Animal Spirits*

George-Marios Angeletos
MIT and NBER

Jennifer La'O
MIT

June 3, 2010
[preliminary and incomplete; please do not circulate]

Abstract

This paper develops a novel theory of fluctuations, one that focuses on their self-fulfilling nature. In this theory, the economy features fluctuations that are orthogonal to the underlying technologies, preferences, and other fundamentals, thus helping capture the informal notions of "animal spirits," "market sentiments", "demand shocks", and the like. Nonetheless, these fluctuations are neither the product of irrationality nor the symptom of equilibrium determinacy. Rather, they rest only on imperfect communication and obtain in an otherwise canonical, unique-equilibrium, rational-expectations economy. What is more, these fluctuations can be consistent with a notion of constrained efficiency that leaves no room for conventional stabilization policies. A new paradigm thus emerges—one whose positive aspects have a genuine Keynesian flavor, and yet on whose policy implications are unsettlingly anti-Keynesian.

*A previous version of this paper was entitled "Sentiments" and contained a different model than the one in the present version, although it shared the same core message. For useful comments we thank seminar participants at MIT, the 2009 SED meeting in Istanbul, the 2009 EEA/ESEM meeting in Barcelona, and the 2010 meeting of the Midwest Economics Association.

1 Introduction

What drives short-run fluctuations in macroeconomic activity and asset prices? And what should the government do about these fluctuations?

The answers provided to these questions by macroeconomic theory are quite different from those often heard in the public arena. Mainstream macroeconomic models attribute short-run fluctuations to high-frequency movements in technologies, preferences, and government policies; the resulting fluctuations are then often, but not always, a symptom of efficiency. In contrast, many practitioners, market pundits, and policy makers, but also some economists, favor an informal argument that goes back to John Maynard Keynes and that attributes short-run fluctuations to more exotic forces such as "animal spirits", "market psychology", "irrational exuberance," and "confidence"; the undesirability of the resulting fluctuations is then taken to be nearly self-evident.

These informal notions find little room in modern macroeconomics. To some, this is our profession's great failure. Commenting on the recent crisis, Robert Shiller (2009, pp. ??) writes:

> *"But lost in the economics textbooks, and all but lost in the thousands of pages of the technical economics literature, is this other message of Keynes regarding why the economy fluctuates as much as it does. Animal spirits offer an explanation for why we get into recessions in the first place—for why the economy fluctuates as it does."*

The same message underlies a recent book by Akerlof and Shiller (2008), with which our paper shares the same title. Similarly, Paul Krugman (2009) argues that it is "foolish" to attribute short-run fluctuations to technology or preference shocks and advocates an all-around return to the old Keynesian view of fluctuations.

But one could also dismiss these informal notions as void of the coherence and discipline that only modern formal theory can afford. For, as Robert Lucas (2001, pp??) has put it most eloquently,

> *"Economic theory is mathematical analysis. Everything else is just pictures and talk."*

One may also be uncomfortable with how conveniently these informal notions pave the way to policy prescriptions that could have otherwise been hard to justify. For example, Shiller argues that the recent US fiscal stimulus would raise "confidence" in the economy. But, as N. Gregory Mankiw (2009) swiftly notes, this argument is based on little more than wishful thinking.

In this paper we contribute to this debate by developing a novel theoretical paradigm for the origins of short-run fluctuations. This paradigm shifts the focus away from technology and preference shocks; it seeks to formalize the informal notions of "animal spirits", "confidence" and the like; and it helps capture the self-fulfilling nature of short-run fluctuations within conventional macroeconomic models. In so doing, it accommodates a view of the business cycle whose positive aspects have a strongly Keynesian flavor. And yet, it does so without abandoning either the elegance or the discipline of modern macroeconomic theory. What is more, it delivers a very different policy message than the one advocated by either old-fashioned or new-born Keynesianism.

In our theory, agents are fully rational; markets are Walrasian; there are no externalities and no non-convexities; the competitive equilibrium is unique; and there is no room for either correlation devices or lotteries. Nevetheless, equilibrium allocations and prices respond to a particular form of "sentiment shocks". These are shocks that do not move either the underlying technologies, preferences and other fundamentals, nor the agents' beliefs about these fundamentals—and yet they move their expectations of economic activity and of market prices.

The resulting fluctuations can thus be interpreted as random self-fulfilling waves of "optimism" and "pessimism". Indeed, these waves appear to be arbitrary and utterly disconnected from fundamentals whether seen from the perspective of an outsider observer or from that of the agents inside our economy. What is more, the response of the economy to these shocks can be hump-shaped, so that these waves may keep building up force for a while before eventually fading off.

Conventional wisdom, and the pertinent macroeconomics literature, would take these arbitrary shifts in "market sentiment" as prima facie evidence of either some form of irrationality or the existence of multiple equilibria—but this conclusion is unwarranted in our model. Rather, what sustains these fluctuations in our model is a distinct form of uncertainty—one that is absent in standard macroeconomic models, and one that requires only that communication is imperfect.

In any large-scale economy, production and exchange takes place in a decentralized manner. As a result, agents face uncertainty about the economic activity of one another and hence about the terms of trade they may face in future market interactions. In practice, investors are uncertain about firm profits; firms are uncertain about consumer demand; consumers are uncertain about labor-market conditions; and one kind of uncertainty often appears to feed the other.

This underscores what, at least in our view, is the self-fulfilling nature of *actual* fluctuations—a nature that standard macro models fail to capture in an effective way. In these models, short-run fluctuations are driven merely by news regarding technologies, preferences, and government policies—the firms' expectations of consumer demand, the consumers' expectations of labor-market conditions, and the investors' expectations of asset prices, are all pinned down by their beliefs about the underlying economic fundamentals, leaving no room for self-fulfilling prophecies.

Whereas previous work has sought to address this failure either by abandoning the axiom of rationality or by introducing multiple equilibria, we argue that this failure rests merely on the convenient, but unrealistic, assumption that communication is perfect. If agents could perfectly communicate with one another or otherwise reach symmetric information—the case assumed in the bulk of the macroeconomics literature—they could also reach agreement on a course of action that depends only on their beliefs of the fundamentals. In contrast, as long as communication is imperfect—the case assumed in this paper—firms, consumers, and investors alike may face significant uncertainty about aggregate economic activity and market prices *beyond* any uncertainty they face about the fundamentals. Moreover, as economic agents find it in their best interest to respond to this uncertainty, equilibrium outcomes are no more pinned down by beliefs of fundamentals, which in turn only reinforces this uncertainty—the fear or the confidence of one agent feeds the fear or the confidence of another agent. It is this distinct form of uncertainty that sustains the sunspot fluctuations in our model and that permits us to capture the self-fulfilling nature of fluctuations within an otherwise canonical, unique-equilibrium macro model.

Turning to the normative properties of our theory, we observe that the *only* reason that the equilibrium fails to be first-best efficient is that communication is imperfect—if communication had been perfect, the "invisible hand" would have worked flawlessly, guaranteeing, not only that our sunspot fluctuations disappear, but also that welfare attains its highest possible level. It follows that improving communication or otherwise reducing the asymmetry of information can certainly raise welfare. Short of doing this, however, there is no way to improve upon the equilibrium. In particular, conventional stabilization policies, which aim at stabilizing activity without influencing communication, can only reduce welfare.

Formally, we show that the equilibrium coincides with the solution to the problem of a planner that faces the same resource and communication constraints as the market mechanism. This result is a contribution on its own right: it establishes, in effect, a variant of the first welfare theorem for a class of economies that feature dispersed information, higher-order uncertainty, and sunspot volatility. But it also instrumental in delivering a very different policy message from the one advocated by Keynesian thinking: even if fluctuations are "evidently" disconnected from fundamentals, and driven by "animal spirits", this is not an invitation for stabilization policies.

Methodological contribution and related literature. Our paper complements the large body of work that has formalized "animal spirits" in models with non-convexities, multiple equilibria and coordination failures.[1] We share with this literature the desire to capture the self-fulfilling nature of fluctuations. But we have provided a novel formalization that achieves the same goal without abandoning either equilibrium uniqueness or the spirit of the welfare theorems.

Our paper also adds to the large, and growing, literature on business cycles with informational frictions.[2] Our contribution, however, is distinct in at least two respects. First, we abandon

[1]Key contributions include Azariadis (1981), Benhabib and Farmer (1994, 1999), Cass and Shell (1983), Gali (1992), Farmer (1993), Farmer and Woodford (1997), Guesnerie and Woodford (1992), Kiyotaki (1988), Shell (1977), and Woodford (1986, 1987, 1991). See also a recent paper by Wang and Wen (2009) that shows the existence of multiple correlated equilibria in the canonical monopolistic Dixit-Siglitz framework.

[2]Phelps (1970), Lucas (1972, 1975), Barro (1976), King (1982), and Townsend (1983) are early seminal contribu-

altogether the theme of "nominal confusion" and instead focus on a real, non-monetary economy. Second, and most importantly, we develop an entirely novel theory of fluctuations—one that helps capture the distinct, self-fulfilling nature of short-run fluctuations within otherwise an canonical, unique-equilibrium macro model.

In delivering this contribution, we build on recent work by Morris and Shin (2002), Angeletos and Pavan (2007), and others on a class of games with incomplete information and strategic complementarity, often referred to as "beauty contest games". This earlier work has not studied micro-founded business-cycle models and has not identified the sunspot fluctuations that are at the core of our contribution; but it has highlighted other related implications of dispersed information, such as the anchoring effect of public news. We uncover the links between our framework and this earlier work by establishing an isomorphism between the competitive equilibrium of our economy and the Perfect Bayesian Equilibrium of a game similar to those studied in this work. We then show that, under the lens of this isomorphism, our sunspot fluctuations can be attributed to random variation in higher-order beliefs ("the forecasts of the forecasts of others") that is independent of either first-order beliefs ("the forecasts of the fundamentals") or the fundamentals themselves.

That being said, it is important to recognize that our model is a Walrasian economy, not a game. In a Walrasian setting, agents do no care per se about the actions and beliefs of other agents. Rather, they only care about the prices they face in competitive markets—in Lucas' (1972) tradition, they only need to know the correct statistical model of these prices. Thus, barring the aforementioned game-theoretic isomorphism, our sunspot fluctuations are best understood as a novel form of self-fulfilling extrinsic uncertainty in allocations and prices.

Moreover, the only sense in which our economy features strategic complementarity is through a plain-vanilla pecuniary externality: while each agent takes the terms of trade as exogenous to his individual choices, these terms of trade ultimately depend on the agents' collective behavior. Part of our contribution is thus to show how such otherwise inconsequential pecuniary externalities can be the source of higher-order uncertainty and sunspot volatility once communication is imperfect.

Finally, note that communication can be imperfect in a Walrasian setting like ours *only* if one relaxes the assumption of centralized, Arrow-Debreu markets. Our analysis thus rests on introducing some segmentation in market interactions. However, we do no introduce any ad-hoc "noise traders" and the like, for this would have impeded the normative contribution of our paper.

More specifically, we formalize the initial dispersion of information by splitting the economy into different "islands"; this is similar to, inter alia, Lucas (1972), Towsend (1983), Lagos and Wright (2005), Lorenzoni (2008), and Amador and Weill (2009). We then model the dynamics of

tions; more recent work includes, inter alia, Mankiw and Reis (2002), Sims (2003), Woodford (2003), Mackowiak and Wiederholt (2009), Lorenzoni (2010), Amador and Weill (2010), and Hassan and Mertens (2010) communication by assuming that islands trade and communicate with one another only through random pair-wise matches. As a result, information gets aggregated only slowly and in a manner that resembles the "percolation of information" in Duffie and Manso (2007) and Golosov, Lorenzoni and Tsyvinski (2009). But whereas in these models—as in most of the finance literature—communication is relevant only in so far it helps learn about an exogenous fundamental (the exogenous dividend), in our model it is relevant because it ultimately helps forecast the endogenous level of economic activity. This difference is crucial and helps explain why communication may, paradoxically, propagate and intensify the impact of "animal spirits" in our economy.

Layout. The rest of the paper is organized as follows. Section 2 introduces our baseline model. Section 3 characterizes the equilibrium. Section 4 contains our core positive lesson regarding the possibility of self-fulfilling fluctuations. Section 5 studies a specific example of waves of optimism and pessimism. Section 6 extends the analysis to a variant model that helps abstract from incomplete risk sharing. Section 7 identifies policies that can remove the sunspot fluctuations. Section 8 contains our core normative lesson regarding the efficiency of our fluctuations. Section 9 concludes with a discussion of our contribution. The Appendix includes any proofs omitted in the main text.

2 The model

Our economy consists of a continuum of islands, indexed by $i \in \mathcal{I} \equiv [0,1]$. Each island is populated by a representative household, which includes a consumer and two workers, and by two representative firms. The one firm specializes in the production of a differentiated consumption good that can be consumed either by the same island or by other islands. The other firm specializes in the production of a differentiated capital good that can be used as an input only locally. Similarly, the one worker specializes in the production of the consumption good, while the other worker specializes in the production of the consumption good.

Time is discrete, indexed by $t \in \{0, 1, ...\}$. Each period contains two stages. The production of the consumption goods takes place in stage 1, while any trading across islands and the production of the capital goods takes place in stage 2. Accordingly, two labor markets operate in every island, one in stage 1 and another in stage 2. Finally, capital is used as an input only in the production of the consumption good, not in its own production. Accordingly, a market for renting the locally available capital operates during stage 1 alongside the corresponding labor market, while the market for purchases of new capital operates in stage 2.

Firms and technologies. Consider, first the firm that produces the differentiated good of island i. The technology of this firm is given by the following:

$$y_{i,t} = A_i (k_{i,t})^{\vartheta_k} (n_{i,t})^{\vartheta_n}, \tag{1}$$

where $y_{i,t}$ is the quantity produced, $k_{i,t}$ is the installed capital stock, $n_{i,t}$ is the first-stage employment, A_i is the local productivity (TFP), and $\vartheta_k, \vartheta_n \in (0,1)$, with $\vartheta_k + \vartheta_n \leq 1$. The realized profit is given by

$$\pi^c_{i,t} = p_{i,t} y_{i,t} - w_{i,t} n_{i,t} - r_{i,t} k_{i,t},$$

where $p_{i,t}$ denotes the local price of the local differentiated commodity, $w_{i,t}$ denotes the local wage for the workers that produce the consumption good, and $r_{i,t}$ denotes the local rental rate of capital. It follows that each firm chooses its labor and capital input so as to maximize its expectation of the local valuation of its profits, namely $\lambda_{i,t}\pi^c_{i,t}$, where $\lambda_{i,t}$ denotes the marginal value of wealth for the representative household in island i.

