inessential for our results.

6.2 Efficiency

Understanding the welfare properties of the equilibrium, and thereby the possible desirability of policy intervention, requires the definition of an appropriate efficiency concept. One possibility is to compare our equilibrium concept with first-best efficiency. In this regard, our equilibrium is clearly inefficient. Note, however, that implementing first-best outcomes would require perfect communication between the entire economy and a “center” (the planner), which is completely at odds with the entire spirit of our exercise—decentralization, and its bite on communication, are at the essence of our approach.

We thus consider a different efficiency concept, one that helps isolate the welfare losses that may obtain from misalignment of private and social incentives from the ones that are the inevitable by-product of decentralization and imperfect communication.³⁰ In so doing, we also isolate the welfare role of institutions that facilitate more communication from that of tax, regulatory, or other policies that seek to manipulate market incentives or otherwise directly impact the allocation of resources in the economy.

Definition 3. *A constrained efficient allocation is a resource-feasible allocation that maximizes ex-ante utility subject to the constraint that information is transferred across any two islands when, and only when, the two islands meet.*

Let $\mathcal{P}_t(\omega, \omega')$ denote the probability that an island of type $\omega \in \mathcal{S}_t$ is matched with an island of type $\omega' \in \mathcal{S}_t$ during period t and note that, once these islands meet, the stage-2 information set of the former becomes $\omega^2 = (\omega, \omega')$, while that of the latter becomes $\omega'^2 = (\omega', \omega)$. We can thus write the planner’s problem whose solution identifies the constrained efficient allocation as follows.

Planning Problem. *The constrained efficient allocation maximizes*

$$\mathcal{W} = \sum_t \beta^t \left\{ \sum_{\omega, \omega' \in \mathcal{S}_t} [U(c_t(\omega, \omega'), c_t^*(\omega, \omega')) - V(n_t(\omega)) + b\tilde{c}_t(\omega, \omega')] \mathcal{P}_t(\omega, \omega') \right\}$$

subject to

$$c_t(\omega, \omega') + c_t^*(\omega', \omega) = y_t(\omega) = A_t(\omega)n_t(\omega)^\vartheta \quad \forall \omega, \omega' \in \mathcal{S}_t, \forall t \quad (11)$$

$$\tilde{c}_t(\omega, \omega') + \tilde{c}_t(\omega', \omega) = 2\tilde{y} \quad \forall \omega, \omega' \in \mathcal{S}_t, \forall t \quad (12)$$

³⁰Angeletos and Pavan (2007, 2009) and Vives (2008) study related concepts within certain classes of games.

Because of the linearity of preferences, the allocation of the numeraire good is indeterminate. The remaining problem, however, is strictly convex and hence has a unique solution, which is pinned down by FOCs. Furthermore, this problem is separable over time, so we can characterize the efficient allocation of one period independently of that of another period.

Let $\lambda(\omega, \omega')\mathcal{P}(\omega, \omega')$ be the Lagrange multiplier on the resource constraint (11). The optimality conditions for consumption give

$$U_c(c_t(\omega, \omega'), c_t^*(\omega, \omega')) = \lambda_t(\omega, \omega') = U_{c^*}(c_t(\omega', \omega), c_t^*(\omega', \omega)),$$

which together with the local resource constraints pins down the efficient levels of consumption for given levels of output. Comparing this result with the corresponding one for the equilibrium, we see immediately that, for any given output levels, the efficient and the equilibrium consumption allocations are the same, and the planner's shadow prices coincide with market prices. Turning to the efficient employment and output levels, these are determined by the following:

$$V'(n_t(\omega)) = \sum_{\omega' \in \mathcal{S}} \lambda_t(\omega, \omega') \mathcal{P}_t(\omega, \omega') \vartheta \frac{y_t(\omega)}{n_t(\omega)}$$

This is the same condition as the corresponding one that characterizes the equilibrium, except that the market prices (p) have now been replaced by the planner's shadow prices (λ). But we already argued that shadow and market prices coincide. The following is thus immediate.

Theorem 2. *The equilibrium is constrained efficient.*

This result contains the core normative lesson of our paper: the ostensibly pathological fluctuations we have documented in this paper are anything but a free call for government intervention. Indeed, the only way that society can improve upon the decentralized equilibrium is by facilitating more communication. It follows that, unless it incidentally does this, no fiscal, monetary, or regulatory intervention could improve welfare. Interestingly, this is true in our setting even though it is common knowledge that the *entire* business cycle is driven by “animal spirits”.

Clearly, this result stands in sharp contrast to conventional wisdom and to previous formalizations of “animal spirits”: in models with multiple equilibria or irrational agents, the need for government intervention is hard-wired. It also offers a new twist to a heated public debate on the state of our science and on the role of the government.

Following the recent crisis, some economists have criticized the dominant macroeconomic paradigm of misguiding policy-making by presuming that the observed fluctuations represent the rational, and coordinated, response of the economy to exogenous disturbances in preferences and technology. Most provocatively, Krugman (2009) declares that this notion is “silly” and that, instead, “Keynesian economics remains the best framework we have for making sense of recessions and depressions.” Similar positions are voiced by Akerlof and Shiller (2008), Shiller (2009), and Solow (2010).

