Sentiments*

George-Marios Angeletos
MIT and NBER

Jennifer La'O
Chicago Booth

September 26, 2010

Abstract

This paper develops a novel theory of fluctuations, one that seeks to capture the self-fulfilling nature of short-run phenomena within an otherwise canonical neoclassical framework. In our model, business-cycle fluctuations obtain without any innovation in technologies, preferences, and other fundamentals. To outside observers, these fluctuations may appear to be driven by random shifts in "market sentiment", by arbitrary waves of "optimism" and "pessimism", or by mystical forces known as "animal spirits" and "demand shocks". Yet, they are neither the product of equilibrium indeterminacy nor the symptom of irrationality. Rather, they rest merely on the decentralization of market interactions and the limits this imposes on market communication. These fluctuations thus uncover a type of "systemic risk" which, at least in our view, is endemic to the market mechanism. What is more, they are consistent with a notion of constrained efficiency that leaves no room for conventional stabilization policy. A new paradigm thus emerges—one that shifts the focus away from technology and preference shocks, one whose positive aspects have a certain Keynesian flavor, and yet one whose foundations and policy implications remain uncompromisingly Neoclassical.

*This paper supersedes a 2009 working paper that contained a different model but a similar methodological message. A previous draft was entitled "Animal Spirits". We thank Tarek Hassan and Per Krusell for their discussions of our paper at the 2010 Bank of Portugal Monetary Economics Conference. We also thank for stimulating discussions and comments Robert Barro, John Geanakoplos, Emmanuel Farhi, Erik Hurst, Ivan Werning, and seminar participants at Harvard, MIT, NYU, the 2009 SED meeting in Istanbul, the 2009 EEA/ESEM meeting in Barcelona, the 2010 Bank of Portugal Monetary Economics Conference, the 2010 Midwest Macroeconomics Meeting at Michigan State, and the 2010 CRETE conference in Tinos.

"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 ... Some models have collective shocks to workers' willingness to work. ... 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?" Kocherlakota (2010)

1 Introduction

What drives, or explains, short-run fluctuations in macroeconomic activity and asset prices?

Mainstream macroeconomic models[1] differ in the emphasis they give to different propagation mechanisms and market frictions, but ultimately reach the same basic answer to the aforementioned question: the bulk of short-run fluctuations is attributed to shocks in preferences and technologies.[2]

We find this answer unsatisfactory. No matter how insightful these models have proved to be for certain issues, or how well they may fit the data, the high-frequency disturbances in preferences and technologies they require are—to use Kochelakota's words—patently unrealistic.

Consider, for example, the ongoing recession. Firms appear reluctant to expand production and employment, largely because they are pessimistic about consumer demand. Consumers, on their part, appear reluctant to spend, largely because they are pessimistic about labor-market conditions. More generally, the recession appears to be sustained, in large part, by a wave of pessimism about aggregate economic activity. But where is any indication of the kind of adverse preference and technology shocks that standard models would assume in order to explain, or fit, the recession?[3]

This issue poses a serious challenge for modern macroeconomic theory—one of evident empirical and policy relevance. Does the standard paradigm helps us understand the *origins* of short-run fluctuations? Should the aforementioned disturbances be treated only as metaphors, as proxies for other shocks that we, economists, have failed to comprehend and model more properly? If so, how could we possibly trust the policy lessons of this paradigm?

[1]By "mainstream" we mean the prototypical RBC and New-Keynesian models, as well as the more recent DSGE models. This excludes models with multiple equilibria or irrational agents, which we discuss in due course.

[2]To give a representative example, in Smets and Wouters (2003), such shocks explain 80% of the variation in consumption and investment at the four-quarter horizon, and 70% of the variation in output and employment.

[3]Of course, financial markets are playing a central role—a role that we leave out of our analysis for pedagogical reasons. But this does not affect our key point. One cannot convincingly attribute the recession to an *exogenous* financial shock. Rather, one should ask what triggered the *endogenous* asset price movements that ultimately caused this "shock". One is then quickly back to the same starting point: was it a shock to preferences and technologies?

Some economists have even argued that modern macro has entirely missed the essence of short-run fluctuations. Most provocatively, Krugman (2009) declares that the neoclassical paradigm is "silly" and that, instead, "Keynesian economics remains the best framework we have for making sense of recessions and depressions." In a similar vein, Akerlof and Shiller (2008) and Shiller (2009) criticize modern macro for ignoring the influence of "animal spirits", and for failing to explain the waves of optimism and euphoria followed by pessimism and distress. In concert with old Keynesianism, these economists then conclude that there is an *obvious* rationale for stabilization policy.

In this paper, we seek to address the aforementioned challenge by developing a novel theory of fluctuations. This theory shifts the focus away from preference and technology shocks; it formalizes a form of "systemic risk" which, in our view, is endemic to the market mechanism; and it helps accommodate a number of mystical notions that rest in the core of Keynesian thinking and that are popular among practitioners and market pundits. And yet, this theory does not share either the methodological or the policy message of the aforementioned vocal critiques.

Preview. In our model, agents are fully rational; preferences and technologies are standard; markets are Walrasian; there are no nominal frictions, no externalities, and no non-convexities; the equilibrium is unique; and there is no room for correlation devices or lotteries. In these respects, our theory has squarely neoclassical foundations. Nonetheless, our theory opens the door to an entirely novel form of uncertainty, one that underscores the self-fulfilling nature of short-run fluctuations.

Formally, we show that equilibrium allocations and prices respond to a type of exogenous shocks that we call "sentiment shocks." By assumption, these shocks are orthogonal to past, current and future fundamentals (preferences, technologies, etc.), as well as to the agents' beliefs about these fundamentals. And yet, along the unique equilibrium of the economy, they cause self-fulfilling variation in the agents' expectations of economic outcomes—be it the firms' expectations of demand, the households' expectations of income, or the investors' expectations of asset returns.

The resulting fluctuations can thus be interpreted as self-fulfilling waves of optimism and pessimism. These waves may initially hit only a small fraction of the population, but may later on build force as they spread through the economy—indeed in a manner akin to the spread of infectious diseases. Sooner or later, however, these waves die off: long-run outcomes are pinned down by fundamentals. In this sense, a boom in our model looks like a period during which more and more people get "exuberant", only to "come back to their senses" after a while; and a recession looks like a period during which people loose their "confidence" in the economy, only to regain it later on.

Furthermore, these fluctuations have a Keynesian feel: while long-run outcomes are dictated by "supply-side" forces (preference and technologies), short-run fluctuations are driven by shocks that cause aggregate demand (consumption and investment) to deviate from the neoclassical benchmark. Alas, isn't this what the original Keynesian notion of "demand shocks" was meant to capture?

Last but not least, these fluctuations are broadly consistent with the data. They feature the right co-movement among all macroeconomic variables—with output, employment, consumption, investment and asset prices all moving in tandem—and they also exhibit hump-shape dynamics similar to those found in structural VAR models.

What sustains these fluctuations in our model is a form of "systemic risk" that, at least in our view, is *endemic* to the market mechanism. Let us elaborate what we mean by this.

The mere fact that agents specialize in different productive activities implies that they must trade with one another. This introduces an important form of interdependence: the choices and fate of each economic agent—be it a firm, a consumer, or an investor—crucially depends on the choices and fate of other agents, simply because the latter determine the terms of trade he himself faces in future market interactions. This interdependence, in turn, opens the door to an important form of uncertainty: each agent can face uncertainty about the choices of other agents, and thereby about his own future terms of trade, *beyond* any uncertainty he may face about the underlying economic fundamentals. In a macro context, this means that the economy can feature a form of macroeconomic risk that is not explained by fundamentals. Moreover, as it is in each individual's best interest to respond to this risk, this risk becomes largely self-fulfilling in the aggregate. This risk therefore has the flavor of equilibrium multiplicity—but its true nature is very different.

Most macroeconomic models have ruled out this kind of risk only inadvertently, because of convenient but unrealistic assumptions: they impose centralized Arrow-Debreu markets or, at the very least, that all agents share the same information about the aggregate state of the economy. When this is the case, it is *as if* all agents get in the same room, talk to one another, and instantaneously communicate with one another any information they may have about their abilities, needs, and other fundamentals. Of course, as Hayek first noted, this communication may happen only indirectly and unintentionally, through the "invisible hand" of the market mechanism. But as long as this communication is perfect, it should be no surprise that market outcomes are pinned down by fundamentals, leaving no room for the type of fluctuations documented in this paper.

In contrast, by introducing dispersed information and decentralized market interactions, our theory seeks to capture the idea that the communication that takes place in the real world through the market mechanism cannot possibly be as perfect as the one assumed in the standard paradigm. Our key contribution is then to show that, as soon as this is the case, fundamentals may no longer pin down equilibrium allocations and prices, leaving room for a distinct form of extrinsic risk. Furthermore, the fact that market interactions facilitate more and more communication over time is what explains, perhaps paradoxically, both the initial buildup and the eventual fading of the waves of optimism and pessimism we document in this paper. In these respects, the type of "systemic risk" we uncover in this paper is, indeed, endemic to the market mechanism.

At a more abstract level, our notion of systemic risk is closely connected to the notions of strategic and higher-order uncertainty in games. Indeed, the unique equilibrium of our Walrasian economy can be mapped to the unique rationalizable outcome of a certain game. The sunspot-like fluctuations we document can then be mapped 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. In this respect, our paper complements Morris and Shin (2002), Woodford (2003), and a growing applied literature that studies various implications of dispersed information and higher-order uncertainty.[4] This prior work has nevertheless ruled out our sunspot-like fluctuations because it has restricted attention to settings in which all the variation in higher-order beliefs is spanned by the one in first-order beliefs and fundamentals. We, instead, bring these fluctuations to life by allowing for richer information structures.

Furthermore, the fact that our model is a Walrasian economy comes with two important observations. First, agents do not care *per se* about the actions and beliefs of one another. Rather, as in Lucas (1972) and any other competitive rational-expectations setting, they only need to form correct expectations about the statistical behavior of prices. From this perspective, the systemic risk we document is best understood as a novel form of extrinsic uncertainty in competitive allocations and prices—one that obtains only in so far information remains dispersed. Second, whereas information is often treated as exogenous in games, communication is a central aspect of the market mechanism. This explains why the modeling of communication take a central place in our theoretical exercise.