Next, consider the firm that produces the local capital good. The technology of this firm is given by

$$x_{i,t} = (\tilde{n}_{i,t})^{\psi} \tag{2}$$

where $x_{i,t}$ is the production of new capital goods, $\tilde{n}_{i,t}$ is the second-stage employment, and $\psi \in (0,1)$. The profit is given

$$\pi^k_{i,t} = q_{i,t} x_{i,t} - \tilde{w}_{i,t} \tilde{n}_{i,t}$$

where $q_{i,t}$ denotes the local price of capital and $\tilde{w}_{i,t}$ the local wage for the workers that produce the capital good. The firm maximizes its expectation of $\lambda_{i,t}\pi^k_{i,t}$.

Households and preferences. The utility of the representative household on island i is given by the following:

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

where $\beta \in (0,1)$ is the discount factor, $U : \mathbb{R}^2_+ \to \mathbb{R}_+$ is a strictly increasing and concave function, $V : \mathbb{R}_+ \to \mathbb{R}_+$ is a strictly increasing and convex function, c_{it} is "home consumption" (consumption of the local good), c^*_{it} is "foreign consumption" (consumption of the good of another island), n_{it} is labor in the first stage, and $\tilde{n}_{it}$ is labor in the second stage. To simplify the exposition, our baseline model assumes the following specifications for U and V:

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

for some $\eta \in [0,1]$. We will relax these assumptions in Section 6.

Trading and budgets. Trading occurs only between islands that get matched together during stage 2 of each period. When the two islands meet, in principle they could trade all the goods in the economy, as well as an arbitrary set of financial claims (or insurance contracts). For expositional simplicity, our baseline model restricts the islands to trade only their differentiated consumption goods. We relax this assumption in Section 6.[3]

Given that—for the time being—we rule out trade in financial claims, the period-t budget constraint of the household can be written as follows:

$$p_{i,t}c_{i,t} + p^*_{i,t}c^*_{i,t} + q_{i,t}x_{i,t} \leq \pi^c_{i,t} + w_{i,t}n_{i,t} + r_{i,t}k_{i,t} + \pi^k_{i,t} + \tilde{w}_{i,t}\tilde{n}_{i,t}$$

The left-hand side of this budget is the total expenditure in consumption and investment goods, with $p^*_{i,t}$ denoting the local price of the differentiated good that island i "imports" from the island it got matched with during stage 2. The right-hand side is the total income from wages, capital returns, and profits. Finally, the law of motion for capital is given by

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

where $\delta \in [0,1]$ is the depreciation rate. For reasons of tractability, our core positive results assume full depreciation ($\delta = 1$); at certain points, however, we allow for the more general case in order to facilitate a more complete discussion of the underlying economics.

Numeraire. We henceforth let the second-stage leisure be the numeraire good and accordingly normalize $\tilde{w}_{i,t} = 1$. The prices $p_{i,t}$, $p^*_{i,t}$, $q_{i,t}$ etc. should thus be interpreted as the local prices in island i relative to the numeraire. Also, note that the relative price of the two tradeable differentiated goods of any two matched islands must of course be the same across the two islands: $p_{i,t}/p^*_{i,t} = p^*_{j,t}/p_{j,t}$. However, as long as the numeraire is not traded, $p_{i,t}$ need not be the same with $p^*_{j,t}$: the relative price of the non-tradeable good may well differ across the two islands.

Fundamentals and information. The only uncertain component of the fundamentals of our economy is the distribution of productivities in the cross-section of islands. We assume that this distribution, while being randomly picked at $t = 0$, stays invariant thereafter. This assumption simplifies the analysis, while also helping clarify that the short-run fluctuations that we will document are not be driven by high-frequency movements in the underlying fundamentals. We then specify the information structure of the economy as follows.

In the beginning of time, Nature draws two independent random variables, θ and ξ. The former is drawn from a finite set $\mathcal{S}_\theta$ according to a probability function $\mathcal{F}_\theta(\theta)$, while the latter is drawn from a finite set $\mathcal{S}_\xi$ according to a distribution function $\mathcal{F}_\xi(\xi)$. The fundamentals of the economy are affected by θ but not by ξ. In particular, local productivities are i.i.d. draws from a finite

[3] In particular, in that section we let the numeraire be traded. The linearity of preferences with respect to the numeraire then guarantees that it is immaterial whether agents across different islands can also trade the capital good, or whether we open markets for financial claims. Indeed, as in Lagos and Wright (2005), this linearity will make sure that any wealth inequality that may be generated during period t gets absorbed by variation in $\tilde{n}_{i,t}$.

set $\mathcal{S}_A \subset \mathbb{R}_+$ according to a conditional probability function $\mathcal{F}_A(A_i|\theta)$, which only depends on θ. We can thus interpret θ as the underlying aggregate fundamental and ξ as a form of noise. We henceforth let $z \equiv (\theta, \xi)$ and refer to it as the underlying aggregate state; its support is given by $\mathcal{S}_z \equiv \mathcal{S}_\theta \times \mathcal{S}_\xi$ and its probability function by $\mathcal{F}(\theta, \xi) \equiv \mathcal{F}_\theta(\theta)\mathcal{F}_\xi(\xi)$.

For any given z, Nature draws a continuum of random variables ω, one for each island, from a finite set $\mathcal{S}_\omega$ according to a conditional probability function $\mathcal{P}(\omega|z)$. These draws are i.i.d. across the islands, so that, by the usual convention, $\mathcal{P}(\omega|z)$ is also the fraction of islands that receive a particular realization of ω when the aggregate state is z. Each island gets to see its own ω_i in the beginning of time, but not the aggregate state z or the ω_j of any other island $j \neq i$. The variable ω thus identifies the initial information set of an island. It can be interpreted as a local signal about the underlying aggregate state and it encodes information, not only about the underlying fundamentals, but also about the likely information sets of other island. In this sense, ω is the analogue in our competitive economy of a Harsanyi type in games, while the probability functions $\mathcal{F}$ and $\mathcal{P}$ define the initial common prior. Finally, to simplify notation, and without any loss of generality, we henceforth assume that ω_i reveals the local productivity: there exists a function $A : \mathcal{S}_\omega \to \mathcal{S}_A$ such that $A_i = A(\omega_i)$. We accordingly refer to ω interchangeably as the "type" or the "local state" of an island.

This completes the description of the underlying fundamentals and the initial information structure. The timing of choices, the trading interactions of the islands, and the diffusion of information are then assumed to be as follows. In stage 1 of each period, each island i operates in isolation. The local labor and rental markets open, households decide how much labor and capital to supply, and the firms producing the differentiated goods make the input and production choices. In stage 2, each island i is randomly matched with another island $j \neq i$. At this point, any two islands in the same match share their information with each other and trade their differentiated commodities. In addition, a local market for new capital (investment) opens in each island, and the firms producing the capital goods make their input and production choices.

Let $\omega_{i,t}$ denote the information set of island i during stage 1 of period t. The (stochastic) sequence of $\omega_{i,t}$ is defined recursively by letting $\omega_{i,0} = \omega_i$ and, for all $t \geq 1$, $\omega_{i,t} = (\omega_{i,t-1}, \omega_{m(i,t),t-1})$, where $m(i,t)$ identifies the (random) island that got matched with island i during stage 2 of period $t-1$. We let $\mathcal{S}_t$ denote the support of $\omega_{i,t}$ and $\mathcal{P}(\omega_{i,t}|z)$ its probability conditional on the underlying aggregate state being z. For expositional simplicity, we assume that the probability that any two islands meet at the trading stage is independent of their type; we will relax this assumption in some examples. Finally, we assume that there is a finite period T (which, though, could be arbitrarily large) such that the aggregate state z, and all the local histories up to this point, become common knowledge thereafter.

Remarks. (i) For any given two islands i and $j \neq i$, the local productivities A_i and A_j can be correlated only through the aggregate fundamental θ, but the local types ω_i and ω_j can also be correlated through the noise ξ. We will later see how this formalization permits us to introduce sunspot fluctuations despite the uniqueness of equilibrium.

(ii) What is crucial for our results is that agents have asymmetric information about the *aggregate* state z; whether they also have private information about their *idiosyncratic* productivities, as we have assumed here, is of secondary importance. This distinguishes our approach from the Mirrleess/New Public Finance literature, which allows agents to have private information about their *idiosyncratic* productivities or tastes but maintains the conventional assumption of symmetric information about the aggregate state.[4]

(iii) In addition to the exogenous flow of information that we have specified above, agents get to observe the prices in the markets they participate. However, prices do not reveal any additional information to them. This is because we have assumed that any two islands that get matched together share their information by "talking" to one another before they trade. Clearly, this is only a convenient abstraction that permits us to rule out the information externalities that could obtain if communication took place only through prices or actions.[5]

(iv) In many of the recent papers on business cycles with information frictions, economic agents suffer from a certain form of "schizophrenia": they have multiple personalities that fail to communicate with one another.[6] While we are quite sympathetic to such abstractions, we also think that this issue poses a non-trivial modeling challenge. In this paper we choose to face this challenge by ruling out any such "schizophrenia" and instead introducing communication frictions only *across* different agents. We think that this adds to the transparency of our theoretical exercise.

3 Equilibrium

To fix language, we henceforth say that two islands i and j ($j \neq i$) are "trading partners", or form a "trading pair", in period t if they are matched together in the second stage of that period. Since communication takes place in stage 2, the information set during state 2 of that period is the same as the one during stage 1 of the next period, that is, it is given by $\omega_{i,t+1}$. It follows that any choices

[4] See Angeletos and Pavan (2009) for a broader discussion of how asymmetric information about aggregate shocks has distinct implications for equilibrium and policy analysis.

[5] See Amador and Weill (2009), Angeletos and Pavan (2009), and Angeletos and La'O (2010) for an analysis of such information externalities in macro settings.

[6] For example, this is the case in Mankiw and Reis (2002), Woodford (2003), and Mackowiak and Wiederholt (2009) alike, where the personality of the firm that adjusts employment and production does not "speak" to the personality that sets prices, as well as in Angeletos and La'O (2009), where each household ("family") has multiple members that cannot communicate with one another at some points of time.

and prices that are determined during stage 1 of period t are contingent on $\omega_{i,t}$, while those that are determined during stage 2 of that period are contingent on $\omega_{i,t+1}$. We can thus represent any allocation and price system with a collection of functions $(p, p^*, w, r, n, k, y, c, c^*, \tilde{n}, \tilde{x})$ such that the following are true:

- $p_{i,t} = p_t(\omega_{i,t+1})$ and $p^*_{i,t} = p^*_t(\omega_{i,t+1})$, with $\omega_{i,t+1} = (\omega_{i,t}, \omega_{j,t})$, are the domestic prices of, respectively, the domestic and the foreign differentiated commodity in island i when its trading partner is j and when their respective information sets before this match are $\omega_{i,t}$ and $\omega_{j,t}$.
- $c_{i,t} = c_t(\omega_{i,t+1})$ is the consumption by island i of its own differentiated good, and $c^*_{i,t} = c^*_t(\omega_{i,t+1})$ is the consumption by island i of its trading partner's differentiated good.
- $\tilde{n}_{i,t} = \tilde{n}_t(\omega_{i,t+1})$ and $\tilde{x}_{i,t} = \tilde{x}_t(\omega_{i,t+1})$ are the corresponding levels of employment and output in the the capital-producing sector, with $\tilde{x}_{i,t} = \tilde{n}^{\vartheta}_{i,t}$.
- $w_{i,t} = w_t(\omega_{i,t})$ and $r_{i,t} = r_t(\omega_{i,t})$ are the wage rate and the rental rate in island i during stage 1 of period t, when the local information of this island is $\omega_{i,t}$.
- $n_{i,t} = n_t(\omega_{i,t})$, $k_{i,t} = k_t(\omega_{i,t})$, and $y_{i,t} = y_t(\omega_{i,t})$ are the corresponding levels of employment, capital, and output in the local differentiated-good sector, with $y_{i,t}(\omega_{i,t}) = A_i k^{1-\vartheta}_{i,t} n^{\vartheta}_{i,t}$.
- $q_{i,t} = q_t(\omega_{i,t+1})$ is the corresponding local price of capital, and $k_{i,t+1} = k_{t+1}(\omega_{i,t+1})$ is the capital stock next period, with $k_{i,t+1} = (1-\delta)k_{i,t} + \tilde{x}_{i,t}$.

There is nothing peculiar with the preceeding notation; it only serves to highlight the information upon which allocations and prices can be contingent. With this in mind, we define a competitive equilibrium in the familiar manner.

Definition 1. *A collection of functions* $(p, p^*, w, r, n, k, y, c, c^*, \tilde{n}, \tilde{x})$ *identifies an equilibrium if and only if:*

(i) the associated quantities are optimal for the households and firms;

(ii) the associated prices clear all markets;

(iii) the trade of differentiated goods is balanced: $p^*_{i,t} c^*_{i,t} = p_{i,t}(y_{i,t} - c_{i,t})$

Condition (iii) highlights that the numeraire is not traded. As anticipated, this is only for reasons of tractability and will be relaxed in Section 6.

We simplify the equilibrium analysis by splitting it into three steps. The first step characterizes the "static equilibrium" that obtains in any given period t for a given level of capital. The second step studies the determination of investment and the price of capital. The third step obtains a tractable representation of the general equilibrium under the restriction of full depreciation.