Lucas (2009), Levine (2009), Cochrane (2009), and Chari (2010) offer pointed responses, reminding us of the fallacies of the old Keynesian paradigm and defending the rules of the game in

modern macroeconomic research. Yet, in certain respects, these responses seem to assume away an important part of the aforementioned criticisms: the dominant methodological framework leaves no room for the kind of forces envisioned by the aforementioned economists, but this does not mean that these forces are not relevant in reality. Our approach, instead, offers a more powerful response to these criticisms by giving a central position to the Keynesian notions of “coordination failure” and “animal spirits” that are so dear to the aforementioned critiques, while at the same time only reasserting the normative lessons of the Neoclassical paradigm.

As with the first welfare theorem, our normative result is likely to break down once one allows for realistic frictions in product, labor, or financial markets. The interaction of these frictions with our notion of self-fulfilling beliefs may thus provide novel insights into both the nature of business cycles and the role of stabilization policy. Yet, the key message of our result will survive—the way in which the “invisible hand” operates in economies with decentralized markets and coordination frictions has hereby been shown to be far more mysterious than what was previously thought.

7 Concluding remarks

When agents meet and trade in markets or otherwise interact with one another, they exchange information, not only regarding the underlying economic fundamentals (such as tastes, abilities, and technologies), but also about their own and others’ likely courses of action (such as firms’ likely employment and investment choices, consumers’ likely spending choices, or investors’ likely portfolio choices). In the Arrow-Debreu framework, and in the vast majority of modern macroeconomic models alike, this type of communication is assumed to be instantaneous and flawless: economic agents are assumed to share common knowledge, not only of the underlying fundamentals, but also of their courses of action (equivalently, of equilibrium allocations and prices).

Our contribution in this paper is to show that, once this convenient but unrealistic assumption is relaxed, our understanding of the workings of the market mechanism take a surprising twist: self-fulfilling phenomena, and seemingly exotic forces such as “coordination failures”, “animal spirits”, and “contagious exuberance” are hereby shown to be endemic to the constraints that decentralization imposes on communication.

For apparent pedagogical reasons, we formalized these insights within the narrow boundaries of a class of competitive, unique-equilibrium, rational-expectations economies. This permitted us to accommodate the aforementioned ideas at the heart of the neoclassical paradigm. It also provided us with an important normative benchmark. It should be clear, however, that our insights hinge only on imperfect communication and can thus be relevant for a wide class of applications in macroeconomics and finance.

With these results we thus hope to push the research frontier in a new direction—one that pays closer attention to the role that self-fulfilling beliefs, and their transmission from one agent to another, appear to play in actual business cycles and in phenomena such as the recent dot-com and housing bubbles (and their subsequent busts). Whether the particular formalization of belief-driven

fluctuations we have proposed in this paper is convincing or appealing is for the reader to decide. One way or another, however, we believe it is time to recognize the distinct, and potentially powerful, role of this kind of forces—and at the very least it is time to experiment beyond the conventional practice of recasting the residuals of our understanding as the preference and technology shocks, or the mysterious wedges, of otherwise elaborate DSGE models.

Appendix

Proof of Proposition 1. Substituting the optimality condition for labor (6) into the production function yields

$$y_{it} = A_i \left(\vartheta \mathbb{E}_{it} [y_{jt}^\eta y_{it}^{1-\eta}] \right)^{\frac{\vartheta_n}{\epsilon}}$$

Rearranging, we get

$$y_{it}^{1-(1-\eta)\left(\frac{\vartheta_n}{\epsilon}\right)} = A_i \vartheta^{\frac{\vartheta_n}{\epsilon}} \left(\mathbb{E}_{it} [y_{jt}^\eta] \right)^{\frac{\vartheta_n}{\epsilon}}.$$

Taking logs, rearranging, and using the definitions of $\hat{\vartheta}$ and $\hat{\alpha}$, we reach condition (8), or equivalently condition (7). *QED*

Proof of Proposition 2. As mentioned in the main text, existence and uniqueness follows from the fact that that operator \mathcal{T} is contraction. To verify the latter fact, we now show that \mathcal{T} satisfies Blackwell's sufficiency conditions.

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

$$\mathcal{T}_t f(\omega) - \mathcal{T}_t g(\omega) = \hat{\alpha} \left\{ H^{-1} \left(\sum_{\omega' \in \mathcal{S}_t} H(f(\omega')) \mathcal{P}_t(\omega'|\omega) \right) - H^{-1} \left(\sum_{\omega' \in \mathcal{S}_t} H(g(\omega')) \mathcal{P}_t(\omega'|\omega) \right) \right\}$$

Note that $\hat{\alpha} > 0$ 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

$$\sum_{\omega' \in \mathcal{S}_t} \eta \exp(f(\omega')) \mathcal{P}_t(\omega'|\omega) \geq \sum_{\omega' \in \mathcal{S}_t} \eta \exp(g(\omega')) \mathcal{P}_t(\omega'|\omega). \quad (13)$$