More specifically, we first split the economy into different "islands" and endow each island with some initial private information. This is similar to, inter alia, Lucas (1972), Lagos and Wright (2005), Amador and Weill (2009), and Lorenzoni (2010), except that the information structure here is rich enough to introduce our sentiment shocks. We then let the islands trade and communicate with one another through random pair-wise matches. It follows that information disseminates gradually in the economy, in a manner that resembles the "percolation of information" in Duffie and Manso (2007) and Duffie, Giroux and Manso (2010). As already mentioned, the resulting information flows are then key to understanding the rich dynamics of our sentiment-driven fluctuations.

Finally, the detailed micro-foundations of our model facilitate a transparent welfare analysis. Our key result in this respect is that, under certain conditions, the equilibrium coincides with the solution to the problem of a planner that faces the same resource and communication constraints as the market mechanism but can otherwise freely dictate behavior. This result is a theoretical contribution on its own right: it establishes, in effect, a variant of the first welfare theorem for a class of economies

[4]The recent literature starts with Mankiw and Reis (2002), Sims (2003), and the aforesaid two works. Key earlier contributions include Phelps (1970), Lucas (1972), Barro (1976), and Townsend (1983). For more references, see Angeletos and La'O (2009). Complementary is also the literature on global games (e.g., Morris and Shin, 2001).

that feature decentralized trading and imperfect communication. But it also offers an important policy lesson, one that contrasts sharply with conventional wisdom and Keynesian thinking. Even if we take for granted that most short-run volatility is disconnected from fundamentals, and perhaps driven by mysterious forces that one may call "market sentiments," "animal spirits", or "demand shocks", this fact does not *by itself* provide a rationale for stabilization policy.

Layout. The paper is organized as follows. Section 2 introduces the model. Section 3 characterizes the equilibrium. Section 4 contains our core positive result. Section 5 presents our waves of optimism and pessimism. Section 6 considers a useful extension. Section 7 studies policy and efficiency. Section 8 concludes. The Appendix includes any proofs omitted in the main text.

2 The model

We start by briefly describing the key ingredients of our model. We then fill in the details.

Key model ingredients. There is 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. All agents are price takers. Time is discrete, indexed by $t \in \{0,1,...\}$, and each period contains two stages. Each island produces two goods. The one is a differentiated consumption good that is produced in stage 1 and can be consumed either at "home" (by the same island) or "abroad" (by another island). The second is a specialized investment good that is produced in stage 2 and can be utilized only at home, in the production of the local consumption good. The one worker-firm pair produces the consumption good; the other pair produces the investment good.

All fundamentals (technologies, preferences, etc) are constant over time at both the aggregate and the individual level. Apart from simplifying the analysis, this assumption means a fortiori that none of the fluctuations in our model are driven by high-frequency in fundamentals. Most crucially for our purposes, some element of the fundamentals—here, the cross-sectional distribution of productivities—is not commonly known. Rater, each island is initially endowed with some private information. Subsequently, islands meet and trade with one another in random pairwise matches. Whenever two islands meet, they exchange any information they may have acquired in the past. Information thus diffuses in the economy, but only slowly.

Firms and technologies. Consider the firm that produces the consumption good in island i. Its technology is given by

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

where $y_{i,t}$ is the quantity produced, $k_{i,t}$ and $n_{i,t}$ are the capital and labor inputs, A_i is the local productivity (TFP), and $\vartheta_k, \vartheta_n \in (0,1)$, with $\vartheta_k + \vartheta_n \leq 1$. Next, consider the firm that produces the local investment good. Its technology is given by

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

where $x_{i,t}$ is the quantity produced, $\tilde{n}_{i,t}$ is the labor input, and $\psi \in (0,1)$. The profits of these two firms are, respectively,

$$\pi^c_{i,t} = p_{i,t}y_{i,t} - w_{i,t}n_{i,t} - r_{i,t}k_{i,t} \qquad \text{and} \qquad \pi^k_{i,t} = q_{i,t}x_{i,t} - \tilde{w}_{i,t}\tilde{n}_{i,t},$$

where $p_{i,t}$ denotes the local price of the local consumption good, $w_{i,t}$ denotes the local wage for the worker that works in the consumption sector, $r_{i,t}$ denotes the local rental rate of capital, $q_{i,t}$ denotes the local price of the investment good, and $\tilde{w}_{i,t}$ the local wage for the worker that works in the investment sector. Finally, each firm maximizes its expectation of the local valuation of its profits, namely $\lambda_{i,t}\pi^c_{i,t}$ for the former and $\lambda_{i,t}\pi^k_{i,t}$ for the latter, where $\lambda_{i,t}$ denotes the marginal value of wealth for the representative household in island i.

Households and preferences. The preferences of the household on island i are given by

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

where $\beta \in (0,1)$ is the discount factor, c_{it} is consumption of the "home" good, c^*_{it} is consumption of the "foreign" good (the consumption good of the island that i gets matched with during stage 2), and n_{it} is the labor of the first worker, and $\tilde{n}_{it}$ is the labor of the second worker. Finally, to simplify the exposition, we assume the following specifications for the functions U and V:

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

for some $\eta \in (0,1)$ and $\epsilon > 1$. We relax these assumptions in Section 6.

Fundamentals and initial information. In the beginning of time, Nature draws two independent random variables θ and ξ, from finite sets $\mathcal{S}_\theta$ and $\mathcal{S}_\xi$, respectively, and with probability functions $\mathcal{F}_\theta : \mathcal{S}_\theta \to [0,1]$ and $\mathcal{F}_\xi : \mathcal{S}_\theta \to [0,1]$. Let $z \equiv (\theta, \xi)$ be the aggregate state, $\mathcal{S}_z \equiv \mathcal{S}_\theta \times \mathcal{S}_\xi$ its support, and $\mathcal{F} : \mathcal{S}_z \to [0,1]$ its probability function, with $\mathcal{F}(\theta,\xi) \equiv \mathcal{F}_\theta(\theta)\mathcal{F}_\xi(\xi)$ for all θ, ξ. Conditional on z, Nature also draws a pair of random variables for each island. The first one, which identifies local productivity A_i, is drawn in an i.i.d. fashion from a finite set $\mathcal{S}_A \subset \mathbb{R}_+$ according to a conditional probability function $\mathcal{F}_A : \mathcal{S}_A \times \mathcal{S}_\theta \to [0,1]$ that depends on z only through θ. This means that the fundamentals are pinned down by θ alone, which in turn permit us to interpret ξ as a form of noise. The second one, which can be interpreted as a local signal about the aggregate stage z and is denoted by $\tilde{\omega}_i$, is drawn, also in an i.i.d. fashion, from a finite set $\mathcal{S}_{\tilde{\omega}}$ and according to a conditional probability function $\mathcal{F}_{\tilde{\omega}} : \mathcal{S}_{\tilde{\omega}} \times \mathcal{S}_z \to [0,1]$ that depends on the entire z.

The aggregate state z is *not* revealed to any island. Rather, at $t = 0$, each island gets to see only its own productivity A_i and signal $\tilde{\omega}_i$. We let $\omega_i \equiv (A_i, \tilde{\omega}_i)$ and refer to it as the "initial local state". By construction, the support of ω_i is $\mathcal{S}_\omega \equiv \mathcal{S}_A \times \mathcal{S}_{\tilde{\omega}}$ and its conditional probability function is $\mathcal{P}_0$, where $\mathcal{P}_0(\omega|z) \equiv \mathcal{F}_A(A|\theta) \cdot \mathcal{F}_{\tilde{\omega}}(\tilde{\omega}|z)$. This variable encodes information, not only about the fundamentals, but also about the information of other islands. In this sense, ω_i is the analogue in our competitive economy of a Harsanyi type in games, and $\mathcal{P}_0$ is the analogue of the common prior.

Timing, markets, and communication. In stage 1, each island operates in isolation. The local labor and rental markets open, households decide how much labor and capital to supply, and the consumption-good firms make the input and production choices. In stage 2, each island i is randomly matched with another island $j \neq i$. At this point, the two islands share their informations and trade their consumption goods in an "integrated" market. In addition, local markets for investment goods open, and the firms producing these goods make their input and production choices.[5]

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 stage 2 of that period. Let $\omega_{i,t}$ denote the information 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 $m(i,t)$ henceforth identifies the trading partner of island i during stage 2 of period $t-1$. Finally, we let $\mathcal{S}_t$ denote the support of $\omega_{i,t}$ and $\mathcal{P}(\omega_{i,t}|z)$ its probability conditional on the aggregate state being z. The latter obtains in the obvious way from the initial prior P_0.

Budget constraints and capital accumulation. Given the aforementioned market structure, the period-t budget constraint of the household on island i can be written as follows:

$$p_{i,t}c_{i,t} + p^*_{i,t}c^*_{i,t} + q_{i,t}x_{i,t} \leq \pi^c_{i,t} + w_{i,t}n_{i,t} + r_{i,t}k_{i,t} + \pi^k_{i,t} + \tilde{w}_{i,t}\tilde{n}_{i,t} \tag{3}$$

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 propositions impose full depreciation ($\delta = 1$). We nevertheless discuss the more general case at various points of the analysis in order to develop more complete intuitions.

[5]When the two islands meet, they could trade, in principle, all the goods in the economy, as well as an arbitrary set of financial claims. For expositional simplicity, our baseline model restricts the islands to trade only their two consumption goods. This assumption, which is embedded in the budget constraint (3) and in condition (iii) of the equilibrium definition later on, is relaxed in Section 6. The communication that takes place directly in our baseline model could then be re-interpreted as the outcome of perfectly revealing markets among the two islands.

Competitive equilibrium. We henceforth let the stage-2 leisure be the local numeraire for each island and accordingly normalize $\tilde{w}_{i,t} = 1$. Since communication takes place only in stage 2, an island's information during this stage is the same as the one during stage 1 of the next period. We can thus represent any allocation and price system with a collection of functions $(p, p^*, w, r, n, k, y, c, c^*, \tilde{n}, \tilde{x})$ that the following hold: for stage-1 variables, $w_{i,t} = w_t(\omega_{i,t})$, $r_{i,t} = r_t(\omega_{i,t})$, $n_{i,t} = n_t(\omega_{i,t})$, $k_{i,t} = k_t(\omega_{i,t})$, and $y_{i,t} = y_t(\omega_{i,t})$; and for stage-2 variables, $p_{i,t} = p_t(\omega_{i,t+1})$, $p^*_{i,t} = p^*_t(\omega_{i,t+1})$, $c_{i,t} = c_t(\omega_{i,t+1})$, $c^*_{i,t} = c^*_t(\omega_{i,t+1})$, $\tilde{n}_{i,t} = \tilde{n}_t(\omega_{i,t+1})$, $\tilde{x}_{i,t} = \tilde{x}_t(\omega_{i,t+1})$, and $q_{i,t} = q_t(\omega_{i,t+1})$. This serves to clarify the information upon which allocations and prices can be contingent. A competitive equilibrium is then defined in an otherwise familiar manner.