Static equilibrium. Consider the behavior of the household of island i. Let $\lambda_{i,t} = \lambda_t(\omega_{i,t})$ denote the marginal value of wealth in stage 2 of period t (the Lagrange multiplier on the period-t budget constraint). The optimality condition for labor supply during stage 1 and 2 give, respectively,

$$V'(n_{i,t}) = \mathbb{E}_{i,t}[\lambda_{i,t}]w_{i,t} \quad \text{and} \quad \lambda_{i,t} = 1. \tag{3}$$

where $\mathbb{E}_{i,t}[\cdot]$ is a short-cut for the rational expectation conditional on $\omega_{i,t}$ (the information set that is available in stage 1 of period t and stage 2 of period $t-1$). The optimal consumption choices, on the other hand, satisfy

$$U_c\left(c_{i,t}, c^*_{i,t}\right) = \lambda_{i,t}p_{i,t} \quad \text{and} \quad U_{c^*}\left(c_{i,t}, c^*_{i,t}\right) = \lambda_{i,t}p^*_{i,t} \tag{4}$$

Combining these conditions with the corresponding conditions for i's trading partner, imposing market clearing and balanced trade, and using the Cobb-Douglas specification of U, we reach the following characterization of the equilibrium consumption allocations:

$$c_{i,t} = (1-\eta)y_{i,t}, \quad c^*_{i,t} = \eta y_{j,t}, \quad c_{j,t} = (1-\eta)y_{j,t}, \quad \text{and} \quad c^*_{j,t} = \eta y_{i,t}, \tag{5}$$

for any trading pair (i,j). Together with the fact that $\lambda_{i,t} = 1$, this implies the following characterization of the local price of the local differentiated good relative to the local numeraire:

$$p_{i,t} = U_c\left((1-\eta)y_{i,t}, \eta y_{j,t}\right) = y_{i,t}^{-\eta}y_{j,t}^{\eta} \tag{6}$$

This last condition plays a central role in the subsequent analysis. Note, in particular, that the price of the local good decreases with local output and increases with the output of the trading partner. This is intuitive: by analogy to international trade, we could say that the price of "exports" decreases with "domestic" output and increases with "foreign" output. More broadly, it captures the idea that the terms of trade faced by any agent in the economy are dependent on the economic choices of other agents. As anticipated in the Introduction, this source of interdependence—a pecuniary externality of a familiar Walrasian type—will alone explain why our economy can be subject to a certain form of sunspot fluctuations.

Consider, next, the behavior of the firms that produce the differentiated consumption goods. These firms choose their labor and capital inputs in stage 1 so as to maximize $\mathbb{E}_{i,t}[\lambda_{i,t}\pi_{i,t}]$. Using the fact that $\lambda_{i,t} = 1$, the first-order condition with respect to labor yields

$$w_{i,t} = \mathbb{E}_{i,t}\left[p_{i,t}\right]\vartheta_n \frac{y_{i,t}}{n_{i,t}}, \tag{7}$$

while the one with respect to capital yields

$$r_{i,t} = \mathbb{E}_{i,t}\left[p_{i,t}\right]\vartheta_k \frac{y_{i,t}}{k_{i,t}}. \tag{8}$$

These conditions have a familiar interpretation: the wage and the rental rate of capital are equated to the expected marginal revenue products of, respectively, labor and capital. The only novelty vis-a-vis more conventional macroeconomic models is that here firms face uncertainty about the prices at which they will sell their products and base their expectations on geographically differentiated information.

Combining conditions (3) and (7), we infer that the equilibrium level of employment satisfies the following optimality condition:

$$V'(n_{i,t}) = \mathbb{E}_{i,t}[p_{i,t}]\vartheta_n \frac{y_{i,t}}{n_{i,t}} \tag{9}$$

This condition has a simple interpretation: the marginal cost of employment (or the disutility of effort) is equated with its marginal revenue product. Furthermore, once this condition is combined with the technology in (1) and the equilibrium price in (6), it permits us to express the equilibrium level of local output $y_{i,t}$ of island i as merely a function of the local productivity A_i, the inherited local capital stock $k_{i,t}$, and the local expectations of the output $y_{j,t}$ of the island's likely trading partner. We thus reach the following result.

Proposition 1. *Let* $\alpha \equiv \frac{\eta}{1/\vartheta_n - 1 + \eta} \in (0,1)$. *For any* $t \geq 0$, *the equilibrium level of output satisfies*

$$\log y(\omega_{i,t}) = (1-\alpha)\left(\tfrac{1}{1-\vartheta_n}\log A(\omega_i) + \tfrac{\vartheta_k}{1-\vartheta_n}\log k_t(\omega_{i,t})\right) + \alpha\,\mathbf{E}\left[\log y(\omega_{j,t})|\,\omega_{i,t}\right] \tag{10}$$

where $\mathbf{E}$ *is a risk-adjusted expectations operator defined by* $\mathbf{E}[X|\omega] \equiv \frac{1}{\eta}\log\mathbb{E}\left[\exp(\eta X)|\omega\right]$.

This result has an appealing game-theoretic interpretation—one that uncovers a certain link between our economy and the more abstract work of Morris and Shin (2002), Angeletos and Pavan (2007) and others on a certain class of games with linear best responses and incomplete information, ofter referred to as "beauty contest" games. To see this, let $f_{i,t} \equiv \frac{1}{1-\vartheta_n}\log A_i + \frac{\vartheta_k}{1-\vartheta_n}\log k_{i,t}$ and rewrite condition (10) as

$$\log y_{i,t} = (1-\alpha)f_{i,t} + \alpha\,\mathbf{E}_{i,t}[\log y_{j,t}] \tag{11}$$

Next, fix a period t and a history up to that point. The equilibrium allocations of our economy in that period can then be obtained as the Perfect Bayesian equilibrium of a fictitious game. The islands of our economy identify the players of this game; the variables $f_{i,t}$ identify their "fundamentals" (as given either by nature or history); the above condition identifies their best responses; and the coefficient α identifies the degree of strategic complementarity. This game theoretic representation—and a related one that we establish in Proposition 2 below—facilitates a certain interpretation of the sunspot fluctuations that we will document in the next section. But whereas this earlier work typically pushes aside the microfoundations of any specific application, in our paper these microfoundations are, not only spelled out fully, but also play a central role in our positive and normative results.

We will return to these points in subsequent discussion. For now, we would like to remind the reader that none of the agents in our economy is truly strategic: all agents are atomistic. Hence, the notion of strategic interaction that is present in our economy is nothing else than the general-equilibrium interaction that is generic to any Walrasian economy. Indeed, the *only* source of strategic complementarity in our setting is the possibility to trade in Walrasian markets.

Finally, what was identified as "fundamental" in the aforementioned game, namely the variable $f_{i,t}$, is predetermined in period t but it is not exogenous. Rather, it depends on past investment choices. To close the characterization of the equilibrium, we thus have to turn attention to the dynamics of investment.

Investment and asset prices. Consider the firms that produce the local capital goods. Given the normalization $\tilde{w}_{i,t} = 1$, the optimality condition for these firms gives

$$1 = q_{i,t}\psi\tilde{n}_{i,t}^{\psi-1}.$$

Since $x_{i,t} = \tilde{n}_{i,t}^{\psi}$, the above gives gross investment as an increasing function of the price of capital:

$$k_{i,t+1} - (1-\delta)k_{i,t} = x_{i,t} = \left(\frac{q_{i,t}}{\psi}\right)^{\frac{\psi}{1-\psi}} \tag{12}$$

This is akin to standard Q theory and in this particular regard there is nothing novel in our model relative to conventional macroeconomic models.

However, there is an important novelty when we look at the determination of the price of capital. From the Euler condition of the households on island i, we have that the price of capital satisfies

$$q_{i,t} = \beta\left\{\mathbb{E}_{i,t+1}[r_{i,t+1}] + (1-\delta)q_{i,t}\right\} \tag{13}$$

Iterating this condition forward, and using the facts that $r_{i,t} = \mathbb{E}_{i,t}[p_{i,t}]\vartheta_k\frac{y_{i,t}}{k_{i,t}}$ and $p_{i,t} = y_{j,t}^{\eta}y_{i,t}^{-\eta}$, we infer that

$$q_{i,t} = \beta\mathbb{E}_{i,t+1}\left[\sum_{s=t+1}^{\infty}(\beta(1-\delta))^{s-t-1}\vartheta_k y_{j,s}^{\eta}y_{i,s}^{1-\eta}k_{i,s}^{-1}\right]$$

This result identifies how higher-order uncertainty can impact asset prices and investment in our economy: the local price of capital $q_{i,t}$ depends on local expectations of future terms of trade $p_{i,s}$ and, through them, on local expectations of the output $y_{j,s}$ of the island's likely trading partners in all future periods $s \geq t$.[7]

A similar result holds for other asset prices in our model. For example, suppose that we introduce a claim on the profits of the firms that produce the local differentiated commodity and let the

[7] Note that the the identity of i's trading partner is random and may change from period to period.

local agents trade this claim in stage 2 of each period. Since realized profits are given by $\pi^k_{i,t} = (1-\vartheta_k-\vartheta_n)p_{i,t}y_{i,t} = y^{\eta}_{j,t}y^{1-\eta}_{i,t}$, the price of this claim is given by

$$Q_{i,t} \equiv \mathbb{E}_{i,t+1}\left[\sum_{s=t+1}^{\infty}\beta^{s-t}\pi^k_{i,t}\right] = \mathbb{E}_{i,t+1}\left[\sum_{s=t+1}^{\infty}\beta^{s-t}(1-\vartheta_k-\vartheta_n)y^{\eta}_{j,s}y^{1-\eta}_{i,s}\right]$$

and once again depends on expectations of the future economic activity in other islands.

We conclude that asset prices—and thereby investment—depend on expectations of the choices of others in the future. But the choices of others in the future depend on their own expectations of the choices of others. This underscores the self-fulfilling nature of the financial and investment side of our economy: high asset prices and investment in one part of the economy (one island) can cause an asset price and investment boom in its likely trading partners, which in turn can justify its own investment and asset price boom.

We will give it a more concrete positive content to this insight in the next two sections, by showing how this opens the door to sunspot volatility in asset prices despite the uniqueness of equilibrium and the absence of noise traders or other irrational forces.[8]

For now, we note that for reasons of tractability we prefer to assume full depreciation for capital ($\delta = 1$). This kills a lot of the interesting action in the price of capital and investment: with $\delta = 1$, the price of capital and investment in period t depend only on expectations of economic activity in period $t+1$, not further in the future. However, we can still understand what the likely behavior of the price of capital would have been with $\delta < 1$ by looking at $Q_{i,t}$, the price of the claims on profits (or the price of land). Indeed, the latter still depends on expectation of economic activity in all periods after t, not only $t+1$.

Complete equilibrium characterization. By assuming full depreciation, we can replace $x_{i,t}$ in condition (12) with $k_{i,t+1}$. Moreover, condition (13) reduces to $q_{i,t} = \beta\mathbb{E}_{i,t+1}[r_{i,t+1}]$. We infer that the equilibrium level of investment satisfies the following condition:

$$k_{i,t}^{\frac{1}{\psi}} = \beta\psi\vartheta_k\mathbb{E}_{i,t}\left[p_{i,t}\right]y_{i,t} \tag{14}$$

This condition has a simple interpretation: the marginal cost of investment is equated with its marginal revenue product. This is similar to the condition we encountered earlier for employment,

[8]Incidentally, this in tun helps capture Keynes' famous "beauty-contest metaphor" for financial markets: instead of asset prices be driven merely by expectations of the underlying fundamentals, they can be driven by expectations of ... the expectations of ... the expectations of .... others. In this respect, our paper complements, inter alia, Allen, Morris and Shin (2005), which also seeks to capture Keynes' beauty contest metaphor, but do so by focusing on the interaction of rational short-term traders and irrational noise traders; and Angeletos, Lorenzoni and Pavan (2009), which shift focus to certain information externalities in the two-way interaction between real investment and asset prices. At the same time, and quite importantly, the normative results of our paper will clarify that, unlike what suggested by Keynes and most funs of his metaphor, the impact of higher-order on asset prices need not be a symptom of inefficiency.

namely condition (9). By combining these two conditions, we can now reach the following characterization of the equilibrium level of output as a function of merely the local productivity and the local belief of others' output.

Proposition 2. *Let* $\hat{\vartheta} \equiv \vartheta_n + \psi\vartheta_k \in (0,1)$ *and*

$$\hat{\alpha} \equiv \frac{\eta}{1/\hat{\vartheta} - 1 + \eta} \in (0,1).$$

For any $t \geq 1$*, the period-*t *equilibrium levels of output solve the following fixed-point problem:*

$$\log y(\omega_{i,t}) = (1-\hat{\alpha})\left(\tfrac{1}{1-\hat{\vartheta}}\log A(\omega_i)\right) + \hat{\alpha}\,\mathbf{E}\left[\log y(\omega_{j,t})|\,\omega_{i,t}\right]. \tag{15}$$

This proposition, which rests on full depreciation, permits us to contain the forward-looking aspects of capital and thereby to treat the dynamic equilibrium of the economy as an essentially static one: equilibrium output depends on beliefs about economic activity only in the present (or the immediate future). Of course, this should not be taken literally; it is only a convenient abstraction. Moreover, by hiding the forward-looking aspects that would have otherwise obtained, this abstraction most likely biases the results *against* us: if we had permitted these forward-looking aspects to operate, we could obtain sunspot fluctuations even when the current fundamentals are common knowledge—for it would suffice that there is dispersed information about uncertain future fundamentals. Nevertheless, this result affords us a level of tractability that we find very valuable. Furthermore, as already mentioned, we will be able to detect some of the forward-looking aspects by looking at the price $Q_{i,t}$.

Proposition 2 admits a similar game-theoretic interpretation as Proposition 1, except that now we have been able to reduce out the capital stock and hence to express the level of output in an island merely as a function of the local productivity and the local expectations of the output in other islands. Also, the associated degree of strategic complementarity is now larger ($\hat{\alpha} > \alpha$); this is thanks to the endogenous adjustment of capital, which alleviates the diminishing returns to labor ($\hat{\vartheta} > \vartheta_n$). Nevertheless, the origin of complementarity remains as before: the complementarity is present only in so far agents trade (i.e., $\hat{\alpha} > 0$ if and only if $\eta > 0$) and is increasing in the intensity of trade (i.e., $\hat{\alpha}$ is increasing in η).