Now, note that $f(\omega) \geq g(\omega)$ for all $\omega \in \mathcal{S}$ implies that $\eta \exp(f(\omega')) \geq \eta \exp(g(\omega'))$ for all $\omega \in \mathcal{S}_t$. This immediately implies that condition (13) 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:

$$\begin{aligned} \mathcal{T}_t [f(\omega) + a] &= (1 - \hat{\alpha}) \left\{ \frac{1}{1-\hat{\vartheta}} \log A_t(\omega) \right\} + \hat{\alpha} \left\{ H^{-1} \left(\sum_{\omega' \in \mathcal{S}_t} H(f(\omega') + a) \mathcal{P}_t(\omega'|\omega) \right) \right\} \\ &= (1 - \hat{\alpha}) \left\{ \frac{1}{1-\hat{\vartheta}} \log A_t(\omega) \right\} + \hat{\alpha} \left\{ H^{-1} \left(\sum_{\omega' \in \mathcal{S}_t} H(f(\omega')) \mathcal{P}_t(\omega'|\omega) \right) \right\} + \hat{\alpha} a \end{aligned}$$

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

As both the monotonicity and the discounting conditions of Blackwell's theorem are satisfied, we conclude that the operator \mathcal{T} is indeed a contraction. *QED*

Proof of Theorem 1. Here we prove the “only if” part of the theorem, namely that perfect communication implies that our economy cannot feature self-fulfilling fluctuations along its unique equilibrium. The converse follows from the Gaussian example of Section 4.3.³¹

Pick an arbitrary match (i, j) in period t , consider stage 1 of this period, and suppose that the islands in this match have reached common knowledge about their output levels (either directly or indirectly by first reaching common knowledge about their terms of trade and then using equilibrium reasoning to translate the latter to their output levels). In this case, the equilibrium fixed-point relation (8) reduces to the following:

$$\begin{aligned}\log y_{it} &= (1 - \hat{\alpha}) \left(\frac{1}{1-\hat{\vartheta}} \log A_i + \frac{\hat{\vartheta}}{1-\hat{\vartheta}} \log K \right) + \hat{\alpha} \log y_{jt} \\ \log y_{jt} &= (1 - \hat{\alpha}) \left(\frac{1}{1-\hat{\vartheta}} \log A_j + \frac{\hat{\vartheta}}{1-\hat{\vartheta}} \log K \right) + \hat{\alpha} \log y_{it}\end{aligned}$$

This has two implications. First, it implies that common knowledge of the two island’s output levels, along with equilibrium reasoning, induces common knowledge of their productivities. And second, it guarantees that equilibrium outcomes are pinned down by productivities. Indeed, the solution to this system is given by $\log y_{it} = g(A_i, A_j)$ and $\log y_{jt} = g(A_j, A_i)$, where

$$g(A, A') \equiv \frac{1}{1-\hat{\vartheta}} \left[\frac{1}{1+\hat{\alpha}} \log A + \frac{\hat{\alpha}}{1+\hat{\alpha}} \log A' \right] + \frac{\hat{\vartheta}}{1-\hat{\vartheta}} \log K$$

That is, the equilibrium output of each island is merely a function of its own productivity, of the productivity of its trading partner, and of the common size of land. It is then immediate that the equilibrium aggregate variables are pinned down by the cross-sectional distribution of productivities, and therefore cannot feature self-fulfilling fluctuations.

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 . It follows that the non-linear expectation $\mathbf{E}_{it}y_{jt}$ and the simple expectation $\mathbb{E}_{it}y_{jt}$ are equal to each other up to a constant that we henceforth ignore for expositional simplicity. We can thus rewrite the key equilibrium condition as

$$\log y(\omega_i) = (1 - \hat{\alpha}) \frac{1}{1-\hat{\vartheta}} a_i + \hat{\alpha} \mathbb{E}_{it} [\log y(\omega_{jt})] \quad (14)$$

where $a_i \equiv \log A_i$.

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 $\log y_t(\omega_{jt}) = \phi_a a_j + \phi_x x_{jt} + \phi_s s_{jt}$, for some coefficients (ϕ_a, ϕ_x, ϕ_s) . It follows that $\log y_t(\omega_{jt})$ is indeed log-normal, with

$$\mathbb{E} [\log y_t(\omega_{jt}) | \omega_{it}] = \phi_a \mathbb{E} [a_j | \omega_{it}] + \phi_x (a_i + \mathbb{E} [\varepsilon_{jt} | \omega_{it}]) + \phi_s (x_{it} + \mathbb{E} [\xi_t | \omega_{it}]) \quad (15)$$

³¹This is true except for one detail: that example introduces an infinite state space, while our model assumed a finite state space. An example with a finite state space can easily be constructed and is omitted here only to economize on space.

where

$$\begin{bmatrix} \mathbb{E}[a_j|\omega_{it}] \\ \mathbb{E}[\varepsilon_{jt}|\omega_{it}] \\ \mathbb{E}[\xi_t|\omega_{P1+}] \end{bmatrix} = \begin{bmatrix} \frac{1}{\rho_\varepsilon+1}x_{it} \\ \frac{\rho_\varepsilon}{\rho_\varepsilon+\rho_\xi}(s_{it}-a_i) \\ \frac{\rho_\xi}{\rho_\varepsilon+\rho_\xi}(s_{it}-a_i) \end{bmatrix}$$