Definition 1. *An equilibrium is collection of functions* $(p, p^*, w, r, n, k, y, c, c^*, \tilde{n}, \tilde{x})$ *such that*

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

(ii) the associated prices clear all markets; and

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

3 Equilibrium characterization

Consider the household on island i. Let $\lambda_{i,t} = \lambda_t(\omega_{i,t})$ denote the Lagrange multiplier on the period-t budget. The optimality condition for labor supply in stages 1 and 2 give, respectively,

$$V'(n_{i,t}) = \mathbb{E}_{i,t}[\lambda_{i,t}]w_{i,t} \quad \text{and} \quad \lambda_{i,t} = 1. \tag{4}$$

where $\mathbb{E}_{i,t}[\cdot]$ is a short-cut for the rational expectation conditional on $\omega_{i,t}$. The optimal consumption choices, on the other hand, satisfy

$$U_c\left(c_{i,t}, c^*_{i,t}\right) = \lambda_{i,t}p_{i,t} \quad \text{and} \quad U_{c^*}\left(c_{i,t}, c^*_{i,t}\right) = \lambda_{i,t}p^*_{i,t} \tag{5}$$

Combining condition (5) with the corresponding one for i's trading partner, imposing market clearing and balanced trade, and using the Cobb-Douglas specification of U and the fact that $\lambda_{i,t} = 1$, we obtain the following:

$$c_{i,t} = (1-\eta)y_{i,t}, \quad c^*_{i,t} = \eta y_{j,t}, \quad c_{j,t} = (1-\eta)y_{j,t}, \quad \text{and} \quad c^*_{j,t} = \eta y_{i,t}, \tag{6}$$

$$p_{i,t} = U_c\left((1-\eta)y_{i,t}, \eta y_{j,t}\right) = y_{i,t}^{-\eta}y_{j,t}^{\eta} \tag{7}$$

The first condition has a familiar interpretation: each island consumes a fraction η of his total expenditure on the "home" good. The second condition reveals that the price of the "home" good increases with "foreign" output. This last property plays a central role in our subsequent results. It captures, more broadly, the interdependence that market interactions induce among agents. This interdependence—here originating from an otherwise inconsequential pecuniary externality of the Walrasian type—is what opens the door to the distinct form of system risk we shall document later.

Consider, now, the firms that produce the differentiated consumption goods. These firms maximize $\mathbb{E}_{i,t}[\lambda_{i,t}\pi_{i,t}]$. Using $\lambda_{i,t} = 1$, the first-order condition with respect to labor yields

$$w_{i,t} = \mathbb{E}_{i,t}\left[p_{i,t}\right]\vartheta_n \frac{y_{i,t}}{n_{i,t}}, \tag{8}$$

while the one with respect to capital yields

$$r_{i,t} = \mathbb{E}_{i,t}\left[p_{i,t}\right]\vartheta_k \frac{y_{i,t}}{k_{i,t}}. \tag{9}$$

These conditions have a familiar interpretation: the wage and the rental rate are equated with the corresponding marginal revenue products. Importantly, though, firms face uncertainty about the prices at which they will sell their products and base their expectations on asymmetric information.

Combining conditions (4) and (8), and replacing $p_{i,t}$ from (7), we reach the following condition, which equates the disutility of labor in each island with the local expectation of its marginal revenue product:

$$V'(n_{i,t}) \;=\; \mathbb{E}_{i,t}[p_{i,t}]\vartheta_n \frac{y_{i,t}}{n_{i,t}} \;=\; \vartheta_n \mathbb{E}_{i,t}[y_{j,t}^{\eta} y_{i,t}^{1-\eta} n_{i,t}^{-1}] \tag{10}$$

This condition helps capture the interdependence of employment and production decisions in the economy: local employment and output depend on expectations of future terms of trade, and thereby on expectations of the level of employment and output in the rest of the economy.

To see this more clearly, take the special case that $\eta = 1$ (which one can interpret as full specialization); the above condition then reduces to

$$n_{i,t}^{\epsilon} = \mathbb{E}_{i,t}[p_{i,t}] = \mathbb{E}_{i,t}[y_{j,t}].$$

Suppose, now, that firms in one island suddenly become "optimistic" in the sense that they raise their expectations about the likely level of employment and output in the rest of the economy even though they believe that nothing has changed in either their own fundamentals or those of others. Suppose further that this expectation is rational. These firms can then infer, correctly, that their own terms of trade—here captured by $p_{i,t}$—are likely to improve. As this increases their expected marginal revenue product of labor, these firms find it optimal to raise local employment and output despite the fact that, at least in their view, nothing changed in the underlying fundamentals. Conversely, if some firms become "pessimistic" about aggregate economic activity, they will rationally respond by cutting down their own employment and production levels.

These intuitions seem to make apparent how booms and recessions could be driven by self-fulfilling waves of optimism and pessimism. But, behold! These intuitions are only *partial*; they fail to take into account the discipline imposed by *general-equilibrium* reasoning. Indeed, this discipline is so powerful within the neoclassical paradigm that the aforementioned notions can make sense only *off* equilibrium: general-equilibrium allocations and prices are ultimately determined by fundamentals, leaving no room for waves of "optimism" and "pessimism" to emerge as equilibrium phenomena. From this perspective, the key contribution of our paper is to show how the aforementioned intuitions can be resurrected, and how waves of "optimism" or "pessimism" can emerge as equilibrium phenomena, once one relaxes the conventional but unrealistic assumption of perfect communication. This will be established in due course. For the time being, we continue with the characterization of the equilibrium, focusing now on the investment and asset-price side of the economy.

Thus consider the firms that produce the local investment good. Their first-order condition with respect to labor gives $\tilde{w}_{i,t} = 1 = q_{i,t}\psi\tilde{n}_{i,t}^{\psi-1}$. Since $\tilde{w}_{i,t} = 1$ and $x_{i,t} = \tilde{n}_{i,t}^{\psi}$, this gives gross investment as an increasing function of the price of capital:

$$x_{i,t} = (\psi q_{i,t})^{\frac{\psi}{1-\psi}} \tag{11}$$

This is similar to standard q theory. However, there is an important novelty when we look at the determination of the price of capital. From the Euler condition of the household, along with 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 get the following pricing equation:

$$\begin{aligned} q_{i,t} &= \beta\mathbb{E}_{i,t+1}\left[r_{i,t+1} + (1-\delta)q_{i,t+1}\right] && (12) \\ &= \beta\sum_{s=t+1}^{\infty}(\beta(1-\delta))^{s-t-1}\vartheta_k\mathbb{E}_{i,t+1}\left[y_{j,s}^{\eta}y_{i,s}^{1-\eta}k_{i,s}^{-1}\right] && (13) \end{aligned}$$

It follows that the price of capital, and thereby investment, depends positively on expectations of future aggregate economic activity, which in turn is endogenous to current investment decisions. A similar result holds for other asset prices in our model, such as those of claims on firm profits. Once again, this underscores the interdependence of economic outcomes in the economy: if asset prices and investment go up in one part of the economy (one island), this causes, ceteris paribus, an asset-price and investment boom in other parts of the economy (the island's likely trading partners). But, as before, this is only a partial intuition; it still remains to see whether this interdependence can actually open the door to self-fulfilling fluctuations in general equilibrium.

Complete characterization. In general, the equilibrium of the economy is obtained as the fixed point of the dynamic system defined by conditions (**??**), (11) and (13). This system has both forward- and backward-looking aspects. When information is asymmetric, solving this system becomes intractable. This motivates the following assumption, which helps reduce the dynamic equilibrium to a sequence of static, and much more tractable, fixed-point problems.

Assumption. *Capital depreciates fully ($\delta = 1$).*

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

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

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

$$\log y(\omega_{i,t}) = (1-\hat{\alpha})\left(\tfrac{1}{1-\hat{\vartheta}} \log A(\omega_i)\right) + \hat{\alpha}\,\mathbf{E}\left[\log y(\omega_{j,t})|\,\omega_{i,t}\right], \tag{14}$$

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

Condition (14) expresses the equilibrium level of output in an island as a simple function of the local productivity and the local expectations of output in the rest of the economy. Once output is determined by solving this fixed-point problem, all other endogenous variables (employment, investment, consumption, wages, asset prices, etc.) are pinned down by the relevant conditions in the preceding analysis. Finally, the fact that $\alpha > 0$ means that there is a positive feedback between local and aggregate economic activity, while the fact that $\alpha \in (0,1)$ guarantees that the aforementioned fixed point is unique and can be obtained by iterating condition (14).

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

We conclude this section by noting that Proposition 1 facilitates a certain game-theoretic interpretation: it is *as if* the islands in our economy play a game in which best responses are given by condition (14). The coefficient α, which conveniently summarizes the degree of interdependence in our economy, can then be interpreted as the degree of strategic complementarity in the aforementioned fictitious game. We will utilize this game-theoretic interpretation at a later point. For now, we stress that the interdependence—or complementarity—in our economy originates merely in specialization and trade. This fact is manifested in the property that α is increasing in η: if agents were consuming their own products ($\eta = 0$), there would be no trade and no interdependence ($\alpha = 0$). All in all, this interdependence is merely an example of the pecuniary externalities that are present in *any* Walrasian setting. This underscores that the type of interdependence we are capturing in this paper is indeed endemic to the market mechanism.

4 Imperfect communication and sunspot fluctuations

We now move to the core positive result of our paper. We accommodate the informal notions of "optimism" and "pessimism" by employing the formal notions of "sunspot fluctuations" and "extrinsic risk". Notwithstanding the fact that we are concerned with a class of convex, unique-equilibrium economies, the latter are defined as in the same way as in the pertinent literature (e.g., Cass and Shell, 1983).