Note that Proposition 2 contains the key information that one needs in order to obtain the entire general equilibrium of the economy: once we have solved the equilibrium output levels from Proposition 2, we can readily recover the equilibrium levels of investment, employment, consumption, and the various prices from the preceding analysis. Finally, the fact that $\alpha \in [0,1)$, which intuitively means that the best response has a slope less than 45 degrees, guarantees that condition (15) admits a unique fixed point, which in turn gives us the following result.

Proposition 3. *The equilibrium exists and is unique.*

4 Imperfect communication and sunspot fluctuations

We are now ready to prove the core positive result of the paper. We first define sunspot fluctuations in a way that is consistent with the standard definition used in the literature—except that in our case we are interested in economies with a unique equilibrium.

Definition 2. *We say that an economy features **sunspot fluctuations** if and only if there exists shocks that cause variation in equilibrium allocations and prices without causing variation in either the underlying fundamentals (here productivities) or any agents' expectations of these fundamentals.*

We can then state our core positive result as follows.

Theorem 1. *The economy can feature sunspot fluctuations along its unique equilibrium if, and only if, information is asymmetric.*

We establish the "if" part—namely, that that sunspot fluctuations can obtain when information is asymmetric—by example. For expositional simplicity, this example allows the state space to be continuous and the matching to be contingent on an island's type.

The fundamental θ is given by a pair of variables a_1 and a_2. The first variable has a degenerate distribution: it always takes value 1. The second is drawn from a Normal distribution with mean zero and variance 1. Conditional on θ, an island's productivity is randomly drawn and can be $A_1 = \exp(a_1)$ with probability $1/2$ and $A_2 = \exp(a_2)$ with probability $1/2$. Thus, half of the islands have productivity A_1 (which is non-random) and the other half have productivity A_2 (which is random). Let us call the first set of islands "group 1" and the second "group 2" and let us assume that matching occurs only across the two groups: each island will trade with an island from the other group.

At $t = 0$, each island gets to see its own productivity, but not the productivities of other islands. Of course, since a_1 is non-random, this means that islands in group 2 are fully informed about the underlying economic fundamentals; but islands in group 1 face uncertainty about the productivity of group 2. We model any information that group 1 may have about the productivity of group 2 by assuming that islands in group 1 get to observe a signal about the productivity of the other group. This is given by $s_1 = a_2 + \varepsilon_1$, were ε_1 is noise, independent of a_2, and drawn from a Normal distribution with mean zero and variance 1. Because s_1 is observed only by group 1, group 2 can now face uncertainty about the likely information of group 1—even though group 2 is perfectly informed about the fundamentals in the economy, it can't be sure about what group 1 knows about the fundamentals. We model any information that group 2 may have about the group 1's knowledge by assuming that group 2 gets to observe a signal about s_1. This signal is given by $s_2 = s_1 + \varepsilon_2$, where ε_2 is noise, independent of either θ or ε_1 (and hence also of s_1), and drawn from a Normal distribution with mean 0 and variance 1.

Consider now equilibrium behavior at $t = 0$. It is obvious that variation in ε_2 does not cause variation in the fundamentals; this is by the assumption that ε_2 is independent of θ. What is more, ε_2 does not affect anybody's expectations about the fundamentals: for islands in group 1, their expectations of the fundamentals are pinned down by s_1; as for islands in group 2, they already know perfectly the fundamentals.

Nevertheless, it can be shown that variation in ε_2 causes variation in equilibrium allocations and prices. In particular, for islands i in group 1, equilibrium output is given by

$$\log y_{i,0} = \kappa_{10} + \kappa_{1,a1} a_1 + \kappa_{1,s} s_1 \tag{16}$$

for some scalars κ_{10}, $\kappa_{1,a1}$, $\kappa_{1,s} > 0$. For islands in group 2, on the other hand, equilibrium output is given by

$$\log y_{i,0} = \kappa_{20} + \kappa_{2,a1} a_1 + \kappa_{2,a2} a_2 + \kappa_{2,s} s_2 \tag{17}$$

for some scalars κ_{20}, $\kappa_{2,a2}$, $\kappa_{2,a1} > 0$, and $\kappa_{2,s} > 0$. We establish this formally in the appendix, but the underlying intuition is quite simple. Because islands from group 1 expect to trade with islands from group 2 , they need to predict their likely terms of trade. Because s_1 is informative about the productivity of group 2, it can help group-1 islands to forecast the likely level of output of group-2 islands, and thereby to forecast their likely terms of trade. As a result, group 1 islands find it optimal to condition their employment and output choices on the signal s_1. But now group-2 islands can use s_2 to help forecast the level of output in group 1 and thereby their own likely terms of trade. It follows that group-2 islands find it optimal to condition their choices on s_2 despite the fact that s_2 is uninformative about the fundamentals of the economy. And as they do so, their employment, output, and investment choices, and equilibrium prices, become sensitive to variation in ε_2, which in turn establishes that the economy features sunspot fluctuations at $t = 0$.

In a certain sense, the variable s_2 ends up serving as a coordination device: it helps some agents forecast the likely level of output in other parts of the economy—and thereby the likely prices they may face in future trading opportunities—despite the fact that it is uninformative about the underlying economic fundamentals. Importantly, though, this happens in an economy where the equilibrium is unique. This coordination device is thus very different from the familiar coordination—or correlation—devices featured in economies with multiple equilibria.

It is worth highlighting two more aspects of our sunspot fluctuations. First, they *have* to occur in equilibrium. This is in contrast to models with multiple equilibria, where sunspot volatility is possible but does not have to take place. Second, sunspot volatility can impact even agents that are fully informed about the fundamentals. In fact, this is precisely the case in the example considered above: islands in group 2 are fully informed about the fundamentals;and yet they are precisely the ones that respond to s_2, thus causing the economy to feature sunspot fluctuations. In this sense, the agents with the highest level of rationality and the most precise information about the fundamentals may well be the ones whose behavior appears as the most arbitrary in the eyes of the outside observer.

We now complete our theorem by establishing that these phenomena are not possible when information is symmetric. Towards this goal, let us rewrite condition (15) as

$$\log y_{i,t} = (1-\hat{\alpha})f_i + \hat{\alpha}\,\mathbf{E}_{i,t}\left[\log y_{j,t}\right]$$

where $f_i \equiv \frac{1}{1-\hat{\vartheta}}\log A_i$. Iterating this condition gives the equilibrium output of an island as a function of the local hierarchy of beliefs:

$$\log y_{i,t} = (1-\hat{\alpha})\left\{f_i + \hat{\alpha}\mathbf{E}_{i,t}[f_j] + \hat{\alpha}^2\mathbf{E}_{i,t}[\mathbf{E}_{j,t}[f_k]] + ...\right\} \tag{18}$$

Because we have assumed that each island knows its own productivity, symmetric information is synonymous to perfect information about the fundamentals. Along with the fact that each island i has an equal probability of being matched with any other island j in the economy, we have $\mathbf{E}_{i,t}[f_j] = \bar{f}$, where $\bar{f} \equiv \frac{1}{\eta}\log\sum_f \exp(\eta f)\phi(f|\theta)$ and where $\phi(f|\theta)$ is the cross-sectional distribution of productivities. It then follows that $\mathbf{E}_{i,t}[\mathbf{E}_{j,t}[f_k]] = \bar{f}$, and so on. Condition (18) thus reduces to

$$\log y_{i,t} = (1-\hat{\alpha})f_i + \hat{\alpha}\bar{f}$$

which pins down the equilibrium level of output in an island as a function of the local productivity and a certain aggregator of the productivities of other islands.

Note, however, that the result extends even in situations where symmetric information is not synonymous to perfect information. As long as information is symmetric, so that first-order beliefs are commonly known, the entire hierarchy of beliefs collapses to first-order beliefs, guaranteeing that equilibrium choices are pinned down by first-order beliefs. This in turn explains why standard macroeconomic models have failed to capture the type of phenomena that we instead captured above by permitting information to be asymmetric.

This argument also suggests a certain reinterpretation of our sunspot fluctuations. By relaxing the assumption of symmetric information, we have permitted higher-order beliefs to be different from first-order beliefs. In so doing, we have also allowed the possibility that there are shocks that move higher-order beliefs without necessarily moving first-order beliefs. It is indeed easy to check that this is precisely what the shock ε_2 does: it moves second- and higher-order beliefs without moving first-order beliefs. It follows that the our sunspot fluctuations can be understood as independent variation in higher-order beliefs.

This in turn brings our contribution closer to that of Morris and Shin (2002) and to other subsequent work that has highlighted other important, and complementary, aspects of higher-order beliefs, such as the anchoring effect of public news (e.g., Morris and Shin, 2002) or the inertia in their response to innovations in fundamentals (e.g., Woodford, 2003, Angeletos and La'O, 2009). This earlier work, however, has restricted attention to settings in which all the signals that the agents receive are correlated with the fundamentals—some of these signals may be private, some of them may be public, but they are all correlated with the fundamentals. In so doing, this work has also imposed that equilibrium outcomes are driven only by shocks that move the agents' beliefs about fundamentals. To the best of our knowledge, the same comment applies to the entire macroeconomics literature on informational frictions.[9] This in turn explains why this literature, which has allowed for asymmetric information, has nevertheless ruled out the sunspot fluctuations we document here.

These points taken, we do not wish the reader to be lost in the wilderness of higher-order beliefs. Rather, we would invite the reader to see this discussion only as one of the different ways of understanding the nature of our sunspot fluctuations. After all, neither the notion of a competitive equilibrium in a Walrasian economy nor the notion of a perfect Bayesian equilibrium in a game require the agents to work through their hierarchy of beliefs; agents only need to make conjectures about the statistical behavior of prices, or about the strategies of other agents, that end up been validated by their collective behavior. In this respect, our contribution is simply to show how such self-consistent conjectures can open the door to a novel type of sunspot fluctuations when, and only when, information is asymmetric.

Similarly, we do not wish the reader to take the signal s_2 that we introduced in the aforementioned example too literally. For us, this is simply a modeling device that permits us to capture a distinct form of uncertainty: as long as information is dispersed and communication is imperfect, agents can face significant uncertainty about economic activity and markets prices *beyond* the uncertainty they might face about the underlying fundamentals. In fact, as the the above example highlighted, some agents may face this distinct form of uncertainty even if they face no uncertainty about the fundamentals. It is this distinct, self-fulfilling type of uncertainty that the pertinent macroeconomics literature has missed; and it is this that our analysis helps capture within otherwise conventional, unique-equilibrium, macroeconomic models.

[9] This includes Lucas (1972), Barro (1976), King (1982), Townsend (1983), Mankiw and Reis (2002), Sims (2003), Woodford (2003), Hellwig (2005), Angeletos and Pavan (2007), Mackowiak and Wiederholt (2009), Lorenzoni (2010), and our own work in Angeletos and La'O (2009).

5 Self-fulfilling waves of optimism and pessimism

In the example studied in the previous section, sunspot fluctuations lasted only one period; this is because, in that example, all agents become fully informed after the first round of trade. We now consider a richer example, one that allows communication to take place more slowly. This example generates fluctuations whose dynamics are (qualitatively) closer to those in the data. It also helps explains how, perhaps paradoxically, communication may intensify the sunspot fluctuations of the economy before it eventually forces them to cease. As a result, the economy can feature seemingly arbitrary waves of optimism and pessimism that keep building up for a while before they eventually fade away. What is more, certain aspects of the equilibrium end up resembling contagion effects, or the spread of fads and rumors. All in all, this example also helps draw a closer connection to many informal descriptions of short-run phenomena, while at the same time offering an entirely novel formal explanation of these phenomena.

Specification. The fundamental θ is given, once again, by a pair of variables a_1 and a_2. But now both of these variables are random, drawn from independent Normal distributions, whose means are zero. The noise ξ, on the other hand, is given by a triple of random variables $\varepsilon_1, \varepsilon_1$ and u. These variables are also Normally distributed and they are independent of one another and of the fundamental. The aggregate state is thus given by the vector

$$z = (\theta, \xi) = (a_1, a_2; \varepsilon_1, \varepsilon_1, u)$$

and it is distributed Normal with mean zero and a diagonal variance-covariance matrix Σ.

Conditional on any particular realization of z, the initial types of the islands are determined as follows. With probability 1/2, an island receives productivity $A_i = \exp(a_1)$; with the remaining probability, it receives $A_i = \exp(a_2)$. Thus, half the island have log-productivity a_1 and half have log-productivity a_1. Let us refer to the former half as "group 1" and the second one as "group 2". Now consider the islands in group 1. Some of these islands get to observe only their own productivities; we refer to them as "uninformed". Some other islands get to see also the following two signals:

$$x_1 = a_2 + \varepsilon_1 \quad \text{and} \quad s_1 = \varepsilon_2 + u;$$

were refer to them as "partially informed". Finally, the remaining islands get to observe the entire state z; were refer to them as "fully informed". Similarly, group 2 is split between uninformed, who only observe a_2; partially informed, who also observe signals

$$x_2 = a_1 + \varepsilon_2 \quad \text{and} \quad s_2 = \varepsilon_1 + u;$$

and fully informed, who observe the entire z. Note that this structures create a "information ladder" within each of the two productivity groups: the uninformed are on the bottom of this ladder, the partially informed are in the middle; and the fully informed are on the top.

The pattern of random matching and trading 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 one step up the ladder. Communication thus takes the form of moving up the ladder. Eventually all islands get to the top, but this may take time, and the ones on the bottom have first to go through the intermediate steps before they can reach the top. This is meant to capture how agents may slowly acquire more and more information about the state of the economy, or otherwise refine their beliefs about both the underlying fundamentals and the likely level of economic activity.

More specifically, we make the following assumptions, which help us maintain tractability. First, an uninformed island can meet either a similarly uninformed island from its own productivity group, in which case it learns nothing, or a partially informed one from it own productivity group, in which it learns the latter's information and hence moves up one step on the ladder. Second, a partially informed island, in turn, can meet either an uninformed one from its own productivity group, in which case it learns nothing itself, or a partially informed one from the *other* productivity group, in which case both learn the aggregate state z and move up their respective information ladders. Third, a fully informed island can only meet with a fully informed from its own productivity group; of course the match then leads to no communication, for these islands already know the entire state. Finally, each island knows beforehand whether it will be matched with an island that is equally or differentially informed; this knowledge is encoded in an idiosyncratic random variable that signals to the island the likely identity of its trading partner.