Substituting these expressions into (21) gives us

$$\log y(\omega_{it}) = (1-\hat{\alpha})\frac{1}{1-\hat{\vartheta}}a_1 + \hat{\alpha} \left[\phi_a \frac{1}{\rho_\varepsilon+1}x_{it} + \phi_x \left(a_i + \frac{\rho_\varepsilon}{\rho_\varepsilon+\rho_\xi}(s_{it}-a_i) \right) + \phi_s \left(x_{it} + \frac{\rho_\xi}{\rho_\varepsilon+\rho_\xi}(s_{it}-a_i) \right) \right]$$

By symmetry, equilibrium output for type ω_{it} must satisfy $\log y(\omega_{it}) = \phi_a a_i + \phi_x x_{it} + \phi_s s_{it}$. For this to coincide with the above condition for every z , it is necessary and sufficient that the coefficients (ϕ_a, ϕ_x, ϕ_s) solve the following system:

$$\begin{aligned} \phi_a &= (1-\hat{\alpha})\frac{1}{1-\hat{\vartheta}} + \hat{\alpha}\phi_x - \phi_s \\ \phi_x &= \hat{\alpha} \left(\phi_a \frac{1}{\rho_\varepsilon+1} + \phi_s \right) \\ \phi_s &= \hat{\alpha} \left(\phi_x \frac{\rho_\varepsilon}{\rho_\varepsilon+\rho_\xi} + \phi_s \frac{\rho_\xi}{\rho_\varepsilon+\rho_\xi} \right) \end{aligned}$$

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

$$\phi_a = \frac{(1-\hat{\alpha})(1+\rho_\varepsilon^2)\left((1+\hat{\alpha})\rho_\varepsilon^2+\rho_\xi^2\right)}{(1-\hat{\vartheta})\left((1+\hat{\alpha})\rho_\varepsilon^2(1+\rho_\varepsilon^2)+(1-\hat{\alpha}^2+\rho_\varepsilon^2)\rho_\xi^2\right)} \quad (16)$$

$$\phi_x = \frac{\hat{\alpha}\rho_\varepsilon^2+(1-\hat{\alpha})\hat{\alpha}\rho_\xi^2}{(1-\hat{\vartheta})\left((1+\hat{\alpha})\rho_\varepsilon^2(1+\rho_\varepsilon^2)+(1-\hat{\alpha}^2+\rho_\varepsilon^2)\rho_\xi^2\right)} \quad (17)$$

$$\phi_s = \frac{\hat{\alpha}^2\rho_\varepsilon^2}{(1-\hat{\vartheta})\left((1+\hat{\alpha})\rho_\varepsilon^2(1+\rho_\varepsilon^2)+(1-\hat{\alpha}^2+\rho_\varepsilon^2)\rho_\xi^2\right)} \quad (18)$$

Given the log-linear structure of equilibrium output, and the log-normal specification for the shock and information structure, we find that aggregate output is given by

$$\log Y_t = \phi_0^Y + \phi_\xi^Y \xi_t \quad (19)$$

where $\phi_0^Y \equiv \frac{1}{2} \left[(\phi_a + \phi_x + \phi_s)^2 + (\phi_x + \phi_s)^2 \rho_\varepsilon \right]$ and $\phi_\xi^Y = \phi_s$. Also, note that from the production function, equilibrium labor is given by

$$\log n(\omega_{it}) = \frac{1}{\vartheta} (\log y(\omega_{it}) - a_i) = \frac{1}{\vartheta} ((\phi_a - 1)a_i + \phi_x x_{it} + \phi_s s_{it})$$

Thus aggregate labor takes the following form

$$\log N_t = \phi_0^N + \phi_\xi^N \xi_t$$

where $\phi_0^N \equiv \frac{1}{2} \left(\frac{1}{\vartheta} \right)^2 \left[(\phi_a - 1 + \phi_x + \phi_s)^2 + (\phi_x + \phi_s)^2 \rho_\varepsilon \right] \sigma^2$ and $\phi_\xi^N = \phi_s$. These results give us the aggregate dynamics stated in the proposition, with $\Phi_0 = (\phi_0^Y, \phi_0^N)$ and $\Phi_\xi = (\phi_\xi^Y, \phi_\xi^N)$. Finally,

note that $\mathbb{E}_{it}[\xi_t] = \frac{\rho_\xi}{\rho_\varepsilon + \rho_\xi} (s_{it} - a_i)$. It follows that the average belief of ξ_t equals $\frac{\rho_\xi}{\rho_\varepsilon + \rho_\xi} \xi_t$, which in turn gives us the average beliefs of aggregate output and employment as stated in the proposition. *QED*

