Synchronization in Sunspot Models

Patir, Assaf

The Hebrew University of Jerusalem

25 August 2019

Synchronization in Sunspot Models

Assaf Patir*

April 25, 2020

Abstract

This paper illustrates how agents' beliefs about economic outcomes can dynamically synchronize and de-synchronize to produce business-cycle-like fluctuations in a simple macroeconomic model. We consider a simple macroeconomic model with multiple equilibria, which represent different ways in which sunspots can forecast future output in a self-fulfilling manner. Agents are assumed to learn to use the sunspot variable through econometric learning. We show that if different agents have varied interpretations of the sunspot, this leads to a complex nonlinear dynamic of synchronization of beliefs concerning the equilibrium played. Depending on the extent of disagreement on the interpretation of the sunspot, the economy will be more or less volatile. The dispersion of agents' beliefs is inversely related to volatility, since low dispersion implies that output is very sensitive to extrinsic noise (the sunspot). When disagreement crosses a critical threshold, the sunspot is practically ignored and output is stable.
The model naturally generates stochastic volatility (although there is no aggregate uncertainty), and explains features found in surveys of professional forecasters.

1 Introduction

This paper has two goals: first, to analyze, in a model with strategic uncertainty, how a sunspot can emerge as a coordination mechanism, even when agents do not agree on what the sunspot is; and, second, to show that the perpetual learning process leads to complex dynamics in terms of agents' beliefs, whereby agents are more or less coordinated on the effect of the

*Department of Economics, The Hebrew University of Jerusalem, email:assaf.patir@mail.huji.ac.il

sunspot and, as a result, both average output and its volatility fluctuate over time.

Sunspots are commonly used in macroeconomics to facilitate correlated equilibria (Cass and Shell, 1983; Aumann et al., 1988). The literal interpretation is that agents commonly observe some extrinsic source of randomness, and incorporate it into their decision-making process in a manner that generates an actual law of motion equivalent to their perceived law of motion. Clearly, there are plenty of processes in the world that are random – or at least quasi-random for all practical purposes – and can therefore be used as sunspots, but is such behaviour ever seen in the real world? Indeed, if economists assert that a certain real-life system is well described by a sunspot model, how is it that the sunspot cannot be identified?

Researchers have looked for evidence of the existence of sunspots in empirical data (e.g. Benhabib and Spiegel, 2019; Nayar and Levchenko, 2017) and in laboratory experiments (e.g. Fehr et al., 2019), but no one has suggested a specific concrete variable that might be termed a sunspot.

Of course, one should not be so literal when interpreting a model: a more plausible interpretation is that agents are continuously observing many stochastic processes that may or may not be related to economic fundamentals, and temporary coordination emerges somehow as an equilibrium result. This interpretation, however, poses a new difficulty: if the particular source of stochasticity is not clearly defined, how can agents learn to coordinate on it? Sunspots in such equilibria are, by definition, only related to outcomes because agents are coordinated in using them. Therefore, unlike signals about fundamentals, it is far less clear if agents can understand its relationship to economic outcomes unless other agents are already coordinated. Previous literature has assumed that agents know that some process, call it $\{z_t\}$, is potentially relevant to outcomes, and have analyzed the circumstances under which they can learn to use it,[1] but this might be considerably more difficult when they do not know what variables are potential sunspots and perhaps have different preconceptions about the source of randomness.

To put this in more formal terms, in the above mentioned papers, and in most sunspot applications, the sunspot process is assumed to be a well-defined iid process. It could be something like the intensity of actual sunspots (i.e. on the sun) at 8AM every day. However, another agent may think that the sunspot is the same measurement performed at a fixed time after sunrise. Both stochastic series would be iid with the same distribution as infinitely

[1] For example, see Woodford (1990); Guesnerie and Woodford (1990); Evans et al. (1994); Evans and Honkapohja (2003a,b); Honkapohja and Mitra (2004).

many other possibilities, and in the non-literal interpretation of sunspots, where the identity of the sunspot is not clearly communicated, it is hard to imagine that agents will not have such different interpretations. Thus, the problem of learning to coordinate on sunspots is, primarily, the problem of learning to construct a stochastic process out of many sources of randomness.