Characterization. Characterizing the equilibrium dynamics requires keeping track of the cross-sectional distribution of types. This is akin to what happens in models with incomplete markets (e.g., Krusell and Smith, 1987). These is of course an important conceptual difference: the idiosyncratic risk that is crucial for our analysis is the one associated with heterogeneous beliefs about the state of the economy, not the one associated with wealth inequality. Nevertheless, the computational challenge is quite similar in general.

It is in this respect that the specific assumptions we have made above help us maintain tractability. In particular, these assumptions guarantee that the support of this distribution is finite and stays constant over time. Indeed, for any given realization of the aggregate state z, for any given period t, and for any given history up to that point, the type $\omega_{i,t}$ of any given island i can take either of the following ten values:

- ω_{U1} and ω_{U2} are uninformed types—of productivity, respectively, a_1 and a_2—that will be matched in stage 2 with an island from their own subgroup and hence remain uninformed;

- ω_{U1+} and ω_{U2+} are uninformed islands that will be matched with a partially informed island

and hence move up the ladder;

- ω_{P1} and ω_{P2} are partially informed islands that will be matched with an island from their own subgroup and hence remain partially informed;
- ω_{P1+} and ω_{P2+} are partially informed islands that will be matched with a fully informed island and hence move up the ladder; and
- ω_{F1} and ω_{F2} are fully informed that will be matched with an island from their own subgroup.

The dynamics of the cross-sectional distribution of types can thus be summarized in a simple law of motion for a vector $\mu_t \in [0,1]^{10}$; this vector measures the fraction of islands in the economy that, as of stage 1 of period t, take each of the aforementioned ten type values.

We can then establish the following result.

Proposition 4. *There exist positive coefficients (ϕ_a, ϕ_x, ϕ_s) and a 10-by-10 matrix M such that, for any realization of the aggregate state z, the following properties hold:*

(i) For any island i in any period t, the equilibrium level of output is given by

$$\log y_{i,t} = \begin{cases} \phi_a a_1 + \phi_x x_1 + \phi_s s_1 & \textit{if } \omega_{i,t} = \omega_{P1}, \\ \phi_a a_2 + \phi_x x_2 + \phi_s s_2 & \textit{if } \omega_{i,t} = \omega_{P2}, \\ \phi_a a_i & \textit{otherwise} \end{cases}$$

(ii) The cross-sectional distribution of types follows

$$\mu_{t+1} = M\mu_t$$

Part (i) gives a closed-form solution of the equilibrium level of output in each island as a log-linear function of its productivity and of any other signals it may observe, while part (ii) gives the equilibrium law of motion for the distribution of types in the cross-section of islands. Together, these results permit us to determine any equilibrium outcome at either the island-wide or the economy-wide level.

Note then that, by construction, the signals (s_1, s_2) are completely uncorrelated with the productivities (a_1, a_2) and hence do not affect the agents' beliefs about the fundamentals; the latter are pinned down merely by the true productivities (a_1, a_2) and the productivity signals (x_1, x_2). Yet, islands of type ω_{P1+} or ω_{P2+} find it optimal to respond to the sunspot-like signals (s_1, s_2) along the unique equilibrium of the economy. This is because these signal end up in equilibrium be informative of the economic activity of their likely trading partners, despite the fact that they are uninformative of the underlying fundamentals. As a result, a positive innovation in u, which other things equal increases both s_1 and s_2, causes a self-fulfilling boom without moving either the

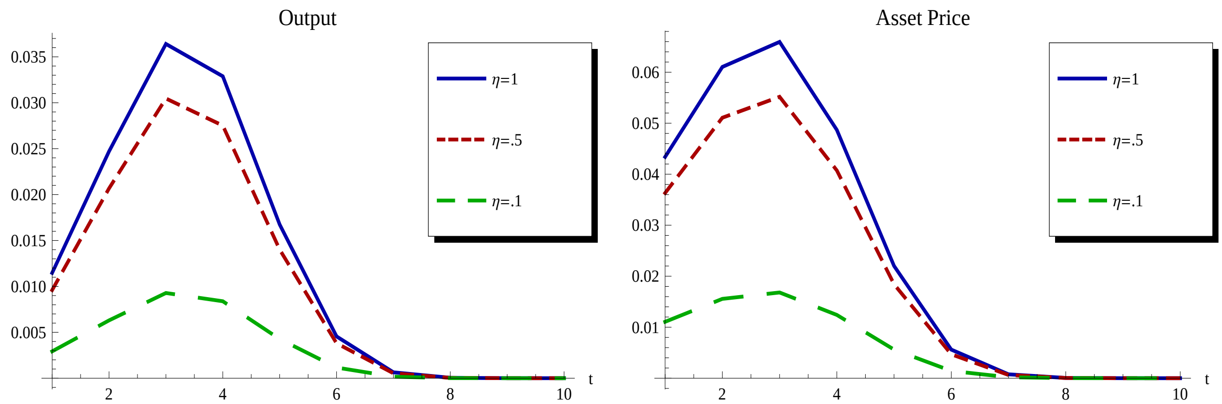

Figure 1: Waves of Optimism. Impulse responses of output Y_t and asset prices Q_t to a positive sentiment shock u, for different values η of the strength of trade interdependence.

fundamentals or any of the agents' beliefs about the fundamentals. Similarly, a negative innovation in u causes a self-fulfilling recession. We henceforth refer to the shock u as a "sentiment shock".

Numerical example. We illustrate the resulting sunspot fluctuations in Figures 1 and 2 for the case of a positive sentiment shock, that is, a positive innovation in u. We focus on the impulse responses of output and asset prices at the aggregate level, which we measure with, respectively, $\log Y_t \equiv \int \log y_{i,t} di$ and $\log Q_{i,t} \equiv \int \log Q_{i,t} di$; the responses of employment and investment are qualitatively similar. The baseline parameterization we use to generate these impulse responses is as follows. We set $\beta = .99$, $\vartheta_n = .6$, $\vartheta_k = .3$, and $\psi = .9$ (which means that we can think of a period as a quarter and that 10% of income is pure profit, or the return to "land"); we let the variances of all the exogenous shocks to equal 1; and we assume that 30% for the population becomes partially informed at $t = 0$, while the rest are uninformed. Finally, we consider three different values for η, namely $\eta \in \{.1, .5, 1\}$, in order to illustrate how the propensity to trade impacts the magnitude of our sunspot fluctuations.

As evident in Figure 1, a positive innovation in u has a positive and persistent effect on aggregate output and asset prices. Moreover, the magnitude of the resulting fluctuations is directly related to the strength of trading interactions, as measured by η. As anticipated, this is because η controls the power of strategic interdependence and thereby the impact of higher-order uncertainty: the higher η, the more sensitive an agent's terms of trades to other agents' activity, and hence the wider the room for self-fulfilling fluctuations. In this respect, our theory also suggests that the processes of economic development, globalization, and financial sophistication, which appear to intensify the level of specialization and trading interdependence in the economy, may well also intensify the self-fulfilling nature of short-run fluctuations.

What is more, the responses of output and asset prices to the shock u are hump-shaped: they

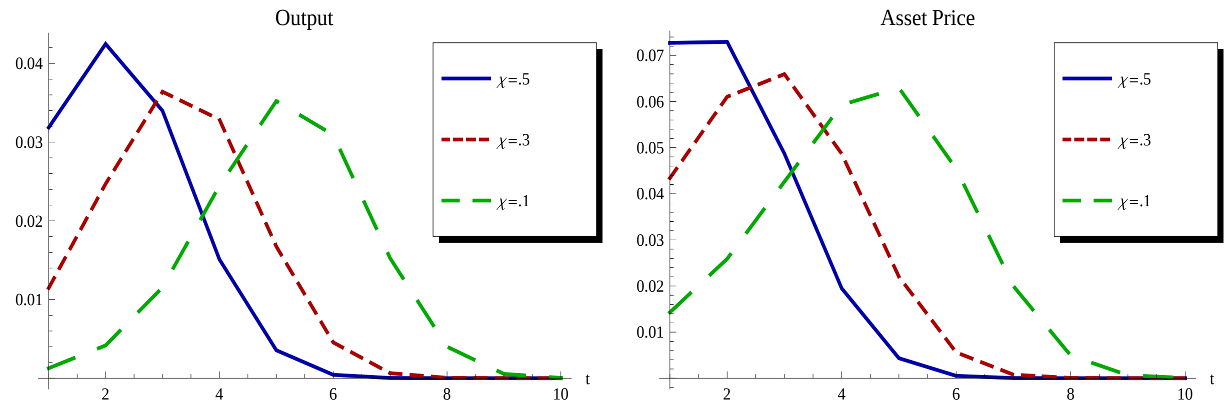

Figure 2: Waves of Optimism. Impulse responses of output Y_t and asset prices Q_t to a positive sentiment shock u, for different values χ of the initial size of "exuberant" agents.

pick few periods after the innovation has taken place. As a result, the resulting fluctuations take the form of self-fulfilling waves of "optimism" and "pessimism"—waves that keep building force for a while before they eventually fade off.

Hump-shaped fluctuations appear to be important features of the data, at least as seen through the lenses of structural VAR models. Here, they obtain quite naturally because of the asymmetry and the endogenous diffusion of information. In the early stages of trade and communication, relative few islands have observed the signals s_1 and s_2 through which the sentiment shock u enters the economy. As these islands meet with uninformed ones, the latter become aware of these sunspot-like signals and themselves start reacting to them. Eventually, however, all islands reach full information and hence the impact of the sentiment shock vanishes.

As evident in Figure 2, the life span of these waves, and the timing at which they peak, is directly related to the fraction of the population that get to observe the sunspot-like signals at $t = 0$. In particular, this figure fixes $\eta = 1$ and considers three values for the fraction of islands that become partially informed (or "exuberant") at $t = 0$, namely $\chi \in \{0.1, .3, .5\}$. We then see that a lower χ tends to push the peak of the wave further in the future. This is due to the fact that, if a lower proportion of the population starts out partially informed, other agents will meet with them only with relatively low probability, and hence the "exuberance" will spread in the economy more slowly. But this also decreases the speed with which agents climb the information ladder, and hence the time that it takes for most of the population to become fully informed. It follows that the boom is likely to persist for longer time, and the bust may take more time to happen.

Finally, it is worth highlighting that, during these self-fulfilling waves, some of the agents that "ride" these waves may have a very good sense that these waves are disconnected from fundamentals. To see this more clearly, suppose that at least some of the fully informed agents faced a positive probability of meeting/trading with a partially informed agent; the latter, in turn, faces uncertainty on whether they will meet a partially or a fully informed agent. In this case, the partially informed agents will continue reacting to the sunspot-like signals for the same reasons as before. But now the fully informed agents will also react to them, because they help them predict the economic activity of some of their likely trading partners. As these fully informed agents start doing so, any other agents who is likely to meet with them, whether partially or fully informed, has an increased incentive to react to the sentiment shock—which, once again, underscores the self-fulfilling nature of fluctuations in our framework. But then note that the economy has agents who are fully informed about the fundamentals, and are thus perfectly confident that the waves could not possibly be justified on the basis of fundamentals, and yet find it in their best interest to ride these waves.

Interpretation and discussion. In more informal terms, we could describe what is going on in our economy after a positive innovation in u as follows. In the beginning, relatively few people have started becoming "exuberant" about the economy. As these people start telling "stories" to other people, or trade with other people, the latter become "infected" by the exuberance of the former, and the optimism spreads out in the economy in a manner that is akin to contagion effects or the spread of fads and fashion. However, as time passes, people start getting more and more into their "senses", realizing that the optimism was not justified by fundamentals. The "irrational exuberance" thus comes and goes, first causing an boom and then a bust.

Clearly, this interpretation of what's going on in our economy is very close to the informal, descriptive explanations of "bubbles" and other short-run fluctuations that often dominate the media and the public arena. It is also very close to the type of arguments one can find in Akerlof and Shiller (2008), Krugman (2009), Shiller (2005). And yet, what lies beneath the surface is very different from the explanation that market pundits and some economists would give to this type of phenomena. According to our theory, the seemingly arbitrary waves of optimism and pessimism, and the associated booms and busts, have nothing to do with irrationality. Rather, they are merely symptoms of the fully rational response of economic agents to the uncertainty they face about economic activity *beyond* the uncertainty they face about the fundamentals. In fact, what would have been truly irrational would be for the agents to ignore this uncertainty and act merely on the basis of fundamentals—doing so would be a symptom of stupidity, not smartness!

At the same time, this is all very different from the perspective that standard macroeconomic models take towards this type of phenomena. The informal arguments about exuberance, self-fulfilling booms and busts, and the like make no sense whatsoever within the context of these models, because all fluctuations are ultimately driven by technology shocks or other fundamentals.[10]

[10]Clearly, this statement does not apply to models that attempt to formalize these notions with multiple equilibria; but these models tend to be a minority in the literature and, as noted earlier, face their own challenges—especially

However, that does not mean that the arguments are completely baseless. Rather, part of the fault is also at the side of modern macroeconomic models. As we have shown here, these models have failed to accommodate these notions only because they have made the convenient, but unrealistic, assumption of perfect communication.

Our self-fulfilling fluctuations show up, not only in real economic activity, but also asset prices. In the specific model of this paper, this is only because asset prices depend on future real economic activity. More broadly, however, our results indicate how the asymmetry of information can open the door to seemingly arbitrary fluctuations in asset markets even when the equilibrium is unique and there are not any noise traders or other irrational forces operating in these markets. Embedding this insight in richer asset-pricing modelsseems an intriguing direction for future research.