Proof of Condition (9). First, note that

$$\Xi_t \equiv \begin{bmatrix} \xi_t \\ \xi_{t-1} \\ \dots \\ \xi_{t-T} \end{bmatrix}$$

is the only component of the underlying aggregate state that has not yet become common knowledge as of period t . Next, consider the information that island i receives during stage 2 of period t . At this point, the island has observed, not only its own current-period signals, (x_{it}, s_{it}) , but also those of its trading partner, (x_{jt}, s_{jt}) . The errors in these signals, namely the noises $u_{it}, u_{jt}, \epsilon_{it}$, and ϵ_{jt} are uncorrelated with one another as well as with the history of the sentiment shock and the history of all the signals that the island has received in the past. It follows that the combination of (x_{it}, s_{it}) and (x_{jt}, s_{jt}) contains the same information about the history of the sentiment shock (and thereby about future terms of trade) as the observation of the following two signals: $\tilde{s}_{it}^1 \equiv s_{it} - x_{jt} = \xi_t + u_{it}$ and $\tilde{s}_{it}^2 \equiv s_{jt} - x_{jt} = \xi_t + u_{jt}$. Finally, since u_{it} and u_{jt} are i.i.d. Normal noises with variance σ_u^2 , the aforementioned two signals are informationally equivalent to a single signal of the form $\tilde{s}_{it} = \xi_t + \tilde{u}_{it}$, where \tilde{u}_{it} is Normal noise with variance equal to $\sigma_u^2/4$. Communication thus involves exchanging the histories of such signals each island has received through past trades.

As the end of period t , this means that the island i has observed the true ξ_{t-T} along with the following series of signals about Ξ_t : 1 signal of the form $\tilde{s}_{it} = \xi_t + \tilde{u}_{it}$; 2 signals of the form $\tilde{s}_{it-1} = \xi_{t-1} + \tilde{u}_{it-1}$; ... ; and $T - 1$ signals of the form $\tilde{s}_{i,t-(T-1)} = \xi_{t-(T-1)} + \tilde{u}_{i,t-(T-1)}$. This is equivalent to observing a vector signal of the form

$$Z_{it} = \Xi_{it} + \nu_{it}$$

where ν_{it} is Normal, independent of ξ_s for all s , i.i.d. across time and islands, with mean 0 and variance-covariance matrix given by

$$\Sigma_\nu = \begin{bmatrix} \frac{\sigma_u^2}{4} & 0 & \dots & 0 & 0 \\ 0 & \frac{\sigma_u^2}{16} & \dots & 0 & 0 \\ \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & \frac{\sigma_u^2}{4(T-1)^2} & 0 \\ 0 & 0 & \dots & 0 & 0 \end{bmatrix}.$$

Now, let

$$X_{it} \equiv \begin{bmatrix} \log A_i \\ x_{i,t} \\ s_{i,t} \\ Z_{i,t-1} \end{bmatrix}.$$

We can then guess and verify that the equilibrium level of local output is given by

$$\log y_{it} = \Phi' X_{it}.$$

for some vector Φ in \mathbb{R}^{T+1} . To see this, note that, as long as the above conjecture holds, i 's forecast of its trading partner's output is given by

$$\mathbb{E}_{it}[\log y_{jt}] = \Phi' \mathbb{E}_{it}[X_{jt}]$$

Next, due to the Gaussian structure,

$$\mathbb{E}_{it}[X_{jt}] = BX_{it},$$

for some $(T+1) \times (T+1)$ matrix B ; this matrix is itself obtained by standard projection methods. It then follows from (8) that

$$\log y_{it} = (1 - \hat{\alpha}) \frac{1}{1 - \hat{\beta}} \log A_i + \hat{\alpha} \Phi' BX_{it},$$

which together with our initial guess gives the following fixed-point relation for the vector Φ :

$$\Phi = (1 - \hat{\alpha}) \frac{1}{1 - \hat{\beta}} e_1 + \hat{\alpha} \Phi' B$$

where e_1 is a vector in \mathbb{R}^{T+1} with 1 in its first element and zeros in the rest. Solving for Φ gives the equilibrium level of output at the island level. Aggregating across islands gives condition (9). *QED*

Proof of Proposition 4. *Part (i).* For any period and any history up to that point, the type of an island belongs to the following set:

$$\bar{\Omega} \equiv \{\omega_{U1}, \omega_{U1+}, \omega_{P1}, \omega_{P1+}, \omega_{F1}; \omega_{U2}, \omega_{U2+}, \omega_{P2}, \omega_{P2+}, \omega_{F2}\},$$

where, for each group $k \in \{1, 2\}$, ω_{Uk} are uninformed islands that are matched with a uninformed island from their group, ω_{Uk+} are uninformed islands that are matched with a partially informed island, ω_{Pk} are partially informed islands that are matched with an uninformed island; ω_{Pk+} are partially informed islands that are matched with a partially informed island from the *other* group; and ω_{Fk} 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_{F1} and ω_{F2} are absorbing states for, respectively, groups 1 and 2. Along with the fact that $\lambda_0 > 0$, this proves that λ_t must eventually decrease and must converge to zero as $t \rightarrow \infty$. Finally, the fact that λ_t must initially increase follows from the assumption $\chi < 1/2$.