This paper demonstrates how, despite the above, it *is* perfectly possible for agents to learn to coordinate, even if they do not agree at any given moment on what the coordination mechanism is. To do this, we consider a system with a continuum of sunspot equilibria that are distinguished by different choices of the sunspot, and show that agents who are using a simple learning rule can converge on identifying an outcome similar to an equilibrium, even while not agreeing on the identity of the sunspot. As the level of disagreement increases up to some critical level, agents put less weight on the sunspot and, beyond the critical level, the sunspot is ignored. This will be defined precisely in the body of the paper.

Beyond proving that such imperfect coordination can emerge, the paper seeks to argue that the dynamics that result have many features that are actually desirable. The analysis demonstrates that the learning process leads to complicated non-linear dynamics in the agents' belief-space, whereby beliefs go through periods of synchronization and of desynchronization, and consequently the strategies neither converge on an equilibrium nor diverge. Instead, there are long periods in which agents' behaviour resembles a sunspot equilibria, and then periods where the agents are not coordinated on the use of the sunspot. Consequently, the economy goes through periods of high volatility (when the agents are coordinated on the sunspot) and periods of low volatility (when the dis-coordination leads to agents' actions cancelling out and output being roughly constant). Thus, this paper is a model of endogenously generated stochastic volatility.

In addition, agents' beliefs display features that are observed empirically in surveys of professional forecasters, and which are not easy to explain in standard models. Specifically, we find that the individual uncertainty about future output has a strong individual fixed-effect; and that output is negatively correlated with both individual uncertainty and dispersion of beliefs (although less strongly with the latter). These are exactly the features that are observed in surveys of professional forecasters (Abel et al., 2016; Bloom, 2014; Dovern et al., 2011; Jurado et al., 2015; Rich and Tracy, 2018, 2010; Mankiw et al., 2004; Ludvigson, 2016; Patton and Timmermann, 2010).

While the results of the paper can easily be adapted to any sunspot model, for concreteness the focus here is on a modified version of a model

from Benhabib, Wang, and Wen (2015). In the model, firms are trying to learn about the relationship between k publicly observed stochastic processes $(z_t^i, i = 1, \ldots, k)$ and total output (y_t). There are specific linear combinations, $y_t = \phi + \xi \cdot z_t$, such that if all agents believed that output fluctuates according to this formula, it would be self fulfilling, and these are the stable equilibria of the model. However, I allow agents to have different notions of the mapping between z_t^i and z_{t+1}^i, and this difference, combined with the learning process, leads to the complex dynamics in the belief space that are described above.

The organization of this paper is as follows: after reviewing some of the relevant literature in 2, the basic model is presented in section3. The analysis of the model in section 4 is divided into three parts: subsection 4.1 derives the dynamic-stochastic equations of the model. Subsection 4.2 shows how exact analytical solutions can be obtained within a special limit, and the properties of these solutions are analyzed in 4.3 using a generalization of impulse-response-functions. Though the exact solutions are only valid within the special limit, numerical analysis discussed in 4.4 suggests that important qualitative features of these solutions remain relevant in the general case. Section 5 discusses the economic implications and the empirical evidence supporting our model. Some additional comments are left to section 6, and appendix A contains some of the mathematical derivations.

2 Related Literature

Learning has a long history in macroeconomics, but the stochastic recursive description in this paper originates from Marcet and Sargent (1989). For a comprehensive account of the state of this field see Evans and Honkapohja (2012). Some papers concerning the presence of multiple equilibria and sunspots are Woodford (1990); Guesnerie and Woodford (1990); Evans et al. (1994); Evans and Honkapohja (2003a,b); Honkapohja and Mitra (2004).

Traditionally, the term *sunspot* is used in macroeconomics to describe a situation where the dynamic equations of a system lead to indeterminacy, and therefore a new stochastic process, the sunspot, can be introduced on which agents coordinate their actions (for example Benhabib and Farmer, 1994; Christiano and Harrison, 1996). In these models the realization of the stochastic process determines the equilibrium being played. In a more recent paper Angeletos and La'O (2013) describe a different situation where there is a unique equilibrium in which the agents use a random variable that they call *sentiment* to choose their actions. While similar in spirit, these are

formally different situations. This paper makes use of the model introduced by Benhabib et al. (2015), which is similar to the latter in that the role of the stochastic process is not to choose between equilibria.

As mentioned above, our model can be seen as a theory of endogenous stochastic volatility. A recent review of this topic is by Fernández-Villaverde and Guerrón-Quintana (2020).