Definition 2. *We say that the economy features "sunspot fluctuations" or "extrinsic risk" if and only if there exist exogenous random shocks that do not cause variation in either the fundamentals or any agent's expectations of these fundamentals, and yet cause variation in equilibrium allocations and prices. Whenever this is the case, we call these shocks "sentiment shocks".*

We next note that, for our purposes, perfect (resp., imperfect) communication is synonymous to symmetric (resp., asymmetric) information. 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, communication is imperfect.*

Consider first the "only if" part. If communication is perfect, and agents reach symmetric information, then the entire cross-sectional distribution of productivities becomes commonly known. It is then straightforward to verify that the unique solution to (14) is the following:

$$\log y_{i,t} = \log y_i^* \equiv \tfrac{1}{1-\hat{\vartheta}}\left[(1-\hat{\alpha})\log A_i + \hat{\alpha}\,\overline{\log A}\right] \tag{15}$$

where $\overline{\log A}$ is an aggregate productivity index, which is itself commonly known.[6] This condition gives the equilibrium output of an island as log-linear combination of the local productivity and an aggregate productivity index. It follows that all macroeconomic variables are pinned down by aggregate productivity. Along with the fact that the latter is time invariant, this rules out, not only sunspot fluctuations, but also *any* time-series variation in aggregate economic outcomes—thus providing a particularly clean benchmark for our subsequent results.

Consider next the "if" part. If communication is imperfect, then information can be asymmetric at least at some point of time. We can then prove the result with a simple example.[7] There is only one period (for simplicity) and the islands are split into two groups, the ones in the "South" and the ones in the "North". Productivities and information differ only across groups. Agents in the South know their own productivity but receive only a noisy private (group-specific) signal about the productivity in the North; let this signal be $x = A_{north} + \varepsilon$, where A_{north} is the productivity in the North and ε is a noise variable. Agents in the North, on the other hand, know the productivities in both locations. In addition, they receive a noisy private signal $s = x+u$, where x is as before and u is an another noise variable; this signal is meant to capture the information (or beliefs) that Northern agents may have about the information (or beliefs) of Southern agents. In equilibrium, Southern agents find it optimal to condition their production choices on x because this signal helps them forecast the North's output, which in turn affects the South's "export prices". But then Northern

[6] This aggregate productivity index is defined as a log-linear aggregator of the cross-sectional distribution of productivities: $\overline{\log A} \equiv \frac{1}{\eta}\log\sum_A \exp(\eta \log A)\mathcal{F}_A(A|\theta)$, where $\mathcal{F}_A(A|\theta)$.

[7] The technical details of this example are spelled out in the Appendix; here we give a brief sketch.

agents find it optimal to condition their own production choices on s because this signal helps them forecast the South's output and hence their own terms of trade. It follows that allocations and prices are sensitive to the noise u. But then note that u is, by assumption, orthogonal to the underlying fundamentals. Furthermore, since only the Northern agents observe s, and since these agents already know the fundamentals, u cannot possibly affect any agent's beliefs about the fundamentals. This proves that the economy features sunspots fluctuations, with u capturing what we earlier defined as a sentiment shock.

Clearly, these insights are not limited to the specific model we have considered in this paper. In any standard, unique-equilibrium macroeconomic model, equilibrium allocations and prices are ultimately pinned down by technologies, preferences, and other fundamentals. As anticipated in the Introduction, this is because these models impose, in effect, that agents can perfectly communicate with one another their information and intended courses of action. Once this unrealistic assumption is relaxed, our distinct type of sunspot volatility is bound to emerge.

From a game-theoretic perspective, this sunspot volatility can be attributed to a certain from of higher-order uncertainty. To see the last point, let $f_i \equiv \frac{1}{1-\hat{\vartheta}} \log A_i$ and iterate condition (14) to get the equilibrium output of an island as a function of the local hierarchy of beliefs:

$$\log y_{i,t} = (1 - \hat{\alpha}) \left\{ f_i + \hat{\alpha} \mathbf{E}_{i,t}[f_j] + \hat{\alpha}^2 \mathbf{E}_{i,t}[\mathbf{E}_{j,t}[f_k]] + ... \right\} \tag{16}$$

where $\mathbf{E}_{i,t}[f_j]$ capture first-order beliefs, $\mathbf{E}_{i,t}[\mathbf{E}_{j,t}[f_k]]$ capture second-order beliefs, and so on. When information is symmetric, first-order beliefs are commonly known, implying that second- and higher-order beliefs collapse to first-order beliefs. It follows that outcomes are driven by first-order beliefs, leaving no room for sunspot fluctuations. When instead information is asymmetric, higher-order beliefs need not coincide with second-order beliefs. Our "sentiment shocks" can thus be interpreted as shocks that are independent of either the fundamentals or first-order beliefs, and nevertheless cause variation in second- and higher-order beliefs.

These points taken, we invite the reader not to take these game-theoretic intuitions too seriously. We do not imagine the agents making all this higher-order reasoning. Rather, we imagine the agents forming beliefs about aggregate economic activity and market prices. Both the sentiment shocks and the aforementioned higher-order beliefs are then merely modeling devices that help formalize how variation in the agents' beliefs about economic activity and market prices can obtain without variation in their beliefs about the fundamentals. It is only this distinct, self-fulfilling variation in expectations of economic activity that our analysis cares to capture.

We finally wish to re-iterate that our model is a Walrasian economy, not a game; that the interdependence that is present in our model originates merely in the specialization of economic activity and the consequent need for trade; and, finally, that information can remain asymmetric only in so far the communication that takes place through market interactions is imperfect. Since specialization, trade, and communication are central aspects of the market mechanism, we conclude that the sunspot volatility we document in this paper is best understood as a form of extrinsic risk—or "systemic risk"—that is endemic to the market mechanism.

5 Sentiments and the Business Cycle

In this section we consider an example that helps illustrate the rich dynamics that our sunspot fluctuations can follow. This example introduces a certain type of sentiment shocks that initially "infects" only only a small portion of the population in the economy. As these agents meet, trade, and communicate with other agents, these shocks propagate in the economy in a manner that resembles contagion effects, or the spread of fads and rumors. The resulting fluctuations take the form of waves of "optimism" and "pessimism" that keep building up force for a while before they eventually fade away. As these fluctuations obtain without any change in the underlying fundamentals, they can be mis-intepreted by outside observers as the product of irrational forces or multiple equilibria. Furthermore, as these fluctuations represent short-run deviations from what predicted by the standard neoclassical paradigm, they appear to capture the Keynesian notion of "demand shocks". All in all, the example of this section draws a close connection to many popular, informal descriptions of short-run phenomena, while at the same time offering an entirely novel, and unconventional, formal explanation of these phenomena. Importantly, these fluctuations also match some key qualitative features of empirical business cycles.

Specification. We modify some of the technical assumptions of our model in order to facilitate a closed-form solution. In particular, we let the exogenous uncertainty be Gaussian; we also the islands receive some information about their likely matches at the trading round.

The fundamental θ is now a pair of random variables a_1 and a_2, which are drawn from independent Normal distributions. 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 islands have log-productivity a_1 and rest have log-productivity a_1. Let us refer to the former half as "group 1" and the second one as "group 2".

Each group is then split in three subgroups of islands. Islands in the first subgroup get to observe only their own productivities; we refer to them as "uninformed". Islands in the second subgroup, which we refer to as "partially informed", get to see two additional signals. The signals are given by

$$x_1 = a_2 + \varepsilon_1 \quad \text{and} \quad s_1 = x_2 + u \quad \text{for group 1}$$

$$x_2 = a_1 + \varepsilon_2 \quad \text{and} \quad s_2 = x_1 + u \quad \text{for group 2}$$

The x signals are thus signals of the productivity of the other group, while the s signals are signals about the information of the other group. Finally, the remaining islands get to observe the entire state z; were refer to them as "fully informed".

This structure creates an "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 exogenous fundamentals and, most crucially, about the endogenous level of aggregated economic activity.

More specifically, we make the following assumptions, which are not crucial for the results but 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 case it learns the latter's information and hence moves up one step on the ladder. Second, a partially informed island can meet either an uninformed one from its own productivity group, in which case it learns nothing itself, or a partially informed one from the *other* productivity group, in which case 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; 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 reveals to the island the particular subgroup from which its trading partner will be randomly drawn.

In this section we also find it useful to introduce a variable input in the production of the specialized goods. This input is meant to capture energy and materials, capital utilization, effort and labor hording, and any other variable input that may be unobserved by the econometrician.

As in the rest of the literature, this input helps capture cyclical movements in labor productivity and measured TFP (Solow residuals); but unlike the pertinent literature, here we will see that the movements could reflect purely extrinsic uncertainty.

We denote the unobserved input by e_{it} and henceforth referred to as capital utilization. We let the gross product of a firm be $y_{i,t} = An_{i,t}^{\vartheta_n}(e_{i,t}k_{i,t})^{\vartheta_k}$; and we specify the cost of this input in terms of the numeraire is $e_{i,t}^{1/\mu}k_{i,t}$, where $\mu \in (0,1)$. The realized profit of a firm is then $\pi_t = p_t y - r_{i,t}k_{i,t} - w_{i,t}n_{i,t} - e_{i,t}^{1/\mu}k_{i,t}$. This specification thus follows closely the modeling of capital utilization in King and Rebelo (2001), although here we do not need to be tied to this particular interpretation of $e_{i,t}$: all that matters for our purposes is that there is a variable input that the econometrician fails to measure and that may move endogenously over the cycle. Under this specification, the optimal level of this input is given by equating its marginal product with its marginal cost:

$$\frac{1}{\mu}{e_{i,t}}^{\frac{1}{\mu}-1}k_{i,t} = \vartheta_k \mathbb{E}_{i,t}\left[p_{i,t}\right]\frac{y_{i,t}}{e_{i,t}}.$$

Using this condition along with the equilibrium conditions given in Section 3, we can show that the equilibrium continues to be pinned down by Proposition 1, except that $\hat{\vartheta}$ must now be redefined as $\hat{\vartheta} \equiv \frac{\vartheta_n}{\epsilon} + (\mu + \psi(1-\mu))\vartheta_k$. Thus, modulo this adjustment in $\hat{\vartheta}$, our analysis remains intact.

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, 1998). There 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.

Let $\bar{\Omega}$ be the set of the aforementioned ten values. The dynamics of the cross-sectional distribution of types can be summarized in a simple law of motion for a vector $\mu_t \in \Delta(\bar{\Omega})$. 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 3. *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

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

(ii) The cross-sectional distribution of types follows

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

Part (i) gives a closed-form solution of the equilibrium output of 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 cross-sectional distribution of types. Together, these results permit us to determine any equilibrium outcome at either the island-wide or the economy-wide level.