Part (ii). We start by showing that there exist positive coefficients (ϕ_a, ϕ_x, ϕ_s) such that the equilibrium level of output in an island is given by

$$\log y_{it} = \begin{cases} \phi_a a_1 + \phi_x x_1 + \phi_s s_1 & \text{if } \omega_{it} = \omega_{P1+}, \\ \phi_a a_2 + \phi_x x_2 + \phi_s s_2 & \text{if } \omega_{it} = \omega_{P2+}, \\ \phi_a a_i & \text{otherwise} \end{cases} \quad (20)$$

We prove this by guessing and verifying. 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 . It follows that the non-linear expectation $\mathbf{E}_{it}y_{jt}$ and the simple expectation $\mathbb{E}_{it}y_{jt}$ are equal to each other up to a constant that we henceforth ignore for expositional simplicity. We can thus rewrite the key equilibrium condition as

$$\log y(\omega_i) = (1 - \hat{\alpha}) \frac{1}{1 - \hat{\vartheta}} a_i + \hat{\alpha} \mathbb{E}_{it} [\log y(\omega_{jt})] \quad (21)$$

Using this condition, we now consider the equilibrium outputs for each of the ten possible types of islands, by considering the equilibrium outcomes for all possible matches.

First, consider matches between two islands of type ω_{U1} . In this case, equilibrium output of type ω_{U1} must satisfy $\log y(\omega_{U1}) = (1 - \hat{\alpha}) \frac{1}{1 - \hat{\vartheta}} a_1 + \hat{\alpha} \log y(\omega_{U1})$. It follows that $\log y(\omega_{U1}) = \phi_a a_1$ for $\phi_a = \frac{1}{1 - \hat{\vartheta}}$. A similar result holds for matches between two islands of type ω_{U2} .

Next, consider matches between two islands of type ω_{U1+} and ω_{P1} . Suppose the equilibrium production strategies of these islands take a log-linear form, that is $\log y(\omega_{U1+}) = \phi_{0U} a_1$ for some coefficient ϕ_{0U} and $\log y(\omega_{P1}) = \phi_{0P} a_1 + \phi_x x_1 + \phi_s s_1$, for some coefficients $(\phi_{0P}, \phi_x, \phi_s)$. It follows that $y(\omega_{U1+})$ and $y(\omega_{P1})$ are indeed log-normal, with

$$\begin{aligned} \mathbb{E}[y(\omega_{U1+}) | \omega_{P1}] &= \phi_{0U} a_1 \\ \mathbb{E}[y(\omega_{P1}) | \omega_{U1+}] &= \phi_{0P} a_1 + \phi_x \mathbb{E}[x_1 | \omega_{U1+}] + \phi_s \mathbb{E}[s_1 | \omega_{U1+}] \end{aligned}$$

where $\mathbb{E}[x_1 | \omega_{U1+}] = \mathbb{E}[s_1 | \omega_{U1+}] = 0$. Substituting these expressions into (21) gives us

$$\begin{aligned} \log y(\omega_{U1+}) &= (1 - \hat{\alpha}) \frac{1}{1 - \hat{\vartheta}} a_1 + \hat{\alpha} \phi_{0P} a_1 \\ \log y(\omega_{P1}) &= (1 - \hat{\alpha}) \frac{1}{1 - \hat{\vartheta}} a_1 + \hat{\alpha} \phi_{0U} a_1 \end{aligned}$$

It follows immediately that the unique solution to this is $\log y(\omega_{U1+}) = \phi_{0U} a_1$ and $\log y(\omega_{P1}) = \phi_{0P} a_1$ with $\phi_{0U} = \phi_{0P} = \frac{1}{1 - \hat{\vartheta}}$. A similar result holds for matches between two islands of type ω_{U2+} and ω_{P2} .

Next, consider matches between two islands of type ω_{P1+} and ω_{P2+} . This case is identical to the equilibrium between two islands in the static Gaussian example considered in Section (4.3). Thus, we may infer that equilibrium output for types ω_{P1+} and ω_{P2+} must satisfy $\log y(\omega_{P1+}) = \phi_a a_1 + \phi_x x_1 + \phi_s s_1$ and $\log y(\omega_{P2+}) = \phi_a a_2 + \phi_x x_2 + \phi_s s_2$ where the coefficients (ϕ_a, ϕ_x, ϕ_s) are given in (16)-(18).

Finally, consider matches between two islands of type ω_{F1} . In this case, equilibrium output of type ω_{F1} must satisfy $\log y(\omega_{F1}) = (1 - \hat{\alpha}) \frac{1}{1 - \hat{\vartheta}} a_1 + \hat{\alpha} \log y(\omega_{F1})$. It follows that $\log y(\omega_{F1}) = \phi_a a_1$ for $\phi_a = \frac{1}{1 - \hat{\vartheta}}$. A similar result holds for matches between two islands of type ω_{F2} .