Finally, it is obvious that the model of this paper was not designed to address quantitative questions. Needless to say, moving in this direction is both important and challenging. We make a (small) step towards this direction in a companion paper, Angeletos and La'O (2009), by embedding dispersed information about technology shocks in the canonical Dixit-Stiglitz framework that underlies modern quantitative macroeconomic models. That paper does not document the sunspot fluctuations that are at the core of the present paper. Instead, like the rest of the pertinent literature, it focuses on fluctuations that are driven either by fundamentals (technology shocks) or by noisy news about the fundamentals. Nevertheless, it establishes a result that is highly complementary to the contribution of the present paper: with dispersed information, most of the business cycle may be dominated by noise even if the agents are well informed about the underlying fundamentals. Of course, that noise originates from news about the fundamentals, thus failing to capture sunspot volatility. However, in the light of the results of the present paper, one could reinterpret the noise-driven fluctuations in our companion paper as a convenient proxy for the sunspot fluctuations we document here. Either way, the combination of the results of the two papers highlight the distinct implications that the asymmetry of information may have for short-run fluctuations.

6 Tradeable numeraire and risk sharing

We now consider the richer version of our model, which allows any two islands that get matched together to trade the numeraire good alongside the specialized goods. Coupled with the assumption that preferences are linear in the numeraire, this guarantees that there is, in effect, complete risk sharing: any idiosyncratic income risk gets absorbed by the consumption of the numeraire, leaving the equilibrium terms of trade, the equilibrium allocation of all the other goods, and the agent's marginal utilities of wealth unaffected. This has only secondary effects on our positive results but,

when it comes to policy analysis.

for obvious reasons, plays an important role once we turn attention to normative questions.

The definition of the equilibrium is the same as before, except that now trades of differentiated goods can be balanced by trades of the numeraire. Accordingly, condition (iii) is dropped from Definition 1 and, instead, all relative prices are equalized across any two trading islands: for any two islands i and j that meet in stage 2 of period t, it must be that $p_{it} = p^*_{jt}$ and $p^*_{it} = p_{jt}$.

The analysis of the equilibrium can then proceed in similar steps as before. The prices and the equilibrium allocation of the differentiated goods are pinned down by the following conditions:

$$U_c\left(c_{i,t}, c^*_{i,t}\right) = p_{i,t} = p^*_{j,t} = U_{c^*}\left(c_{j,t}, c^*_{j,t}\right)$$

$$c_{i,t} + c^*_{j,t} = y_{i,t} \quad c_{j,t} + c^*_{i,t} = y_{j,t}$$

These conditions can be solved for the equilibrium prices as functions of the output levels of the differentiated goods:

$$p_{i,t} = P(y_{i,t}, y_{j,t}) \tag{19}$$

for some function $P : \mathbb{R}^2_+ \to \mathbb{R}_+$. This condition is the analogue of condition (6) in the baseline model. The only difficulty is that P may not have a closed-form solution.[11] Modulus this difference, the rest of the equilibrium characterization is identical to the one in the baseline model. We thus reach the following variants of Proposition 2 and 3.

Proposition 5. *The equilibrium exists and is unique. Furthermore, there exists a monotone function $G : \mathbb{R}_+ \to \mathbb{R}$ such that, for any $t \geq 1$, the period-t equilibrium levels of output solve the following fixed-point problem:*

$$\mathbb{E}\left[G\left(y(\omega_{i,t}), A(\omega_i), y(\omega_{j,t})\right) \middle| \, \omega_{i,t}\right] = 0 \tag{20}$$

Condition (20) is the analogue of condition (15) in the baseline model: it gives the equilibrium level of output in an island as a function of the local productivity and the local beliefs about the likely output level of other islands. As in the baseline model, the origin of the dependence on the beliefs of the output of other islands is the dependence of the terms of trade. Before this dependence was captured in an explicit way by condition (6); now it is captured in an implicit way by condition (19). Either way, however, the underlying economics are essentially the same.

[11] This is with one exception: the special case where there is no "home" or "foreign" bias in preferences. Formally, this is the case when $U(c, c^*) = U(c^*, c)$ for all c, c^*. In this case, the two goods are equally split across the two islands and hence the equilibrium price function is simply given by $P(y, y^*) \equiv U_c(y/2, y^*/2)$. Clearly, the latter is decreasing in its first argument and, as long as $U_{c,c^*} > 0$, it is increasing in its second argument. One can show that this property generalizes beyond this special case. It follows that a key qualitative property of the baseline model continues to hold in the variant model of this section: the price of the local differentiated good increases with the level of output (or "demand") of the trading partner.

It is then also immediate that the core of our positive results remains unaffected: as long as communication is imperfect, the equilibrium can feature essentially the same type of sunspot fluctuations as the one in the baseline model. This is evident in the following special case.

Proposition 6. *Suppose* $U(c, c^*) = (c^{1/2}c^{*1/2})^\gamma$ *and* $V(n) = n^\epsilon$, *where* $\gamma < 1 < \epsilon$. *Then, there exist scalars* $\varphi > 0$ *and* $\alpha \in (0, 1)$ *such that*

$$\log y(\omega_{i,t}) = (1 - \alpha)\varphi \log A(\omega_i) + \alpha \mathbf{E}\left[\log y(\omega_{j,t})|\, \omega_{i,t}\right]$$

Modulus the precise values of the scalars φ and α, this is identical to Proposition 2. It follows that the self-fulfilling waves of optimism and pessimism that we documented in the previous section can be re-casted *exactly* in this special case. Away from this special case, we lose the tractability afforded by the baseline model, but this does not affect the core of our positive results: sunspot fluctuations continue to possible if and only if information is asymmetric.

The only economics that are affected is the precise channel through prices depend on others' output. In the baseline model, part of the reason that $p_{i,t}$ increase with $y_{j,t}$ was a wealth effect: other things equal, a higher $y_{j,t}$ means that i's trading partner is richer, which raises her demand for i's good. Here, instead, such wealth effects play no role. Rather, the only reason that $p_{i,t}$ increase with $y_{j,t}$ is the complementarity of the goods in preferences (namely $U_{c,c^*} > 0$).[12]

From an empirical perspective, such wealth effects can be quite appealing. We nevertheless opt to abstract from them in order to insulate the policy analysis that follows in the subsequent two sections from considerations regarding social insurance and redistribution.

7 Policies that can "tame" animal spirits

The preceding analysis has documented how dispersed information and imperfect communication can open the door to a novel type of sunspot fluctuations, one that is akin to those associated with the informal notions of "animal spirits", "market psychology," and the like. This pegs two policy

[12]To expand on what we mean by these wealth effects, consider a variation of the model that imposes exogenous endowments for the differentiated goods. When the numeraire is not traded, any redistribution of the given aggregate endowment of the differentiated goods across the two islands impacts equilibrium prices precisely because it induces wealth effects. When, instead, the numeraire is traded, any redistribution of the given aggregate endowments of the differentiated goods is offset by variation in the consumption of the numeraire, thus leaving the prices and the consumption allocation of the differentiated goods completely unaffected—these prices and consumption allocations depend only on the total endowments of the two goods, not the distribution of these endowments across the islands. Of course, in our model neither the aggregate endowment of the differentiated goods nor its distribution is exogenous; they are both determined by employment and investment choices at earlier stages of economic activity. As a result, differences in productivities and/or information induce distributional effects whenever the numeraire is not traded. It is these effects that vanish once we let the numeraire be traded.

questions, one positive and one normative. The first is whether the government could do something to stabilize the economy against these fluctuations—or to "tame" the market's animal spirits. The second is whether such policy interventions would be socially desirable. We address the positive question in this section and the normative one in the next section.

For the positive exercise, we consider is a linear, state-contingent tax (or subsidy) on firm profit. Alternatively, we could consider a linear, state-contingent tax on the labor and capital income of the households. These taxes are collected at stage 2 of each period, are allowed to be specific to each pair of islands that trade in that stage, and are allowed to be contingent on any information that is commonly known to these islands at that stage—but cannot depend on any information that would not otherwise be available to them. In this respect, the policy instruments we consider do not require any informational advantage on the side of the government and the resulting allocations could be implemented as a part of an incentive-compatible mechanism that constrains the planner under the same communication constraints as the market.[13]

In particular, we let the period-t tax rate faced by island i be given by

$$\tau_{i,t} = T(A_i, A_j, y_{i,t}, y_{j,t}), \tag{21}$$

where $j = m(i,t)$ is the island that gets matched with island i, A_i and A_j are their productivities, $y_{i,t}$ and $y_{j,t}$ are their output levels, and T is a function from $\mathbb{R}^4_+$ to $\mathbb{R}$. These taxes thus depend only on information that is commonly known to islands i and j at stage 2 of period t. Finally, any taxes collected are assumed to be distributed back to the agents in these islands in a lump sum fashion.[14]

We can then show the following result.

Proposition 7. *There exists a policy as in (21) such that the following are true:*

(i) The equilibrium is unique.

(ii) There is no room for sunspot fluctuations: equilibrium allocations and prices depend only on actual and expected fundamentals.

The intuition for this result is simple. The only reason that our economy features sunspot fluctuations is that agents care to predict one another's economic activity so as to predict their future terms of trades. By appropriately choosing the design of the state-contingent taxes in (21), the government can guarantee that the after-tax terms of trade depend only on the fundamentals, not

[13] We make this idea precise in the next section.

[14] We allow lump-sum transfers for two reasons. First, there is no good reason in our model that would rule them out or make them undesirable; note, in particular, that households are risk-neutral, leaving no room for distributional concerns. And second, we prefer not to confuse the incentive effects of taxes that we wish to highlight in this section with the restrictions imposed by the unavailability of lump-sum transfers.

on the activity of other agents.[15] In so doing, the government can also guarantee that agents cease to care to predict the endogenous level of economic activity—in order to predict their after-tax terms of trade, they only need to predict the fundamentals. But then agents have no more incentive to react to information about each other's economic activity and the entire sunspot volatility vanishes.

Of course, the specific taxes utilized in this result might be hard to implement in practice, because they require information about who trades with whom, about the choices of these agents, and about their fundamentals.[16] However, one should not take this too literally. Rather, the broader insight here is that taxes, or other policies, that depend on both signals of the exogenous fundamentals and signals of the endogenous level of economic activity can induce the agents to care more about the fundamentals and less about the choices of other agents—and in so doing to dampen any self-fulfilling aspects of short-run fluctuations. This insight is borrowed from Angeletos and Pavan (2009), who show within an abstract but flexible framework how such state-contingent taxes can help the government control the equilibrium effects of higher-order uncertainty.

More concretely, in our business-cycle context this insight means that the government should promise to tax (respectively, subsidize) economic activity whenever there is a boom (respectively, recession) that is not justifiable on the basis of ex post public information about the underlying fundamentals. These taxes would reduce the incentives of agents to "ride" any waves of optimism or pessimism that are not driven by fundamentals. And if one considers a monetary variant of our model, one may mimic the incentive effects of such tax instruments with more realistic monetary-policy instruments. In either case, our analysis suggests that conventional stabilizations policies may give the government the power to "tame" the market's animal spirits.

At this point, it is worth highlighting how this policy lesson contrasts with the one provided by previous formalizations of animal spirits. In models with multiple equilibria, the response of the economy to many policies—including conventional fiscal or monetary policies—is often indeterminate. As a result, policy analysis often rests on ad hoc assumptions about equilibrium selections. In models with irrational agents, on the other hand, there is little discipline on how these agents may react to any particular policy intervention: depending on what one assumes about the response of these agents to policy changes, one could pretty much justify any policy. This is most eloquently highlighted in N. Gregory Mankiw's reaction to Robert Shiller's claim that a fiscal stimulus would—by sheer magic—improve confidence in the economy:

[15] Indeed, it suffices to choose $T(.)$ so that, in equilibrium, the after-tax sale price of the commodity produced by island i, namely $(1-\tau_{i,t})p_{i,t}$, depends at most on A_i and A_j.

[16] This information is no more than what, in our model, is already available to the agents that trade, and in this sense they do not require any informational advantage on the side of the government. Nevertheless, this information may still be a lot to ask in practice.

"Yale's Bob Shiller argues that confidence is the key to getting the economy back on track. I think a lot of economists would agree with that. The question is what it would take to make people more confident. Bob thinks that confidence would rise if the government borrowed more and spent more. Other economists think that confidence would rise if the government committed itself to, say, lower taxes on capital income. The sad truth is that we economists don't know very much about what drives the animal spirits of economic participants. Until we figure it out, it is best to be suspicious of any policy whose benefits are supposed to work through the amorphous channel of 'confidence'."

To recap, prior formalizations of animal spirits pose serious difficulties for policy analysis and provide little guidance, if any, on how conventional stabilization policies could help "tame" the market's animal spirits. In sharp contrast, our approach permits these notions to be accommodated within otherwise conventional policy exercises without the need to rely either on ad hoc equilibrium selections or on convenient but arbitrary assumptions regarding the "amorphous channel of confidence".

8 Constrained efficiency

We now turn to the normative question of whether the policies identified in the previous section, or any other stabilization policies that might try to "tame" animal spirits, are socially desirable.

Understanding the welfare properties of the equilibrium, and thereby the desirability of policy intervention, requires the definition of an appropriate efficiency concept. One possibility is to compare equilibrium outcomes with first-best outcomes. Note, however, that implementing the first-best outcomes requires perfect communication.[17] The first-best efficiency concept would thus be inconsistent with taking the communication frictions as part of the exogenous primitives of the economy and would effectively assume that the government has an informational advantage vis-a-vis the market mechanism. What is more, comparing the equilibrium under imperfect communication with the first-best allocation—or equivalently with the equilibrium under perfect communication—would help us understand the welfare costs of imperfect communication, but would not address the policy question that we are interested in here. For conventional stabilization policies are meant to affect budget sets and incentives, not the level of communication.

[17] Indeed, for the first best outcomes to obtain, it is necessary that the agents either talk to each other up to the point that they reach common knowledge about the underlying fundamentals or that otherwise achieve essentially the same goal. For example, the agents could first communicate all their information to a center. This center could then aggregate this information and communicate it back to the agents, at which point the market mechanism would induce the complete-information equilibrium, which is of course first-best efficient. Alternatively, once the center has aggregated all relevant information, it could directly dictate the first-best allocation.