Note that that variation in u, which causes variation only in the signals s_1 and s_2, does not cause variation in either the fundamentals or any of the agents' beliefs about the fundamentals. For fully informed islands, this is simply because these islands already know the entire aggregate fundamentals, and hence their beliefs are independent of those signals; for partially informed islands, this is because these islands know their own productivity and their beliefs about the productivity of other islands are pinned by the signals x_1 and x_2. Yet, variation in u causes self-fulfilling variation in equilibrium allocations and prices. In particular, islands of type ω_{P1+} and ω_{P2+} find it optimal to respond to the sunspot-like signals s_1 and s_2 along the unique equilibrium of the economy. This is because these signals 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. Similarly, a negative innovation in u causes a self-fulfilling recession. The shock u therefore plays the role of an aggregate "sentiment shock".

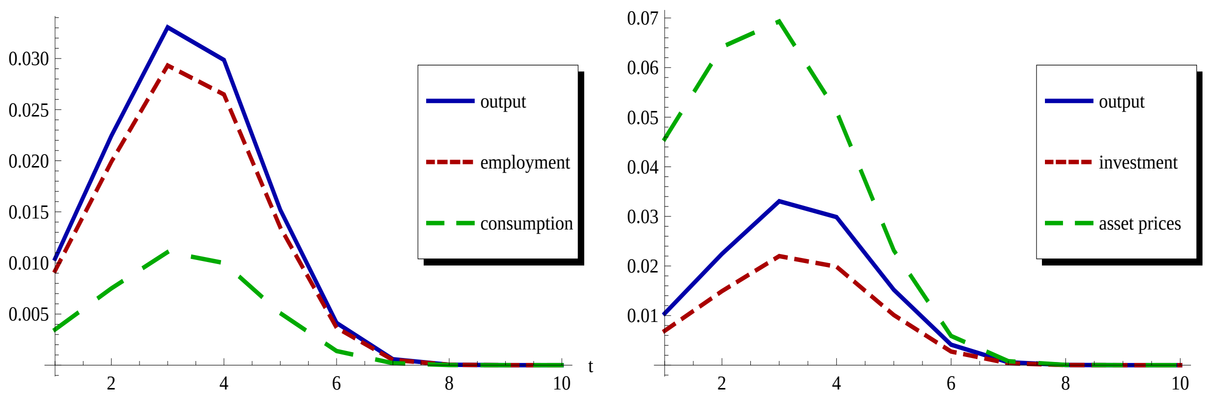

Figure 1: Impulse responses of macroeconomic variables to a positive sentiment shock.

Numerical example. We now illustrate the dynamic effects of a positive sentiment shock (that is, a positive innovation in u) with a numerical example.This example is not intended to represent a serious quantitative assessment; this would require different modeling choices. It is nevertheless based on a somewhat plausible paramaterization of our model.

In particular, we set $\beta = .99$, $\epsilon = 1.5$, $\vartheta_n = .7$, $\vartheta_k = .3$, which are plausible values for preferences and technologies. Next, we set $\eta = 1$ so as to capture the idea that the typical consumer only consume the good produced by their trading partners, not the good they produce themselves.[8] We also set $\psi = .5$, which guarantees that investment moves moderately over the cycle, and $\mu = .9$, which guarantees that capital utilization is highly variable. Finally, we let $\sigma_a = 0.25$ for the standard deviation of the island-specific productivity, and set the standard deviation of the noise in the x signal to $\sigma_\varepsilon = \sigma_a$ and the standard deviation of the u shock in the s signal, or the "sentiments" shock, to $\sigma_u = .5\sigma_a$; we do not have a particularly good justification for these values, except for the observation that they help us generate sizable fluctuations while not been completely crazy.[9]

In what follows, we then consider the impulse response of the economy to the following exogenous shock. At $t = 0$, a fraction $\chi = 30\%$ of the population becomes partially informed, while the rest are uninformed. After that initial shock, not other exogenous change occurs in the economy, either in the fundamentals or the information structure. All the rich dynamics we document below are thus the product of the endogenous communication that takes place in the economy over time.

Impulse responses. Figure 1 illustrates the dynamic effects of a positive innovation in u on the key macroeconomic variables: aggregate output, employment, consumption, investment, and asset prices. As evident in this figure, the innovation in u has a sizable and persistent effect on all these

[8]This assumption excludes most economists, who obviously get a lot of consumption from their own production!

[9]As we discuss in the end of this section, there is no reason to tie σ_a with the volatility of aggregate productivity. Rather, σ_a may be tied to cross-sectional heterogeneity, in which case a high value for σ_a would be possible.

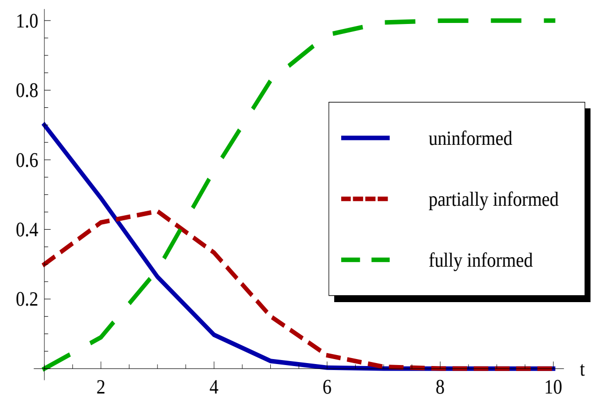

Figure 2: Population dynamics and the diffusion of information.

macroeconomic variables: notwithstanding the quantitative limitations of our exercise, we see that the boom lasts for about 6 or 8 quarters and that output increases by 3% at the peak of the boom. More importantly, all these aggregates move in the same direction, as it is the case in the data. What is more, the responses of these macroeconomic aggregates to the shock u are hump-shaped: they peak a 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 like the ones in Figure 1 appear to be important features of the data, at least as seen through the lenses of structural VAR models. Typically, generating such hump-shaped impulse responses in standard, representative-agent macro models requires the addition of habit formation, adjustment costs, and the other sources of inertia. Here, instead, they obtain quite naturally because of the endogenous diffusion of information.

To see this more clearly, Figure 2 illustrates the dynamics of μ_t, the cross-sectional distribution of types in the population. When the shock u hits the economy, only few agents react to it. This is because, by assumption, relatively few islands initially observe the signals s_1 and s_2 through which the shock u enters the economy. Over time, however, some of these partially informed islands meet, trade and communicate with uninformed islands. As the latter then turn into partially informed islands, they themselves start reacting to the sunspot-like signals s_1 and s_2. This explains why the fraction of partially informed islands tends to increase for some time after the initially shock hits the economy, which in turn explains why the impact of the shock on aggregate economic activity keeps building up for a while. At the same time, however, some of the partially informed islands meet, trade and communicate with other partially informed islands, thereby turning into fully informed. Thus, after some time, the fraction of partially informed types begins to decrease and, eventually, all islands reach full information. This explains why the impact of the shock starts fading off after

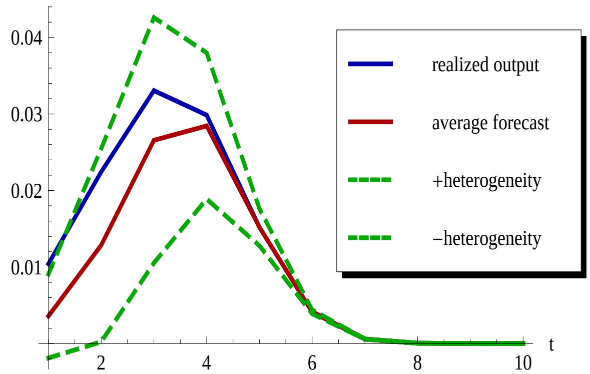

Figure 3: Self-fulfilling expectations and heterogeneity.

some point, and indeed vanishes eventually.

We now turn attention to the dynamics of the agents' expectations of economic activity. As we have discussed, the fluctuations we model here are meant to capture the self-fulfilling nature of short-run phenomena—the shock u is a modeling device that permits us to formalize how agents could suddenly get optimistic about the prospects of the economy, and how this optimism can be largely self-fulfilling. To illustrate this, Figure 4 relates the dynamics of realized aggregate output, given by the blue line, with the dynamics of the average forecast of aggregate output in the population, given by the red line. Evidently, the boom is associated with a self-fulfilling wave of optimism, that is, with a persistent, positive shift in expectations of economic activity. During the early phases of this wave, expectations become more and more optimistic over time, and actual output moves in tandem. Eventually, however, the optimism fades away, and actual output returns to "normal".

Because not all agents share the same information, the extent of optimism varies in the cross-section of the economy— there is heterogeneity in expectations of economic activity. This heterogeneity is illustrated by the two green dashed lines in Figure 3. These lines represent the "confidence band" around the average forecast that obtains by adding to, and subtracting from, the latter one standard deviation of the cross-sectional dispersion of forecasts. As evident from this figure, some agents may overestimate the boom in economic activity (these are the partially informed in our example), while others tend to underestimate it (these are the uninformed in our example). Furthermore, this heterogeneity varies over time; it tends to be higher in the early phases of the boom, and of course vanishes eventually as all agents become fully informed. This indicates how our model can accommodate the heterogeneity one observes in surveys of forecasts of economic activity, and perhaps even match the fact that this heterogeneity seems to pick in the early phases of the cycle.