Subsection 5.1 gives empirical support to the model from the empirical literature on forecasting, focusing on papers that analyze surveys of professional forecasters (e.g. Abel et al., 2016; Bloom, 2014; Dovern et al., 2011; Jurado et al., 2015; Rich and Tracy, 2010, 2018; Mankiw et al., 2004; Ludvigson, 2016; Patton and Timmermann, 2010). Other mechanisms to generate the same phenomena are discussed in Bloom (2014), Fajgelbaum et al. (2017) and Zohar (2020).

The dynamics of the model are closely related to the Kuramoto model (Kuramoto, 1975), which has been used to describe synchronization phenomena across different disciplines and subject areas including synchronization of flashing fireflies, phase lock in metronomes, and synchronized applause at the end of a concert. The Kuramoto model describes a set of oscillators whose phases are nonlinearly coupled, not unlike how the learning process in our model links agents' beliefs about the equilibrium being played. This is, to our knowledge, the first time that this link has been made, and potentially opens the door to incorporating into macroeconomics the rich phenomena that the Kuramoto model can describe. Thorough introductions to the model and reviews of the current state of the literature include Strogatz (2000) and Acebrón et al. (2005). The full model in this paper can be seen as a Kuramoto model with three modifications: stochastic coupling, nonlinear corrections, and amplitude interactions. Similar types of modifications have been studied in previous literature as extensions of the Kuramoto model: multiplicative stochastic coupling has been studied in Park and Kim (1996), generalized nonlinear interactions between phases in (Daido, 1993, 1994, 1996a,b), and amplitude interactions in (Ermentrout, 1990; Matthews and Strogatz, 1990; Matthews et al., 1991).

3 Model Setup

The model setup is based on Benhabib, Wang, and Wen (2015).

3.1 Households and Firms

3.1.1 Households

A representative household values streams of consumption $C_t \geq 0$ and labor $N_t \geq 0$ according to

$$U = \sum_{t=0}^{\infty} \beta^t [\log C_t - \psi N_t], \quad \beta \in (0,1), \quad \psi > 0,$$

and is subject to the budget constraint

$$P_t C_t \leq W_t N_t + \Pi_t,$$

where P_t, W_t and Π_t are the prices of the consumption good, the nominal wage, and the profits from ownership of firms, respectively.

The household's first-order conditions are

$$C_t = \frac{1}{\psi} \cdot \frac{W_t}{P_t}, \qquad N_t = \frac{1}{\psi} - \frac{\Pi_t}{W_t}. \tag{1}$$

3.1.2 Final Good Producers

The consumption good is produced by competitive final good producers using a continuum of intermediate goods indexed by $j \in [0,1]$, with the stochastic technology

$$Y_t = \left[\int_0^1 \epsilon_{jt}^{\theta} Y_{jt}^{1-\theta} dj \right]^{\frac{1}{1-\theta}}, \quad \theta > 0 \tag{2}$$

where ϵ_{jt} are iid random variables, and can be interpreted as preference shocks. We shall assume throughout that $\log \epsilon_{jt} \sim \mathrm{N}(0, \sigma_\varepsilon^2)$.

Denoting the price of good j at time t by P_{jt}, the demand for intermediate good j is given by

$$\left(\frac{Y_{jt}}{Y_t} \right)^{\theta} = \frac{P_t}{P_{jt}} \varepsilon_{jt}^{\theta}.$$

From this we also get the relationship:

$$P_t^{1-1/\theta} = \int_0^1 \epsilon_{jt} P_{jt}^{1-1/\theta} dj.$$

3.1.3 Intermediate Goods Producers

Each variety of intermediate good j is manufactured by a monopolist using labor as the only input and with the production function: $Y_{jt} = AN_{jt}$. The intermediate good manufacturers must decide on their level of production simultaneously at the beginning of the period without observing the shocks ϵ_{jt}. After these decisions have been made, prices are set so that markets clear, similarly to a Cornout competition.