We thus consider a different efficiency concept, one that subjects the planner under the same communication constraints as those imposed on the market mechanism.[18]

Definition 3. *A constrained efficient allocation is an allocation that maximizes ex-ante utility subject to the following two sets of constraints:*

(i) ***Feasibility.*** *The functions $y, n, \tilde{n}, k, c, c^*$ satisfy the following:*

$$y_t(\omega_{i,t}) = A(\omega_i) k_t(\omega_{i,t})^{1-\vartheta} n_t(\omega_{i,t})^{\vartheta} \tag{22}$$

$$k_{t+1}(\omega_{i,t+1}) = \tilde{n}_t(\omega_{i,t+1})^{\psi} \tag{23}$$

$$c_t(\omega_{i,t+1}) + c_t^*(\omega_{j,t+1}) = y(\omega_{i,t}) \tag{24}$$

where j stands for trading partner of island i.

(ii) ***Communication.*** *Information sets evolve according to the following law of motion: if i and j are trading partners during stage 2 of period t,*

$$\omega_{i,t+1} = (\omega_{i,t}, \omega_{j,t}) \quad \text{and} \quad \omega_{j,t+1} = (\omega_{j,t}, \omega_{i,t}). \tag{25}$$

Recall that $P_t(\omega_{i,t})$ denotes the probability that an island has type $\omega_{i,t}$ in stage 1 of period t, while $P_t(\omega_{i,t+1}|\omega_{i,t})$ denotes the probability that that its type transits to $\omega_{i,t+1}$ in stage 2 of that period (and hence also in stage 1 of the next period). We can thus write ex ante utility as in the following statement of the planner's problem.

Planning Problem. *The efficient allocation maximizes*

$$\mathcal{W} = \sum_t \beta^t \left\{ \sum_{\omega_{i,t}} \left[\sum_{\omega_{i,t+1}} \left[U\left(c_t(\omega_{i,t+1}), c_t^*(\omega_{i,t+1})\right) - \tilde{n}_t(\omega_{i,t+1}) \right] P_t(\omega_{i,t+1}|\omega_{i,t}) - n_t(\omega_t) \right] P_t(\omega_{i,t}) \right\}$$

subject to (22)-(25).

Clearly, this is a strictly convex problem; it has a unique solution, which is pinned down by FOCs. Let $\lambda(\omega_{i,t+1}) P(\omega_{i,t+1}|\omega_{i,t}) P(\omega_{i,t})$ to be the Lagrange multiplier on the resource constraint (24). Using (22) and (23) to drop $y_t(\omega_{i,t})$ and $\tilde{n}(\omega_{i,t+1})$, and after some rearrangement, we can write the Lagrangian of this problem as follows:

$$\mathcal{L} = \sum_t \beta^t \left\{ \sum_{\omega_{i,t}} \left[\sum_{\omega_{i,t+1}} (\ldots)\, P_t(\omega_{i,t+1}|\omega_{i,t}) - n_t(\omega_t) - \frac{1}{\beta} k_t(\omega_{i,t})^{\frac{1}{\psi}} \right] P_t(\omega_{i,t}) \right\}$$

[18] See Angeletos and Pavan (2007, 2009) and Vives (2010) for similar concepts within a certain class of linear-quadratic games.

where (...) stands for

$$U\left(c_t(\omega_{i,t+1}), c_t^*(\omega_{i,t+1})\right) + \lambda_t(\omega_{i,t+1})[A(\omega_i)k_t(\omega_{i,t})^{1-\vartheta}n(\omega_{i,t})^{\vartheta} - c_t(\omega_{i,t+1}) - c_t^*(\omega_{j,t+1})]$$

The efficient allocation of consumption is thus pinned down by the following optimality conditions (along with the resource constraints):

$$U_c\left(c_t(\omega_{i,t+1}), c_t^*(\omega_{i,t+1})\right) = \lambda_t(\omega_{i,t+1}) = U_{c^*}\left(c_t(\omega_{j,t+1}), c_t^*(\omega_{j,t+1})\right)$$

Comparing these conditions with the corresponding one for the equilibrium, we see immediately that, for any given output levels, the efficient and equilibrium consumption allocations are the same, and the planner's shadow prices coincide with the market prices. Turning to the efficient employment and investment decisions, these are determined by the following optimality conditions:

$$\begin{aligned} 1 &= \sum_{\omega_{i,t+1}} \lambda_t(\omega_{i,t+1}) P_t(\omega_{i,t+1}|\omega_{i,t}) \vartheta \frac{y_t(\omega_{i,t})}{n_t(\omega_{i,t})} \\ \frac{1}{\beta\psi} k_t(\omega_{i,t})^{1/\psi-1} &= \sum_{\omega_{i,t+1}} \lambda_t(\omega_{i,t+1}) P_t(\omega_{i,t+1}|\omega_{i,t})(1-\vartheta) \frac{y_t(\omega_{i,t})}{k_t(\omega_{i,t})} \end{aligned}$$

Clearly, these are the same conditions as the corresponding ones that characterize the equilibrium, except that the market prices have now been replaced by the planner's shadow prices. But we already argued that shadow and market prices coincide. The following is thus immediate.

Theorem 2. *As long as the numeraire is traded across islands, the competitive equilibrium is constrained efficient.*

This result contains the core normative lesson of our paper. In effect, it establishes a variant of the first welfare theorem for a class of economies that introduce asymmetric information and communication imperfections. As such, this result is an important contribution on its own right.[19]

What is more, this result implies that our sunspot fluctuations are *not* an invitation for stabilization policy: despite the fact that, in our framework, such policies give the government the power to "tame" the market's animal spirits, exercising this power is anything but desirable. This is true no matter how arbitrary, disconnected from fundamentals, and "insane" the resulting fluctuations may appear to either the outside observer or the insiders of our economy.

This lesson stands in sharp contrast, not only to conventional wisdom and old-Keynesian thinking, but also to previous formalizations of "animal spirits". In models with multiple equilibria, animal spirits are modeled a coordination failure at the aggregate level. In models with irrational

[19] It is safe to conjecture that this result extends to more general preferences and technologies provided that we abstract from incomplete risk-sharing, as we have done here. See Angeletos and La'O (2009) for a similar result in a variant economy that is closer to the canonical Dixit-Stiglitz framework.

agents, animals spirits are modeled as a form of "stupidity" at the individual level. Either way, the resulting fluctuations are socially undesirable and the door for government intervention is left wide open. Here, instead, the door is shut down.

Of course, our efficiency result rests on abstracting from incomplete risk sharing. As anticipated in the previous section, we have achieved this in a similar manner as in Lagos and Wright (2005): by assuming that preferences are linear in the numeraire and that the numeraire is tradeable. If we relax these assumptions without otherwise insulating the economy from the resulting idiosyncratic income risk, our efficiency result will almost surely perish. Nonetheless, the inefficiency that emerges in this case has little to do with the presence of animal spirits: incomplete risk-sharing would cause inefficiency even if information had been symmetric and our sunspot fluctuations were absent. What is more, this is almost certainly not the reason that underlies the Keynesian view that the government should try to "tame" the market's animal spirits. Rather, this view appears to be based on the simple reasoning that, as long as fluctuations are disconnected from fundamentals, they are "evidently" undesirable. Our analysis has shown that this informal reasoning, while appearing "obviously true" at first glance, is reduced to "clearly wrong" once subjected to the test of our theory.

To recap, the overall picture that emerges from combining the normative results of this section with the positive results of the previous sections strikes a delicate balance between the two sides of the debate that we mentioned in our introduction. On the one hand, the positive properties of our theory have a strong Keynesian flavor, emphasizing the self-fulfilling nature of short-run fluctuations and formalizing the notions of "animal spirits", "sentiments", "confidence", and "demand shocks". On the other hand, the normative properties of our theory are as Neoclassical as it gets, leaving no room for the policy message advocated by either old-fashioned or new-born Keynesianism.

9 Discussion

We would like to close our paper with a discussion on the empirical and methodological content of our contribution. Towards this goal, we find it useful to consider the following extract from Narayana Kocherlakota's (2010) recent article on the state of macroeconomics.

> *Why does an economy have business cycles? Why do asset prices move around so much? At this stage, macroeconomics has little to offer by way of answers to these questions. The difficulty in macroeconomics is that virtually every variable is endogenous, but the macroeconomy has to be hit by some kind of exogenously specified shocks if the endogenous variables are to move.*
>
> *The sources of disturbances in macroeconomic models are (to my taste) patently unrealistic. Perhaps most famously, most models in macroeconomics rely on some form of*

large quarterly movements in the technological frontier (usually advances, but sometimes not). Some models have collective shocks to workers' willingness to work. Other models have large quarterly shocks to the depreciation rate in the capital stock (in order to generate high asset price volatilities). To my mind, these collective shocks to preferences and technology are problematic. Why should everyone want to work less in the fourth quarter of 2009? What exactly caused a widespread decline in technological efficiency in the 1930s? Macroeconomists use these notions of shocks only as convenient shortcuts to generate the requisite levels of volatility in endogenous variables.

Of course, macroeconomists will always need aggregate shocks of some kind in macro models. However, I believe that they are handicapping themselves by only looking at shocks to fundamentals like preferences and technology. Phenomena like credit market crunches or asset market bubbles rely on self-fulfilling beliefs about what others will do. For example, during an asset market bubble, a given trader is willing to pay more for an asset only because the trader believes that others will pay more. Macroeconomists need to do more to explore models that allow for the possibility of aggregate shocks to these kinds of self-fulfilling beliefs.

The theory that we have developed in this paper has attempted to address precisely the kind of issues that Kocherlakota discusses in an eloquent and effective way. Like him, we are not satisfied with theories that attribute the bulk of short-run fluctuations to high-frequency movements in preferences and technologies. Instead, we have sought to provide an entirely novel theory, one that focuses on the apparent self-fulfilling nature of these fluctuations.

To achieve this, our model, too, had to introduce some exogenous shocks—the ones that we interpreted as a form of sunspots or "sentiment shocks". It may be hard to find a literal analogue of these shocks in the real word—just as it is hard to find a literal analogue for most of the technology, preference or other shocks featured in standard macroeconomic model. However, we think that a literal interpretation of our sentiment shocks would be too narrow. Rather, as anticipated in the Introduction, we would like to see these shocks as modeling devices that permit us to capture a distinct type of uncertainty—the uncertainty that economic agents may face about the endogenous economic activity *beyond* the uncertainty they face about the exogenous fundamentals.

While we have made no attempt in the present paper to quantify this particular type of uncertainty, we believe that it plays a central role in practice. Consider for example the recent recession. During this recession, as in many others, firms appear to cut down on employment and investment because they are worried about future consumer demand. Consumers, in turn, appear to cut down on consumption because they are afraid of future labor-market conditions—if they are not already unemployed. One kind of fear seems to feed the other—a vicious cycle that ultimately seems to determine how deep or long the recession might be.

By formalizing this self-fulfilling nature of short-run fluctuations within an otherwise standard, neoclassical, unique-equilibrium, rational-expectations model of the macroeconomy, our approach also captures what, at least in our eyes, is the core *positive* aspect of the Keynesian view of the world: short-run fluctuations are explained as self-reinforcing waves of optimism and pessimism, or a distinct form of "demand shocks", rather than the response to high-frequency changes in preferences and technologies.

And yet, the policy implications of our analysis are strongly anti-Keynesian. Even if we were to reach a consensus on the view that most short-run fluctuations in asset markets and the macroeconomy are "obviously" disconnected from fundamentals and can only be explained as the product of "animal spirits", this does *not* mean either that these seemingly arbitrary fluctuations are self-evident symptoms of inefficiency or that the government should try to insulate the economy from these fluctuations.

Of course, it would only be naive on our side to argue that the current recession is a symptom of efficiency. Also, it should be clear that the normative results of our paper are certain to be less robust than the positive ones. For example, we already highlighted that the constrained efficiency of the equilibrium ceases to hold when risk-sharing is incomplete. The same may well be true once one introduces credit frictions, nominal rigidities, and other important features that are present in reality but not in our model—extending the analysis in these directions would thus be worthwhile, not only for quantifying the type of uncertainty we have highlighted in this paper, but also for providing more guidance on policy. However, none of this could vindicate some recent calls for a return to old-Keynesian thinking. For, as John Cochrane (2009) puts it in his response to Krugman,

> *"A science that moves forward almost never ends up back where it started. Einstein revises Newton, but does not send you back to Aristotle. At best you can play the fun game of hunting for inspirational quotes, but that doesn't mean much."*

Needless to say, we have played our fair portion of the quotes game in this paper! But we believe that we have also provided a novel perspective on the positive and normative properties of short-run fluctuations—one that seeks to push the frontier of macroeconomic research away from models that attribute the entire fluctuations to technology and preference shocks and towards models that better capture the self-fulfilling nature of fluctuations.