This completes the characterization of local outcomes. By aggregating (20), we then obtain the following characterization for aggregate output:

$$\log Y_t = \phi_a \bar{a} + \lambda_t [\phi_x \bar{\varepsilon} + \phi_s \xi]$$

where $\bar{a} \equiv \frac{1}{2}(a_1 + a_2)$ and $\bar{\varepsilon} \equiv \frac{1}{2}(\varepsilon_1 + a_2)$, and where λ_t is the fraction of islands with types either ω_{P1+} or ω_{P2+} . The result then follows immediately by letting $\Phi \equiv \phi_s$. *QED*

Proof of Proposition 5. Note that the variance of aggregate output is simply given by

$$\text{Var}(\log Y_t) = \phi_s^2 \text{Var}(\xi_t),$$

where $\text{Var}(\xi_t) = \gamma_\xi^2 \sigma^2$ is the variance of the sentiments shock. Using the equilibrium value for ϕ_s , we can compute the variance of Y_t as a function of the primitive parameters, $\beta, \epsilon, \eta, \theta, \psi, \gamma_\xi, \gamma_\epsilon$ and σ . This is given by the following expression.

$$\text{Var}(\log Y_t) = \frac{\hat{\alpha}^4 \gamma_\xi^2 \gamma_\epsilon^4}{\left(1 - \hat{\vartheta}\right)^2 \left[\gamma_\xi^2 (1 - \hat{\alpha}^2 + \gamma_\epsilon^2) + (1 + \hat{\alpha}) \gamma_\epsilon^2 (1 + \gamma_\epsilon^2)\right]^2} \sigma^2$$

where $\hat{\alpha}$ and $\hat{\vartheta}$ are given in the statement of Proposition 1. From the above expression, it is immediate that the variance of output is proportional to σ^2 .

Proof of Proposition 6. Because of the quasi-linearity of preferences, the marginal value of wealth is $\lambda_{it} = 1$. The optimality labor demand of the firm yields

$$w_{it} = \mathbb{E}_{it}[p_{it}] \vartheta \frac{y_{it}}{n_{it}},$$

while the optimal labor supply of the household gives

$$w_{it} = V'(n_{it})$$

Combining the above conditions we get

$$V'(n_{it}) n_{it} = \mathbb{E}_{it}[p_{it}] \vartheta y_{it}$$

By combining this condition with the production function (1), we can express y_{it} as a function of A_i and $\mathbb{E}_{it}[p_{it}]$. Replacing the latter with $p_{it} = P(y_{it}, y_{jt})$ gives

$$y_{it}^{\epsilon/\vartheta - 1} = A_i^{\epsilon/\vartheta} \mathbb{E}_{it}[P(y_{it}, y_{jt})],$$

which coincides with condition (10) once we let

$$G(y, A, y') \equiv y^{1 - \epsilon/\vartheta} A^{\epsilon/\vartheta} P(y, y').$$

Finally, the existence and the uniqueness of the equilibrium follow from Theorem 2, which establishes that the equilibrium coincides with the constrained efficient allocation, which in turn exists and is unique thanks to the convexity of the planner's problem. *QED*

Proof of Theorem 2. This follows from the discussion in the main text.

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] Amador, Manuel, and Pierre-Olivier Weill (2010), "Learning from Prices: Public Communication and Welfare," *Journal of Political Economy*, forthcoming.
- [3] Amador, Manuel, and Pierre-Olivier Weill (2011), "Learning from Private and Public Observations of others actions," Stanford/UCLA mimeo.
- [4] Angeletos, George-Marios, and Jennifer La'O (2009), "Noisy Business Cycles," *NBER Macroeconomics Annual 2009*, 24, 319-378.
- [5] Angeletos, George-Marios, and Alessandro Pavan (2007), "Efficient Use of Information and Social Value of Information," *Econometrica* 75(4), 1103-1142.
- [6] Angeletos, George-Marios, and Alessandro Pavan (2009), "Policy with Dispersed Information," *Journal of the European Economic Association* 7, 11-60.
- [7] Azariadis, Costas (1981), "Self-fulfilling Prophecies," *Journal of Economic Theory* 25 380-96.
- [8] Beaudry, Paul, and Franck Portier (2006), "Stock Prices, News, and Economic Fluctuations," *American Economic Review* 96, 1293-1307.
- [9] Benhabib, Jess, and Roger Farmer (1994), "Indeterminacy and Increasing Returns," *Journal of Economic Theory* 63, 19-41.
- [10] Benhabib, Jess, and Roger Farmer (1999), "Indeterminacy and Sunspots in Macroeconomics," in John Taylor and Michael Woodford, eds., *Handbook of Macroeconomics*, Amsterdam: North-Holland.
- [11] Blanchard, Olivier J., Jean-Paul L'Huillier and Guido Lorenzoni (2009), "News, Noise, and Fluctuations: An Empirical Exporation," MIT mimeo.
- [12] Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo (2011), "Understanding Booms and Busts in Housing Markets," Duke/Northwestern mimeo.
- [13] Cass, David, and Karl Shell (1983), "Do Sunspots Matter?" *Journal of Political Economy* 91, 193-227.
- [14] Chari, V. V. (2010), "Testimony before the Committee on Science and Technology, Subcommittee on Investigations and Oversight, U.S. House of Representatives," July 2010. (http://democrats.science.house.gov/Media/file/Commdocs/hearings/2010/Oversight/20july/Chari_Testimony.pdf)
- [15] Chari, V. V., Patrick Kehoe, and Ellen McGrattan (2007), "Business Cycle Accounting," *Econometrica* 75, 781-836.