The intermediate good producers' problem is therefore

$$\max_{Y_{jt}} \mathbb{E}_{jt}[(P_{jt} - W_t/A)Y_{jt}],$$

where $\mathbb{E}_{jt}$ represents the firm's expectation operator conditioned on the information (and beliefs) available to firm j at time t, which will be described below. The first-order-condition is

$$Y_{jt} = \mathbb{E}_{jt}\left[A(1-\theta)\frac{P_t}{W_t}Y_t^{\theta}\epsilon_{jt}^{\theta}\right]^{1/\theta}.$$

Substituting (1) into the above, we get

$$Y_{jt} = \mathbb{E}_{jt}\left[\frac{A(1-\theta)}{\psi}Y_t^{\theta-1}\epsilon_{jt}^{\theta}\right]^{1/\theta} = \mathbb{E}_{jt}\left[Y_t^{\theta-1}\epsilon_{jt}^{\theta}\right]^{1/\theta},$$

where in the last step, without loss of generality, units of output are chosen such that $\psi = A(1-\theta)$. Finally, it is useful to redefine $y_t = \log Y_t$, and $\varepsilon_{jt} = \log \epsilon_{jt}$ so that we have

$$y_{jt} = \frac{1}{\theta}\log \mathbb{E}_{jt}\left[e^{\theta\epsilon_{jt}-(1-\theta)y_t}\right]. \tag{3}$$

3.2 Information

A large number of 'forecasters' observe both firm specific shocks ϵ_{jt} and a "sunspot" variable z_t. The process $\{z_t\}_{t=0}^{\infty}$ is a standard Gaussian white noise vector: that is, for all t, $z_t \in \mathbb{R}^k$ is multivariate normal $\mathrm{N}(0, I_k)$ $(k > 1)$, and independent across t.

The intermediate-good firms do not get to see z_t directly. Instead, they rely on a survey of forecasters to estimate their demand curves. However, the firm is limited in its ability to conduct market research, so it eventually obtains a signal that mixes the information that the forecasters have:

$$s_{jt} = \lambda\varepsilon_{jt} + (1-\lambda)\mathbb{E}_t^f y_t, \quad \lambda \in (0, 1/2), \tag{4}$$

where $\mathbb{E}_t^f y_t$ is the forecasters' estimate average for y_t. The value of z_t is revealed at the end of each period after all decisions have been made.

Benhabib et al. (2015) show that it is always an equilibrium for the agents to ignore z_t and believe that

$$y_t = \theta\sigma_\varepsilon^2/(2(1-\theta)) \equiv \phi^C. \tag{5}$$

In this case s_{jt} reveals ε_{jt} to the firms, so the firms each produce the efficient amount and overall output is constant. However, when $\lambda < 1/2$, they also find a sunspot equilibrium. In our notation, the sunspot equilibrium is obtained when all agents assume that output follows $y_t = \phi^S + \xi^S \cdot z_t$ with

$$\|\xi^S\|^2 = \frac{\theta\lambda(1-2\lambda)}{(1-\lambda)^2}\sigma_\varepsilon^2, \qquad \phi^S = \phi^C\left(1 - \frac{(1-\theta)(1-2\lambda)}{1-\lambda}\right). \tag{6}$$

These are sunspot equilibria: when the projection of z_t on a certain vector ξ^S is high the firms get a high signal, but since they do not know whether the signal is high due to ε_{jt} or y_t being high, they overproduce, and y_t ends up high as a result. Any vector ξ^S that satisfies the norm condition above can serve as an equilibrium.

3.3 Learning

The sunspot process ($z_t \in \mathbb{R}^k$) in our model represents nondescript amorphous variables that may affect agents expectations about economic outcomes. These are so called 'sentiments' variables. The stochastic equilibrium above only imposes a condition on the variance of output. Therefore, any stochastic process $\tilde{z}_t = M_t z_t$ can serve as the sunspot as long as $\{M_t\}_{t=0}^\infty$ are orthogonal matrices. Different agents having different understandings of what the sunspot is means that we should allow for different choices of $\{M_t\}_{t=0}^\infty$ for each agent.

For concreteness, it is assumed that each agent is considering a sunspot series of the form $\tilde{z}_t^j = (M^j)^t z_t$, that is, agent j's 'sentiment' at time $t+1$ is affected by a realization of z_{t+1} in the same way that it would be affected by the realization $z_t = M^j z_{t+1}$ at time t. Limiting our attention to time-independent M^j is assuming a constant drift that creates differences in how agents perceive the sunspot.[2] This assumption is somewhat arbitrary and, although it leads to tractable results, other options are certainly worth

[2]Notice that we are not discussing how the sentiment affects the agents' expectations, only how they perceive the latent sunspot process.

studying. Choosing M^j to be identical for all agents is simply the model of Benhabib et al. (2015).