We now consider the implications of our model for three other macroeconomic variables of interest: labor productivity, TFP as measured by the Solow residual, and the labor wedge. As

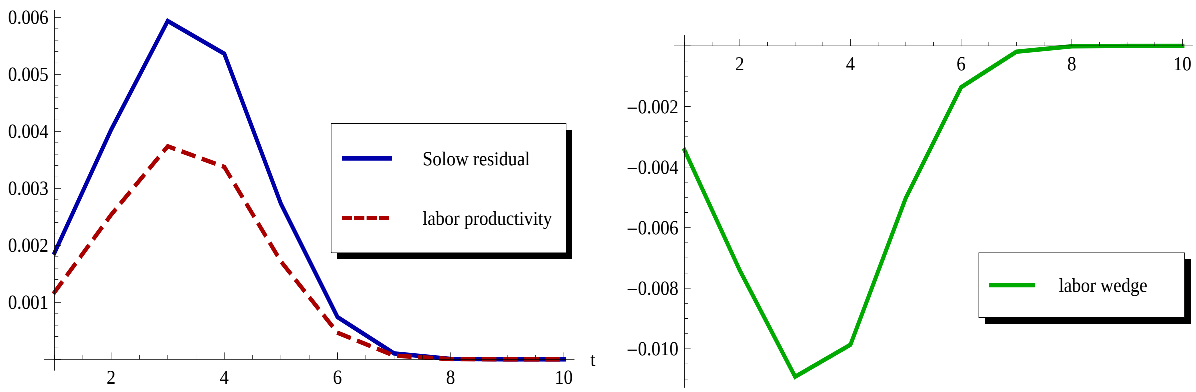

Figure 4: Labor productivity, Solow residual, and labor wedge.

mentioned previously, by introducing an unobserved variable input, we can now show how our waves can feature procyclical labor productivity and measured TFP. This is evident in the left panel of Figure 4, which plots the dynamic response of aggregate labor productivity and the Solow residual to a positive shock in u. The resulting wave of optimism has positive and persistent effects on both measures of productivity. It follows that both measures of productivity are procyclical in our model, much like in the data or in the standard RBC model. Importantly, though, these procyclical movements in productivity do not reflect exogenous variation in technology. Rather, they reflect the endogenous response of unmeasured inputs such as capital utilization and labor effort to the extrinsic uncertainty we have uncovered here.

Turning to the labor wedge, we denote it by τ_t and define it by the following identity:

$$N_t^{\epsilon-1} = (1-\tau_t)\,\vartheta_n \frac{Y_t}{N_t},$$

where Y_t and N_t are the aggregate output and employment levels. That is, the labor wedge τ_t is the implicit "tax" between the marginal product of labor and the marginal rate of substitution of consumption and labor, as seen through the lenses of a representative-agent model. The right panel of Figure 4 plots the impulse response of τ_t to the positive shock in u. As evident from this figure, the labor wedge is countercyclical: the "tax" is negative during booms. This is because, at least for the example considered here, the sentiment shock causes a stronger response in employment than in labor productivity: the boom in employment is larger than the one justified by the concurrent increase in labor productivity.[10]

[10] In fact, if we had not introduced capital utilization, we would necessarily have that the wedge is countercyclical: a positive shock in u would raise employment N, reduce labor productivity Y/N due to diminishing returns, and hence necessarily reduce the implied τ. Now, the endogenous response of capital utilization raises labor productivity, but this effect is not strong enough to undo the countercyclicality in the labor wedge.

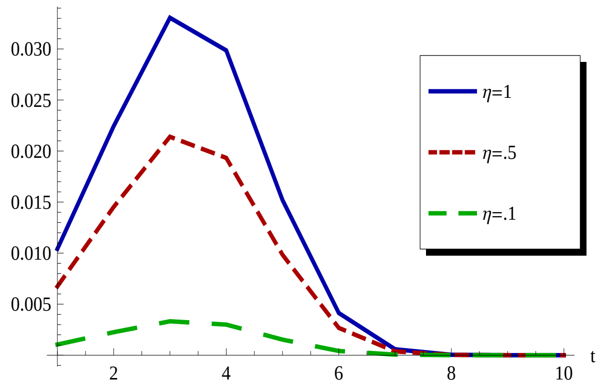

Figure 5: Interdependence (as measured by η) and the dynamic response of output.

Needless to say, countercyclical labor wedges are a central feature of actual business cycles (e.g., Hall, 1997; Chari, Kehoe and McGrattan, 2007; Shimer, 2009). Here, they emerge as the byproduct of the response of the economy to self-fulfilling waves of optimism and pessimism. An interesting direction for future work is then to examine how search frictions in labor markets—an important aspect of reality that we abstracted from—may interact with the extrinsic uncertainty we have introduce in this paper. Most likely, this will help explain why unemployment may vary over the business cycle largely because of self-fulfilling forces rather than because of shocks to the efficiency of firms or the "laziness" of workers.

To recap, the self-fulfilling fluctuations we have documented here appear to match a number of key qualitative features of actual business cycles. Whether a richer variant of our model would succeed also in the quantitative dimension is, of course, an important open question—a question, however, that is beyond the scope of the present paper.

Interpretation and discussion. While one may seek to find direct analogues in reality of either the sentiment shock u or the signals s_1 and s_2, our preference is to interpret these shocks and signals as modeling devices that permit us to capture a distinct form of uncertainty—the uncertainty that agents can face about economic activity and market prices *beyond* the uncertainty they face about the underlying fundamentals. It is this distinct, self-fulfilling type of uncertainty that the pertinent macroeconomics literature has failed to capture; it is the symptoms of this uncertainty that some observers may call "animal spirits" and may misinterpret as the product of irrational forces or multiple equilibria; and it this uncertainty that our analysis helps capture within otherwise conventional, unique-equilibrium, macroeconomic models.

In 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 "exuberance" thus comes and goes, first causing an boom and then a bust.

Clearly, this interpretation of the behavior of our economy is very close to the informal explanations of "bubbles" and other short-run phenomena that often dominate the media and the public arena. It is also reminiscent of some of the arguments that Akerlof and Shiller (2008), Krugman (2009), and Shiller (2005) use to make a case that short-run phenomena are driven by "animal spirits", "irrational exuberance", and other irrational forces that are absent in the standard paradigm. Furthermore, the fluctuations we document here are reminiscent of the ones that obtain in models with multiple equilibria: booms and recessions are "driven" by self-fulfilling variation in expectations of aggregate economic activity.

And yet, what lies beneath the surface is very different. According to our theory, the seemingly arbitrary waves of optimism and pessimism, and the associated booms and busts, have nothing to do with either irrationality or equilibrium indeterminacy. Rather, they are the symptoms of a distinct form of extrinsic uncertainty that emerges when, and only when, communication is imperfect.

At the same time, the perspective we offer here is all very different from the one that standard macroeconomic models offer. The informal arguments about "market psychology," "animal spirits", and the like, find no room within the context of these models, simply because all fluctuations are ultimately explained by shocks to technology, preferences, or other fundamentals. However, this does not mean that these arguments are completely baseless. As we have shown here, modern macroeconomic models have left no room for these notions—and more generally have failed to capture the distinct self-fulfilling nature of short-run phenomena—only because they have made the convenient, but unrealistic, assumption of perfect communication.[11]

Interestingly, the sentiment shocks we have formalized here can also be interpreted as "demand shocks". Let us elaborate. The archetypical Keynesian notion of "demand shocks" appeared to refer to some mystical forces that moved the economy away from what predicted by the neoclassical paradigm: "supply-side forces" explained the long run, "demand shocks" explained the short run. These ideas find little room in New-Keynesian models. The notion of "demand shocks" has been degenerated to the more mundane formal notion of preference shocks and their presence does not per se represents deviations form what predicted by the neoclassical paradigm. In contrast, our

[11]That being said, we cannot possibly understate the need for formalism: without a proper formalization, the aforementioned popular ideas can quickly be reduced to mere non-sense. As Robert E. Lucas, Jr. (2001) has warned us, *"Economic theory is mathematical analysis. Everything else is just pictures and talk."*

sentiment shocks are disconnected from preferences and technologies and cause fluctuations truly away from what predicted by the standard neoclassical paradigm. Furthermore, these fluctuations are necessarily of short-run nature, for they have to fade away eventually as agents communicate more and more over time. In these respects, our theory seems, not only to accommodate the archetypical notion of demand shocks, but also to do so in a more effective, and more faithful, manner than modern New-Keynesian models.

Finally, it is worth highlighting that some of the agents that "ride" our waves of optimism and pessimism may have a very good sense that these waves are disconnected from fundamentals. To see this more clearly, suppose that some of the fully informed agents in our preceed example 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 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.

Interdependence and systemic risk. We now study how the magnitude of the fluctuations we have documented depends on the extent of specialization and interdependence in the economy, as captured by η. Figure 6 illustrates the response of aggregate output to the same positive shock in u for three different values of η, namely $\eta \in \{.1, .5, 1\}$. Clearly, the magnitude of the resulting fluctuations is directly related to the strength of interdependence, as captured by η. As anticipated, this is because η controls how much each individual agent—or each individual island—cares about aggregate economic activity. In particular, as shown in Proposition 1, a higher η (and hence a higher α) means a higher sensitivity of individual economic activity to expectations of aggregate economic activity relative to fundamentals. When communication is perfect (information is symmetric), this ends up playing no role for aggregate fluctuations, simply because expectations of economic activity are ultimately pinned down by fundamentals. Indeed, one can easily check that, when information is symmetric, the response of economic activity to any shock in technologies or preferences is independent of the magnitude of η. But once communication is imperfect, η controls how sensitive economic activity is to the particular form of extrinsic risk we have identified in this paper: the more extensive the specialization and interdependence in the economy, the more each individual agent cares about any uncertainty he faces in economic activity *beyond* the one he faces in fundamentals.

The following interesting insight thus emerges. One may perhaps expect the processes of economic development, globalization, and financial sophistication to be conductive to more stability, and more efficiency. But as these processes appear to intensify the level of specialization and trading interdependence in the economy, our theory predicts that these processes make the economy more vulnerable to the particular form of systemic risk we have identified in this paper—and thereby more vulnerable to exotic forces that some may wish to call as "animal spirits", "market psychology", "irrational exuberance", and the like.

Finally, we seek to clarify that, whereas our results require lack of common knowledge about the underlying aggregate fundamentals, they do not require a high level of uncertainty about them. To see this more clearly, consider an N-replica of the economy we have studied in this section, with $N \geq 2$. Suppose that the productivity shocks (a_1, a_2) are independent across these economies, while the u shock is perfectly correlated across them. Suppose further that no trade takes place across these economies. It follows that each of these economies continue to behave in exactly the same way as in isolation. But now consider the behavior of the "mega-economy" that consists of the union of these N identical economies. Because the sentiment shock is perfectly correlated across them, the response of the mega-economy to this shock is exactly the same as in the example we have studied in this section, no matter the value of N. At the same time, because the productivity shocks are uncorrelated, the overall uncertainty in aggregate productivity decreases monotonically with N and vanishes as $N \to \infty$. We thus reach the following important observation.