Appendix

Proof of Proposition 1. Using (6) in (9), along with the fact that $V'(n) = 1$, we get

$$n_{i,t} = \vartheta_n \mathbb{E}_{i,t}[y_{j,t}^{\eta}] y_{i,t}^{1-\eta}$$

Solving $y_{i,t} = A_{i,t} k_{i,t}^{\vartheta_k} n_{i,t}^{\vartheta_n}$ for $n_{i,t}$, substituting the solution into the above, and rearranging yields

$$y_{i,t}^{\frac{1}{\vartheta_n} - (1-\eta)} = \left(A_i k_{i,t}^{\vartheta_k} \right)^{\frac{1}{\vartheta_n}} \mathbb{E}_{i,t}[y_{j,t}^{\eta}]$$

Taking the log of both sides of this condition, solving for $\log y_{i,t}$, and using the definition of the operator $\mathbf{E}$ and the more precise notation $y_{i,t} = y(\omega_{i,t})$, $y_{j,t} = y(\omega_{j,t})$ and $A_i = A(\omega_i)$, we reach condition (10). *QED*

Proof of Proposition 2. Substituting the optimality condition for capital (14) into the production function yields

$$y_{i,t} = A_i \left(y_{i,t}^{1-\eta} \mathbb{E}_{i,t}[y_{j,t}^{\eta}] \right)^{\psi \vartheta_k} n_{i,t}^{\vartheta_n}$$

Using this to eliminate $n_{i,t}$ in (9), replacing $p_{i,t}$ in the latter from (6), and rearranging, we get

$$y_{i,t}^{1-(1-\eta)(\vartheta_n + \psi\vartheta_k)} = A_i \left(\mathbb{E}_{i,t}[y_{j,t}^{\eta}] \right)^{\vartheta_n + \psi \vartheta_k}.$$

Taking logs, rearranging, and using the definitions of $\hat{\vartheta}$ and $\hat{\alpha}$, we reach condition (15). *QED*

Proof of Proposition 3. Iterating condition (15) gives output as a function of the hierarchy of beliefs:

$$\log y_{i,t} = (1-\hat{\alpha}) \left\{ f_i + \hat{\alpha} \mathbf{E}_{i,t}[f_j] + \hat{\alpha}^2 \mathbf{E}_{i,t}[\mathbf{E}_{j,t}[f_k]] + ... \right\}, \tag{26}$$

where $f_i \equiv \frac{1}{1-\hat{\vartheta}} \log A_i$. Because the fundamentals are bounded (and this fact is common knowledge), the beliefs of all orders are also bounded. Along with the fact that $\hat{\alpha} \in (0,1)$, this guarantees that the above infinite sum converges, proving the existence and uniqueness of equilibrium. *QED*

Proof of Theorem 1. The main argument towards this Theorem is presented in the main text. The only part that remains to be proved is the pair of conditions (16) and (17), which characterize the equilibrium for the specific example with sunspot fluctuation studied in the main text.

Suppose that, conditional on ω_i, output of the matched island $\log y(\omega_j)$ is log-normal, with variance independent of $\omega_{i,0}$; that this is true under the log-normal structure for the underlying shocks and signals we will prove shortly. Using log-normality of $\log y(\omega_j)$ in condition (15) in Proposition 2, we infer that the equilibrium production strategy of island i must satisfy the following condition

$$\log y(\omega_i) = (1-\hat{\alpha}) \frac{1}{1-\theta} a_i + \hat{\alpha}\, \mathbb{E}_{it} [\log y(\omega_{j,t})] \tag{27}$$

We now guess and verify a log-linear equilibrium under the log-normal specification for the shock and information structure. Suppose the equilibrium production strategies of island 1 and island 2 take the log-linear form given in (16) and (17), respectively, for some coefficients $(\kappa_{10}, \kappa_{1,a1}, \kappa_{1,s})$ and $(\kappa_{20}, \kappa_{2,a1}, \kappa_{2,a2}, \kappa_{2,s})$. It follows that y_1 and y_2 are indeed log-normal, with

$$\begin{aligned} \mathbb{E}\left[\log y_1 | \omega_2\right] &= \kappa_{10} + \kappa_{1,a1} a_1 + \kappa_{1,s}\left(a_2 + \mathbb{E}\left[\varepsilon_1 | \omega_2\right]\right) \\ \mathbb{E}\left[\log y_2 | \omega_1\right] &= \kappa_{20} + \kappa_{2,a1} a_1 + \kappa_{2,a2} \mathbb{E}\left[a_2 | \omega_1\right] + \kappa_{2,s} \mathbb{E}\left[s_2 | \omega_1\right] \end{aligned}$$

where $\mathbb{E}\left[\varepsilon_1 | \omega_2\right] = \frac{1}{2}\left(s_2 - a_2\right)$, $\mathbb{E}\left[a_2 | \omega_1\right] = \frac{1}{2} s_1$, and $\mathbb{E}\left[s_2 | \omega_1\right] = s_1$. Substituting these expressions into (27) gives us

$$\begin{aligned} \log y\left(\omega_2\right) &= (1-\hat{\alpha}) \frac{1}{1-\theta} a_2 + \hat{\alpha}\left[\kappa_{10} + \kappa_{1,a1} a_1 + \kappa_{1,s} a_2 + \kappa_{1,s} \frac{1}{2}\left(s_2 - a_2\right)\right] \\ \log y\left(\omega_1\right) &= (1-\hat{\alpha}) \frac{1}{1-\theta} a_1 + \hat{\alpha}\left[\kappa_{20} + \kappa_{2,a1} a_1 + \kappa_{2,a2} \frac{1}{2} s_1 + \kappa_{2,s} s_1\right] \end{aligned}$$

For this to coincide with the system of equations in (16) and (17) for every $(a_2, \varepsilon_1, \varepsilon_2)$, it is necessary and sufficient that the coefficients $(\kappa_{10}, \kappa_{1,a1}, \kappa_{1,s}, \kappa_{20}, \kappa_{2,a1}, \kappa_{2,a2}, \kappa_{2,s})$ solve the following system:

$$\begin{aligned} \kappa_{2,a1} &= \hat{\alpha} \kappa_{1,a1} \\ \kappa_{2,a2} &= (1-\hat{\alpha}) \frac{1}{1-\theta} + \hat{\alpha} \frac{1}{2} \kappa_{1,s} \\ \kappa_{2,s} &= \hat{\alpha} \kappa_{1,s} \frac{1}{2} \\ \kappa_{1,a1} &= (1-\hat{\alpha}) \frac{1}{1-\theta} + \hat{\alpha} \kappa_{2,a1} \\ \kappa_{1,s} &= \hat{\alpha}\left(\kappa_{2,a2}/2 + \kappa_{2,s}\right) \end{aligned}$$

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

Proof of Proposition 4. *Part (i)* Using the conjectured log-normality of $\log y\left(\omega_j\right)$, we again infer that the equilibrium production strategy of island i must satisfy the best-response condition given in (27). Using this condition, we now consider the equilibrium strategy for each of the ten types $\omega_{i,t}$, by considering the equilibrium for all possible matches.

$$\log y\left(\omega_i\right) = (1-\hat{\alpha}) \frac{1}{1-\theta} a_i + \hat{\alpha}\, \mathbb{E}_{it}\left[\log y\left(\omega_{j,t}\right)\right] \tag{28}$$

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

- 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 $\log y\left(\omega_{U1+}\right) = \phi_{0U} a_1$ for some coefficient ϕ_{0U} and $\log y\left(\omega_{P1}\right) = \phi_{0P} a_1 + \phi_x x_1 + \phi_s s_1$, for some coefficients $(\phi_{0P}, \phi_x, \phi_s)$. It follows that $y(\omega_{U1+})$ and $y(\omega_{P1})$ are indeed log-normal, with

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

where$\mathbb{E}\left[x_1|\omega_{U1+}\right] = \mathbb{E}\left[s_1|\omega_{U1+}\right] = 0$. Substituting these expressions into (27) gives us

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

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

- Matches between two islands of type ω_{P1+} and ω_{P2+}. 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_{P2+} takes log-linear form given by $\log y\left(\omega_{P2+}\right) = \phi_0 a_2 + \phi_x x_2 + \phi_s s_2$, for some coefficients (ϕ_0, ϕ_x, ϕ_s), for $i = 1, 2$. It follows that $y\left(\omega_{P2+}\right)$ is indeed log-normal, with

$$\mathbb{E}\left[\log y\left(\omega_{P2+}\right)|\omega_{P1+}\right] = \phi_a\mathbb{E}\left[a_2|\omega_{P1+}\right] + \phi_x\left(a_1 + \mathbb{E}\left[\varepsilon_2|\omega_{P1+}\right]\right) + \phi_s\mathbb{E}\left[\varepsilon_1 + u|\omega_{P1+}\right]$$

where

$$\begin{bmatrix}\mathbb{E}\left[a_2|\omega_{P1+}\right] \\ \mathbb{E}\left[\varepsilon_1|\omega_{P1+}\right] \\ \mathbb{E}\left[\varepsilon_2|\omega_{P1+}\right] \\ \mathbb{E}\left[u|\omega_{P1+}\right]\end{bmatrix} = \begin{bmatrix}\frac{\kappa_\varepsilon}{\kappa_a+\kappa_\varepsilon}x_1 \\ \frac{\kappa_a}{\kappa_a+\kappa_\varepsilon}x_1 \\ \frac{\kappa_u}{\kappa_u+\kappa_\varepsilon}s_1 \\ \frac{\kappa_\varepsilon}{\kappa_u+\kappa_\varepsilon}s_1\end{bmatrix}$$

Substituting these expressions into (27) gives us

$$\log y\left(\omega_{P1+}\right) = (1-\hat{\alpha})\frac{1}{1-\theta}\,a_1 + \hat{\alpha}\,\left[\phi_a\frac{\kappa_\varepsilon}{\kappa_a+\kappa_\varepsilon}x_1 + \phi_x\left(a_1 + \frac{\kappa_u}{\kappa_u+\kappa_\varepsilon}s_1\right) + \phi_s\left(\frac{\kappa_a}{\kappa_a+\kappa_\varepsilon}x_1 + \frac{\kappa_\varepsilon}{\kappa_u+\kappa_\varepsilon}s_1\right)\right]$$

By symmetry, equilibrium output for type ω_{P1+} must satisfy $\log y\left(\omega_{P1+}\right) = \phi_a a_1 + \phi_x x_1 + \phi_s s_1$. For this to coincide with the above condition for every z, it is necessary and sufficient that the coefficients (ϕ_a, ϕ_x, ϕ_2) solve the following system:

$$\begin{aligned}\phi_a &= (1-\hat{\alpha})\frac{1}{1-\theta} + \hat{\alpha}\phi_x \\ \phi_x &= \hat{\alpha}\left(\phi_a\frac{\kappa_\varepsilon}{\kappa_a+\kappa_\varepsilon} + \phi_s\frac{\kappa_a}{\kappa_a+\kappa_\varepsilon}\right) \\ \phi_s &= \hat{\alpha}\left(\phi_x\frac{\kappa_u}{\kappa_u+\kappa_\varepsilon} + \phi_s\frac{\kappa_\varepsilon}{\kappa_u+\kappa_\varepsilon}\right)\end{aligned}$$

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

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

Part (ii). The vector $\mu_t \in [0,1]^{10}$ records the fraction of islands in the economy that, as of stage 1 of period t, take each of the aforementioned ten type values.

$$\mu_t = \left[\begin{array}{cccccccccc} \mu_{U1} & \mu_{U1+} & \mu_{P1} & \mu_{P1+} & \mu_{F1} & \mu_{U2} & \mu_{U2+} & \mu_{P2} & \mu_{P2+} & \mu_{F2} \end{array} \right]'$$

That is, μ_x corresponds to the fraction of islands of type ω_x. The dynamics of the cross-sectional distribution of types depends on the match probabilities. We assume that the probability an island of type ω matches with an island of type ω' at the trading stage is proportional to population size. This generates a transaction matrix M so that the cross-sectional distribution of types can thus be summarized by a simple law of motion given by $\mu_{t+1} = M\mu_t$. *QED*

Proof of Proposition 5. The optimality conditions for the firms that produce the differentiated goods give

$$w_{i,t} = \mathbb{E}_{i,t}\left[p_{i,t}\right] \vartheta \frac{y_{i,t}}{n_{i,t}} \quad \text{and} \quad r_{i,t} = \mathbb{E}_{i,t}\left[p_{i,t}\right](1-\vartheta)\frac{y_{i,t}}{k_{i,t}}.$$

Next, the optimality condition for the firms that produce the capital good gives

$$1 = q_{i,t}\psi \tilde{n}_{i,t}^{\psi-1},$$

with $k_{i,t} = \tilde{n}_{i,t}^{\psi}$. Finally, the households labor-supply and Euler conditions give

$$w_{i,t} = V'(n_{i,t}) \quad \text{and} \quad q_{i,t} = \beta \mathbb{E}_{i,t+1}[r_{i,t+1}].$$

Combining the above we get

$$\begin{aligned} V'(n_{i,t}) &= \mathbb{E}_{i,t}[p_{i,t}]\vartheta \frac{y_{i,t}}{n_{i,t}} \\ k_{i,t}^{1/\psi-1} &= \beta\psi \mathbb{E}_{i,t}[p_{i,t}](1-\vartheta)\frac{y_{i,t}}{k_{i,t}} \end{aligned}$$

which simply equate the marginal costs of effort and investment with their corresponding expected marginal revenue products. By combining these two conditions with the production function (1), we can express $y_{i,t}$ and a function of A_i and $\mathbb{E}_{i,t}[p_{i,t}]$. Finally, replacing the latter from condition (19) gives condition (20), which completes the equilibrium characterization. Finally, the existence and uniqueness of the equilibrium follow from Theorem 2, which establishes that the equilibrium coincides with planner's solution, which in turn exists and is unique thanks to the usual convexity properties of the planner's problem/ *QED*

Proof of Proposition 6. It is easy to check that this case implies

$$p_{i,t} = P(y_{i,t}, y_{j,t}) = y_{i,t}^{-\frac{\gamma}{2}} y_{j,t}^{\frac{\gamma}{2}}.$$

Clearly, this is the same as condition (6) if we let $\eta = \gamma/2$. The result then follows essentially from Proposition 2 once we replace η with $\gamma/2$. *QED*

Proof of Proposition 7. Let $\hat{p}_{i,t} \equiv (1 - \tau_{i,t})p_{i,t}$ denote the after-tax price faced by firms in island i. Clearly, it suffices to pick $\tau_{i,t}$ so that, in the unique equilibrium, $\hat{p}_{i,t}$ depends only on the island's own productivity A_i and its trading partner's productivity A_j. Note then that $p_{i,t}$ is always given by a fixed function P of $y_{i,t}$ and $y_{j,t}$. It follows that it suffices to pick a function T such that

$$[1 - T(A_i, A_j, y_{i,t}, y_{j,t})]P(y_{i,t}, y_{j,t})$$

reduces to a function of A_i and A_j alone. For example, the government could guarantee both a unique equilibrium and and not sunspot volatility by setting

$$(1 - T(A_i, A_j, y_{i,t}, y_{j,t}) = 1 - \frac{\exp(\zeta \log A_i + \zeta^* \log A_i)}{P(y_{i,t}, y_{j,t})}$$

for arbitrary scalars ζ and ζ^*. The government could then fashion the response of equilibrium outcomes to the underlying fundamentals by appropriately choosing the coefficients ζ and ζ^*. *QED*

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

References

[ADD REFERENCES]