At any point in time, all agents are assumed to have the point-belief that output is related to the sunspot via $y_t = \log Y_t = \phi^j + \tilde{\xi}^j \cdot \tilde{z}_t^j$, with $(\phi^j, \tilde{\xi}^j) \in \mathbb{R}^{k+1}$. In other words, we limit the belief space of each agent to points in $\mathbb{R}^{k+1}$. At the end of the period, the variable z_t and y_t are revealed, and firms update their beliefs. This non-Bayesian form of learning is sometimes called econometric learning, and has been used extensively in macroeconomics.[3]

The updating process can be written recursively:

$$\begin{aligned}
\begin{pmatrix} \phi_{t+1}^j \\ \tilde{\xi}_{t+1}^j \end{pmatrix} &= \begin{pmatrix} \phi_t^j \\ \tilde{\xi}_t^j \end{pmatrix} + g_t {\Upsilon_{t+1}^j}^{-1} \begin{pmatrix} 1 \\ \tilde{z}_t^j \end{pmatrix} (y_t - \phi_t^j - \tilde{\xi}_t^j \cdot \tilde{z}_t^j), \\
\Upsilon_{t+1}^j &= \Upsilon_t^j + g_t \left[\begin{pmatrix} 1 \\ \tilde{z}_t^j \end{pmatrix} \cdot \begin{pmatrix} 1 \\ \tilde{z}_t^j \end{pmatrix}' - \Upsilon_t^j \right] = \\
&= (1 - g_t)\Upsilon_t^j + g_t \begin{pmatrix} 1 & 0 \\ 0 & M^j \end{pmatrix} \begin{pmatrix} 1 & z_t' \\ z_t & z_t z_t' \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & M^{j\prime} \end{pmatrix}
\end{aligned}$$

where g_t is the gain sequence and Υ_t^j is the estimated variance-covariance matrix. The gain sequence $g_t = 1/t$ corresponds to least-square learning (RLS), and replicates the OLS estimator. This paper employs the RLS gain sequence as well as the sequence $g_t = (1-q)/(1-q^t)$, that corresponds to weighting past observations with a factor of $q \in (0,1)$ per-time-period. It is more reasonable to assume that agents who live in an environment that seems to keep changing would prefer to employ the latter gain sequence, in order to react faster to changes.

The estimator Υ_t^j depends on the initial prior Υ_0^j, on M^j, and on the realizations of z_t. Using the strong law of large numbers, it is straightforward to show that if $g_t = 1/t$ or $g_t \to 1 - q$, the estimators $\lim_{t\to\infty} \Upsilon_t^j = I_{k+1}$ uniformly over j. This is simply stating that all agents, regardless of their M^j, must come to agree on the variance-covariance matrix of the sunspot regardless of how they interpret it. Thus, for simplicity, it is assumed that $\Upsilon_t^j = I_{k+1}$ throughout. Furthermore, by redefining $\tilde{\xi}_t^j = (M^j)^t \xi_t^j$, the learning process simplifies to

$$\phi_{t+1}^j = \phi_t^j + g_t(y_t - \phi_t^j - \xi_t^j \cdot z_t), \tag{7a}$$

$$\xi_{t+1}^j = M^{j\prime} \cdot (\xi_t^j + g_t z_t (y_t - \phi_t^j - \xi_t^j \cdot z_t)). \tag{7b}$$

[3] For a comprehensive account of this approach, see Evans and Honkapohja (2012)

Finally, the beliefs of the forecasters at the beginning of every period are assumed to be identically distributed to those of the firms. This simplifying assumption is similar to assuming that firms do get to observe z_t but with a very large error, so that this information is not useful for making their own prediction about output, and that the surveys are conducted by polling representatives of other firms.