Proposition 4. *The self-fulfilling fluctuations we have documented are consistent with arbitrarily small level of uncertainty in the underlying aggregate productivity.*

One may find the example we have used in order to prove this result to be too special or too ad hoc. For instance, one may question why the cross-sectional correlation in the sentiment shock may be higher than the one in productivities. A possible answer is that, in richer versions of our model, this correlation may itself be the product of communication. In fact, the preceding numerical example already embeds this idea: the sentiment shock initially hits only a small fraction of the population, and then spreads out in the rest of the economy through trade and communication.[12]

In any event, there is a broad, and powerful, insight that underlies Proposition 5. A high level of *aggregate* extrinsic uncertainty does not require a high level of uncertainty about the aggregate fundamentals. Rather, it can be sustained by the *idiosyncratic* uncertainty that each individual

[12] Whereas this fraction was set to 30% for our numerical example, it could be arbitrarily smaller without affecting the essence of our results. Other things equal, the smaller this initial fraction is, the longer it takes for the wave of optimism to build force and reach its peak. But if some of the initial few "optimists" have a more central position the economy, either in the sense that they trade with many other agents or in the sense that their voices are heard by many other agents, then their "optimism" could spread quickly in the economy.

agent may face about the idiosyncratic fundamentals and the likely actions of his future trading partners. In this sense, it suffices to have a sufficiently high level of cross-sectional heterogeneity.[13]

We conclude that there is no tight connection between the exogenous volatility in aggregate fundamentals that one may observe in the data and input in the model, and the endogenous volatility that our distinct form of extrinsic uncertainty can sustain. To put tight bounds on the latter, one would need more detailed evidence about the heterogeneity in fundamentals, the heterogeneity of information, and the network structure of the economy.

6 Extension: tradeable numeraire and risk sharing

We now consider a variant of our model that allows any two islands that get matched together to trade the numeraire good alongside the specialized goods. This extension serves two goals. First, it clarifies that our results are not sensitive to the simplifying assumption we made earlier on. Second, it 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, for the 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. The analysis 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:

$$p_{i,t} = P(y_{i,t}, y_{j,t}) \tag{17}$$

for some function $P : \mathbb{R}^2_+ \to \mathbb{R}_+$. This condition is the analogue of condition (7) in the baseline model. The only difficulty is that P may not have a closed-form solution. modulo this difference, the rest of the equilibrium characterization is identical to the one in the baseline model. We thus reach the following variant of Proposition 2 and 3.

[13]Indeed, in the aforementioned N-replica example, the value of σ_a is tied to the magnitude of cross-sectional heterogeneity as $N \to \infty$

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:*

$$\mathbb{E}\left[G\left(y(\omega_{i,t}), A(\omega_i), y(\omega_{j,t}) \right) \middle| \, \omega_{i,t} \right] = 0 \tag{18}$$

Condition (18) is the analogue of condition (14) in the baseline model: it gives equilibrium output 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 (7); now it is captured in an implicit way by condition (17). This, however, does not affect the core of our positive results: sunspot fluctuations continue to be possible if and only if information is asymmetric.[14]

7 Policy and Efficiency

The preceding analysis has documented how dispersed information and imperfect communication can open the door to a novel type of systemic risk, one that can help capture the informal notions of "animal spirits", "market psychology," and the like. This pegs two policy 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.

Consider first the positive question. The fact that the equilibrium of our economy is unique implies that policy analysis remains, at least conceptually, almost as simple as in any other unique-equilibrium macro models. For example, a tax that increases with aggregate economic activity can induce economic agents to react less to any shock the correlates with aggregate economic activity, whether this shock is fundamental or non-fundamental. More sophisticated policy interventions can be done along the lines considered in the more abstract work of Angeletos and Pavan (2009): one can devise retroactive state-contingent taxes that eliminate the sentiment-driven fluctuations in our model without reducing the sensitivity of the economy to fundamentals. Finally, in monetary variants of our model, a similar role could be played by monetary policy.

The point we are making here is conceptually simple. But it is worth highlighting how this contrasts with competing formalizations of the notions of "self-fulfilling fluctuations", "animal spirits"

[14]This is evident in the following special case. Suppose $U(c, c^*) = (c^{1/2} c^{*1/2})^\gamma$ and $V(n) = n^\epsilon$, where $\gamma < 1 < \epsilon$. Then, it is easy to check that $P(y, y^*) = y^{-\gamma/2} y^{*\gamma/2}$ and, by implication, there exist scalars $\varphi > 0$ and $\alpha \in (0, 1)$ such that condition (18) can be restated as $\log y(\omega_{i,t}) = (1-\alpha)\varphi \log A(\omega_i) + \alpha \mathbf{E}\left[\log y(\omega_{j,t}) | \, \omega_{i,t} \right]$. modulo 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 documented in the previous section can be re-casted *exactly* in this special case.

and the like. In models with multiple equilibria, the response of the economy to conventional fiscal or monetary policies is often indeterminate. Policy analysis then may have to rest either on ad hoc equilibrium selections, or with strong off equilibrium effects. In models with irrational agents, on the other hand, one could justify very different policies depending on what one assumes about the how these agents react, or fail to react, to any particular policy intervention. Either way, one is left with little discipline on policy analysis.[15] In contrast, our framework retains, not only the convenience, but also the discipline that standard unique-equilibrium, rational-expectations macroeconomic models provide for policy analysis.

We now turn to the normative question, namely whether stabilization policy is desirable in our economy. 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. We thus consider a different efficiency concept, one that subjects the planner under the same communication constraints as those imposed on the market mechanism.[16]

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:*

$$y_t(\omega_{i,t}) = A(\omega_i) k_t(\omega_{i,t})^{1-\vartheta} n_t(\omega_{i,t})^{\vartheta} \tag{19}$$

$$k_{t+1}(\omega_{i,t+1}) = \tilde{n}_t(\omega_{i,t+1})^{\psi} \tag{20}$$

$$c_t(\omega_{i,t+1}) + c_t^*(\omega_{j,t+1}) = y(\omega_{i,t}) \tag{21}$$

where j stands for trading partner of island i.

[15] This lack of discipline is eloquently pointed out in N. Gregory Mankiw's (2009) reaction to Robert Shiller's (2009) claim that a big fiscal stimulus would raise confidence in the economy:

> *"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'."* Mankiw (2009)

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

(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,*

$$\omega_{i,t+1} = (\omega_{i,t}, \omega_{j,t}) \quad \textit{and} \quad \omega_{j,t+1} = (\omega_{j,t}, \omega_{i,t}). \tag{22}$$

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 (19)-(22).

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 (21). Using (19) and (20) 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:

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

where $\mathcal{V}(\omega_{i,t+1})$ 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.

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.

Once again, this result is in sharp contrast to models with multiple equilibria or irrational agents, where the need for government intervention is hard-wired. It also unsettles conventional wisdom: even if we take for granted that most short-run volatility is disconnected from fundamentals, and perhaps driven by mysterious forces that one may call "market sentiments," "animal spirits", or "demand shocks", this fact does not *by itself* provide a rationale for stabilization policy.

8 Conclusion

Standard macroeconomic models attribute the bulk of short-run fluctuations to high-frequency disturbances in preferences and technologies. We find this highly unsatisfactory. If taken literally, these shocks seem empirically implausible. Instead, short-run phenomena appear to have a largely self-fulfilling nature—one that leads many practitioners to attribute these phenomena to more exotic forces such as "animal spirits", "sentiments", or "market psychology", and one that standard macroeconomic models have failed to capture. As Kocherlakota (2009) puts it more eloquently,

> *"[Macroeconomists] 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.*

Of course, many macroeconomists may wish to see these shocks only as "as convenient shortcuts to generate the requisite levels of volatility in endogenous variables" (Kocherlakota). But doing so is anything but reassuring. How could one insist on the need for precise micro-foundations, and then let the most core aspect of these models—the origins of fluctuations—be a black box of unknown consequences? And how could one then trust the policy guidance provided by these models?

Note, in particular, that the efficiency of real business cycles and the optimality of price stability—the most central normative benchmarks in modern macroeconomic theory—crucially rest on a literal interpretation of the preference and technology shocks featured in these models. How could we trust these results if these shocks were only proxies for the more mystical forces others prefer to associate with "animal spirits", "market pshycology", "consumer confidence", and the like? And how could we respond to the critiques raised by Krugman, Akerlof, Shiller and others that the need for stabilization policies becomes self-evident once one recognizes the importance of these forces, forces that move the economy away from what captured by the neoclassical paradigm?

Or, to give a more topical example, consider the desirability of expansionary fiscal policy during the ongoing recession. Recent work has tried to capture the recession with a negative preference shock that triggers the natural (efficient) real interest rate to turn negative, forces the nominal interest rate to hit the zero bound, and keeps the equilibrium real interest rate above its natural value. This work then proceeds to show that it is optimal for the government to bring the real interest rate down by raising government spending—hence providing a normative basis for the recent fiscal stimulus. At the same time, most would readily acknowledge that the aforementioned preference shock is only a short-cut, probably one trying to capture the more mystical notion of a negative "demand shock". But if this preference shock is not to be taken literally, doesn't the entire normative basis for the fiscal stimulus collapse?

These observations define the context and scope of our contribution. In this paper, we are not interested in studying yet another canonical macroeconomic model, of either RBC or the New-Keynesian variety. Rather, we revisit the very foundations of this class of models and seek to develop an alternative formalization of the origins of short-run fluctuations. This formalization shifts the focus away from the usual preference and technology shocks and, instead, identifies an entirely novel form of uncertainty—one that, at least in our view, is endemic to the market mechanism and helps capture an important aspect of reality.

At the core of our theory is a simple, and powerful, idea. Any real-world market economy is characterized by a high level of specialization in economic activity and a complicated network of decentralized market interactions. Furthermore, while these market interactions may facilitate significant communication throughout the economy, this communication is unlikely to be instantaneous and perfect. As long as this is the case, the uncertainty agents may face about macroeconomic activity—for example, the uncertainty firms may face about consumer demand, the one consumers face about labor market conditions, or the one investors face about firm profitability—is not spanned by the uncertainty they face about the underlying preferences, technologies and other fundamentals. It is this residual, extrinsic uncertainty that our theory seeks to bring to the forefront of macroeconomic research.

This extrinsic uncertainty has very similar flavor as the one encountered in models with multiple equilibria: it captures the self-fulfilling nature of short-run fluctuations. Importantly, though, it does not rest on the severe externalities, non-convexities and missing markets that are most often needed to sustain multiple equilibria. Rather, it rests only on the imperfection of the communication that is likely to take place through the market mechanism.