- [16] Cochrane, John (2009), "How did Paul Krugman get it so Wrong?," (http://faculty.chicagobooth.edu/john.cochrane/research/Papers/krugman_response.htm)
- [17] Cooper, Russell, and Andrew John (1988), "Coordinating Coordination Failures in Keynesian Models," *Quarterly Journal of Economics* 103, 441-463.
- [18] Diamond, Peter (1982), "Aggregate Demand Management in Search Equilibrium," *Journal of Political Economy* 90(5), 881-894.
- [19] Diamond, Peter, and Drew Fudenberg (1989), "Rational Expectations Business Cycles in Search Equilibrium," *Journal of Political Economy* 97(3), 606-619.
- [20] Duffie, Darrell, and Gustavo Manso (2007). "Information Percolation in Large Markets, " *American Economic Review* 97, 203-209.
- [21] Grossman, Sanford J. (1981), "An Introduction to the Theory of Rational Expectations under Asymmetric Information," *Review of Economic Studies* 48(4), 541-559.
- [22] 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: Cambridge Univ. Press.
- [23] Hayek, Friedrich (1945), "The Use of Knowledge in Society," *American Economic Review* 35(4), 519-530.
- [24] Hellwig, Christian, and Venky Venkateswaran (2009), "Setting the Right Prices for the Wrong Reasons," *Journal of Monetary Economics* 56, S57-S77.
- [25] Howitt, Peter, and Bruce McAfee (1992), "Animal Spirits," *American Economic Review* 82, 493-507.
- [26] Jaimovich, Nir, and Sergio Rebelo (2009), "Can News About the Future Drive the Business Cycle?," *American Economic Review*, 99(4), 1097-1118.
- [27] Kiyotaki, Nobuhiro, and Randall Wright (1993), "A Search Theoretic Approach to Monetary Economics," *American Economic Review* 83(1), 63-77.
- [28] Kocherlakota, Narayana (2010), "Modern Macroeconomic Models as Tools for Economic Policy," *The Region*, Federal Reserve Bank of Minneapolis, May 2010, 5-21.
- [29] Lagos, Ricardo, and Randall Wright (2005), "A Unified Framework for Monetary Theory and Policy Analysis," *Journal of Political Economy* 113, 463-484.
- [30] La'O, Jennifer (2010), "Collateral Constraints and Noisy Fluctuations," Chicago Booth mimeo.
- [31] Lorenzoni, Guido (2010), "A Theory of Demand Shocks," *American Economic Review* 99(5), 2050-2084.
- [32] Levine, David K. (2009), "An Open Letter to Paul Krugman," *The Huffington Post*, September 18, 2009. (http://www.huffingtonpost.com/david-k-levine/an-open-letter-to-paul-kr_b_289768.html)

- [33] Lucas, Robert E., Jr. (2009), "In Defence of the Dismal Science," *The Economist*, August 6, 2009. (<http://www.economist.com/node/14165405>)
- [34] Lucas, Robert E., Jr. (1972), "Expectations and the Neutrality of Money," *Journal of Economic Theory* 4, 103-124.
- [35] Lucas, Robert E., Jr. and Prescott, Edward C. (1974), "Equilibrium search and unemployment," *Journal of Economic Theory* 7(2), 188-209.
- [36] Mackowiak, Bartosz, and Mirko Wiederholt (2009), "Optimal Sticky Prices under Rational Inattention," *American Economic Review*, 99, 769-803.
- [37] Mankiw, N. Gregory and Ricardo Reis (2002), "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *Quarterly Journal of Economics* 117, 1295-1328.
- [38] Matsuyama, Kiminori (1991), "Increasing Returns, Industrialization, and Indeterminacy of Equilibrium," *Quarterly Journal of Economics* 106(2):617-650.
- [39] Morris, Stephen, and Hyun Song Shin (2002), "The Social Value of Public Information," *American Economic Review* 92, 1521-1534.
- [40] 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.
- [41] Shell, Karl (1977) "Monnaie et allocation intertemporelle," mimeo prepared for the Roy-Malinvaut seminar at the Centre National de la Recherche Scientifique; published in *Macroeconomic Dynamics*.
- [42] Shiller, Robert (2005), *Irrational Exuberance*, 2nd Ed., Princeton: Princeton University Press.
- [43] Shiller, Robert (2009), "Animal Spirits Depend on Trust: The proposed stimulus isn't big enough to restore confidence," *The Wall Street Journal*, January 27, 2009, A15. (<http://online.wsj.com/article/SB123302080925418107.html>).
- [44] Shimer, Robert (2009), "Convergence in Macroeconomics: The Labor Wedge," *American Economic Journal: Macroeconomics* 1, 280-297.
- [45] Vives, Xavier (2008), *Information and Learning in Markets*, Princeton: Princeton University Press.
- [46] Woodford, Michael (1991), "Self-Fulfilling Expectations and Fluctuations in Aggregate Demand," in N. Gregory Mankiw and David Romer, eds., *New Keynesian Economics*, Cambridge: The MIT Press.
- [47] 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.