Furthermore, this extrinsic uncertainty seems to capture what is the core *positive* aspect of the archetypical Keynesian view of fluctuations: fluctuations are attributed to mysterious "demand shocks" that are disconnected from preferences and technologies, that move the economy away from "supply-side forces", and that tend to get "multiplied" as agents react to them. And yet, the *normative* lesson that comes from our analysis is anything but Keynesian: the extrinsic uncertainty we uncover in this paper is a symptom of constrained inefficiency, leaving no room for conventional stabilization policies. What, instead, our theory does is to open the door to non-conventional policies that seek to facilitate more communication in the economy, either by providing more information directly, or perhaps by helping centralize market interactions.

To deliver these insights in a transparent manner, we considered a highly stylized model of the macroeconomy that assumed away, among other things, any frictions in financial, labor and product markets. While we certainly think that these frictions play a central role in the real world, abstracting from them helped provide a clean methodological benchmark. We hope that this benchmark will in turn fuel more research in the direction of models that better capture the self-fulfilling nature of short-run phenomena. In particular, an obvious direction for future research is to study how the aforementioned frictions may interact with the form of extrinsic uncertainty we have identified here. This will almost surely lead to novel policy insights, while at the same time helping develop a more complete picture of the self-fulfilling nature of business cycles, credit crunches, asset-price bubbles, and other short-run phenomena.[17]

[17] None of this, however, would vindicate the recent calls by Krugman (2009) and others for an all-around return to old Keynesian thinking. 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 hope that we have also pushed the frontier of macroeconomic research.

Appendix

Proof of Proposition 1. When $\delta = 1$, condition (13) reduces to $q_{i,t} = \beta \mathbb{E}_{i,t+1}[r_{i,t+1}]$ and $x_{i,t} = k_{i,t}$. Along with (11), this gives

$$k_{i,t}{}^{\frac{1}{\psi}} = \beta\psi\vartheta_k \mathbb{E}_{i,t}\left[p_{i,t}\right] y_{i,t}. \tag{23}$$

This condition has an analogous interpretation as that of condition (10) for labor: it equates the marginal private costs and benefits of capital.

Substituting the optimality condition for capital (23) 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 (10), replacing $p_{i,t}$ in the latter from (7), 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 (14). *QED*

Proof of Proposition 2. Iterating condition (14) gives output as a function of the hierarchy of beliefs:

$$\log y_{i,t} = (1-\hat{\alpha})\left\{f_i + \hat{\alpha}\mathbf{E}_{i,t}[f_j] + \hat{\alpha}^2\mathbf{E}_{i,t}[\mathbf{E}_{j,t}[f_k]] + ...\right\}, \tag{24}$$

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. This follows from the discussion in the main text.

Proof of Proposition 3. *Part (i).* In the proposed equilibrium, the period-t output of island j is log-normally distributed conditional on the information of island i, for any i, j, and t. It follows that the non-linear expectation $\mathbf{E}_{i,t}y_{j,t}$ and the simple expectation $\mathbb{E}_{i,t}y_{j,t}$ are equal to each other up to a constant that we henceforth ignore. We can thus rewrite the key equilibrium condition as

$$\log y\left(\omega_i\right) = (1-\hat{\alpha})\frac{1}{1-\hat{\vartheta}}\,a_i + \hat{\alpha}\,\mathbb{E}_{it}\left[\log y\left(\omega_{j,t}\right)\right] \tag{25}$$

Using this condition, we now consider the equilibrium outputs for each of the ten types $\omega_{i,t}$, by considering the equilibrium outcomes for all possible matches.

- Matches between two islands of type ω_{U1}. In this case, equilibrium output of type ω_{U1} must satisfy $\log y\left(\omega_{U1}\right) = (1-\hat{\alpha})\frac{1}{1-\hat{\vartheta}}\,a_1 + \hat{\alpha}\,\log y\left(\omega_{U1}\right)$. It follows that $\log y\left(\omega_{U1}\right) = \phi_a a_1$ for $\phi_a = \frac{1}{1-\hat{\vartheta}}$. 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 $\log y(\omega_{U1+}) = \phi_{0U} a_1$ for some coefficient ϕ_{0U} and $\log y(\omega_{P1}) = \phi_{0P} a_1 + \phi_x x_1 + \phi_s s_1$, for some coefficients $(\phi_{0P}, \phi_x, \phi_s)$. It follows that $y(\omega_{U1+})$ and $y(\omega_{P1})$ are indeed log-normal, with

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

where$\mathbb{E}\left[x_1|\omega_{U1+}\right] = \mathbb{E}\left[s_1|\omega_{U1+}\right] = 0$. Substituting these expressions into (25) gives us

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

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

- 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 $\log y(\omega_{P2+}) = \phi_0 a_2 + \phi_x x_2 + \phi_s s_2$, for some coefficients (ϕ_0, ϕ_x, ϕ_s), for $i = 1, 2$. It follows that $y(\omega_{P2+})$ is indeed log-normal, with

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

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 (25) gives us

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

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

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

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 $\log y(\omega_{F1}) = (1-\hat{\alpha})\frac{1}{1-\hat{\vartheta}} a_1 + \hat{\alpha} \log y(\omega_{F1})$. It follows that $\log y(\omega_{F1}) = \phi_a a_1$ for $\phi_a = \frac{1}{1-\hat{\vartheta}}$. A similar result holds for matches between two islands of type ω_{F2}.

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.

$$\mu_t = \begin{bmatrix} \mu_{U1} & \mu_{U1+} & \mu_{P1} & \mu_{P1+} & \mu_{F1} & \mu_{U2} & \mu_{U2+} & \mu_{P2} & \mu_{P2+} & \mu_{F2} \end{bmatrix}'$$

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 4. This follows from the main text.

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

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

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 (17) gives condition (18), 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 Theorem 2. This follows from the discussion in the main text.

References

[1] Akerlof, George, and Robert Shiller (2008), *Animal Spirits*, Princeton University Press.

[2] Amador, Manuel, and Pierre-Olivier Weill (2009), "Learning from Prices: Public Communication and Welfare," Stanford University/UCLA mimeo.

[3] Angeletos, George-Marios, and Jennifer La'O (2009), "Noisy Business Cycles," NBER Macroeconomic Annual.

[4] Angeletos, George-Marios, and Alessandro Pavan (2009), "Policy with Dispersed Information," *Journal of the European Economic Association* 7, 11-60.

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

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

[7] Barro, Robert (1976), "Rational Expectations and the Role of Monetary Policy," *Journal of Monetary Economics* 2, 1-32.

[8] Barsky, Robert, and Eric Sims (2009), "News Shocks," University of Michigan mimeo.

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

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

[11] Benhabib, Jess, and Roger Farmer (1999), "Indeterminacy and Sunspots in Macroeconomics," in John Taylor and Michael Woodford, eds., *Handbook of Macroeconomics*, Amsterdam: North-Holland.

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

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

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

[15] Cochrane, John (2009), "How did Paul Krugman get it so Wrong?," (http://faculty.chicagobooth.edu/john.cochrane/research/Papers/krugman_response.htm)

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

[17] Duffie, Darrell, and Gustavo Manso (2007). "Information Percolation in Large Markets, " *American Economic Review* 97, 203-209.

[18] Duffie, Darrell, Gaston Giroux, and Gustavo Manso (2009). "Information Percolation. " *American Economic Journal: Microeconomics* 2.

[19] Farmer, Roger (1993), *The Macroeconomics of Self-fulfilling Prophecies*, Cambridge, MA: MIT Press.

[20] Farmer, Roger, and Michael Woodford (1997), "Self-fulfilling Prophecies and the Business Cycle," *Macroeconomic Dynamics* 1, 740-769.

[21] Hassan, Tarek, and Thomas M. Mertens (2010), "The Social Cost of Near-Rational Investment: Why we should worry about volatile stock markets," University of Chicago/NYU mimeo.

[22] Hellwig, Christian (2005), "Heterogeneous Information and the Welfare Effects of Public Information Disclosures," UCLA mimeo.

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

[24] Kiyotaki, Nobuhiro (1988), "Multiple Expectational Equilibria under Monopolistic Competition," *Quarterly Journal of Economics* 103, 695-713.

[25] King, Robert (1982), "Monetary Policy and the Information Content of Prices," *Journal of Political Economy* 90, 247-279.

[26] Krugman, Paul (2009), "How Did Economists Get It So Wrong?," *New York Times*, September 2, 2009. (http://www.nytimes.com/2009/09/06/magazine/06Economic-t.html)

[27] Krusell, Per and Anthony Smith (1998) "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, vol. 106(5), 867-896.

[28] Lorenzoni, Guido (2010), "A Theory of Demand Shocks," *American Economic Review*, forthcoming.

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

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

[31] Mackowiak, Bartosz, and Mirko Wiederholt (2009), "Optimal Sticky Prices under Rational Inattention," *American Economic Review*, 99, 769-803.

[32] Mankiw, N. Gregory and Ricardo Reis (2002), "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve," *Quarterly Journal of Economics* 117, 1295-1328.

[33] Mankiw, N. Gregory (2009), blog entry entitled "Shiller on Animal Spirits" (http://gregmankiw.blogspot.com/2009/01/shiller-on-animal-spirits.html).

[34] Mankiw, N. Gregory, and Ricardo Reis (2006), "Pervasive Stickiness," *American Economic Review* 96, 164-169.

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

[36] Phelps, Edmund (1970), "Introduction: The New Microeconomics in Employment and Inflation Theory," in A.A. Alchian, C.C. Holt et al., *Microeconomic Foundations of Employment and Inflation Theory,* New York: Norton.

[37] Rondina, Giacomo (2008), "Incomplete Information and Informative Pricing," UCSD mimeo.

[38] Sims, Christopher (2003), "Implications of Rational Inattention," *Journal of Monetary Economics* 50, 665-690.

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

[40] Shiller, Robert (2009), "Animal Spirits Depend on Trust: The proposed stimulus isn't big enough to restore confidence," *The Wall Street Journal,* January 27, 2009, A15. (http://online.wsj.com/article/SB123302080925418107.html).

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

[42] Veldkamp, Laura (2009), *Information Choice in Macroeconomics and Finance*, book in progress.

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

[44] Woodford, Michael (1991), "Self-Fulfilling Expectations and Fluctuations in Aggregate Demand," in N. Gregory Mankiw and David Romer, eds., *New Keynesian Economics.*