4 Analysis

4.1 Solving the Firm's Problem

In the appendix I show that a firm whose beliefs are given by (ϕ^j, ξ^j), and receives the signal s_{jt} will choose to produce:

$$y_{jt} = (1-\theta^{-1})\phi^j + \theta^{-1}\left[m(\|\xi^j\|^2)s_{jt} + \frac{1}{2}\hat{\Sigma}(\|\xi^j\|^2)\right],$$

where

$$m(\|\xi^j\|^2) = \frac{\theta\lambda\sigma_\varepsilon^2 - (1-\lambda)(1-\theta)\|\xi^j\|^2}{\lambda^2\sigma_\varepsilon^2 + (1-\lambda)^2\|\xi^j\|^2}, \tag{8a}$$

$$\hat{\Sigma}(\|\xi^j\|^2) = \frac{(\theta+\lambda-2\theta\lambda)^2\|\xi^j\|^2\sigma_\varepsilon^2}{\lambda^2\sigma_\varepsilon^2 + (1-\lambda)^2\|\xi^j\|^2}. \tag{8b}$$

By integrating over all firms we get (the full derivation is in appendix A.1)

$$y_t = \frac{1}{1-\theta}\log\int_0^1 e^{(1-\theta)(A^j + B^j\cdot z_t)}dj, \tag{9}$$

where

$$A^j = \frac{\sigma_\varepsilon^2}{2(1-\theta)}\left[\theta + \frac{1-\theta}{\theta}\lambda m(\|\xi^j\|^2)\right]^2 + \frac{1}{2\theta}\hat{\Sigma}(\|\xi^j\|^2) - \frac{(1-\theta)}{\theta}\phi^j,$$
$$B^j = \frac{(1-\lambda)}{\theta}m(\|\xi^j\|^2)\langle\xi^i\rangle.$$

4.2 Exact Stationary Solutions of the Associated ODE

Equation (9) describes the mapping from the full belief space to actual output. Equations (7) describe how beliefs are updated after each period. Together (7) and (9) form a stochastic difference equation that fully describes the dynamics of the model.

In standard learning models, one can obtain an associated ordinary differential equation by fixing the beliefs, and taking expectations of (7) (see Evans and Honkapohja, 2012, chap. 6). The intuitive explanation for this is that when the gain g_t is small, agents are observing many realizations of (y_t, z_t) while their beliefs are not changing significantly. In the limit of $g \to 0$, the stochastic difference equation can be approximated by taking a continuous-time limit and replacing the stochastic right-hand-side of (7) with its expectation with respect to z_t.

However, in our model the drift in the perception of beliefs also scales with time, so taking the continuous-time limit requires also replacing the matrix M^j with its continuous time counterpart. For concreteness, consider the three-dimensional model ($k = 2$), and let the M^j matrices be $M^j = M(\alpha^j)$, where

$$M(\alpha) = \begin{pmatrix} \cos\alpha & \sin\alpha \\ -\sin\alpha & \cos\alpha \end{pmatrix},$$

and α^j is distributed symmetrically around zero: that is, we include all the special-orthogonal matrices and choose some distribution around the unit matrix. In the continuous-time limit we take $g \to 0$ and $\alpha^j \to 0$, while g/α^j remains constant. In this limit, to first order

$$M^j = M(\alpha^j) = I_2 + \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \alpha + O(\alpha^2).$$

Substituting the above into (7), taking expectations with respect to z_t, and taking the continuous-time limit, we have

$$\dot{\phi}_t^j = g_t(\mathbb{E}_z y_t - \phi_t^j), \tag{10}$$

$$\dot{\xi}_t^j = \frac{\alpha^j}{g_t} \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \cdot \xi_t^j + g_t \mathbb{E}_z[z_t(y_t - \xi_t^j \cdot z_t)]. \tag{11}$$

The next natural step is to ask if the ODE system has stationary solutions. From (10), $\dot{\phi}_t^j = 0$ implies $\phi^j = \mathbb{E}_z y_t$, which means that all agents have a common belief $\phi^j = \phi$. Next, assume that $\langle \xi^j \rangle$ is in the 1-direction[4] Thus, in (9), the $B^j \cdot z_t$ term is independent of z^2, and therefore $\mathbb{E}_z[yz^2] = 0$. Equation (11) becomes

$$0 = \frac{\alpha^j}{g} \begin{pmatrix} -{\xi^j}^2 \\ {\xi^j}^1 \end{pmatrix} + \begin{pmatrix} \mathbb{E}_z[yz^1] - {\xi^j}^1 \\ -{\xi^j}^2 \end{pmatrix}.$$

[4]This is without loss of generality since we can always rotate the axes to achieve this.

The second component implies that $\xi^j = R^j(\cos\psi^j, \sin\psi^j)$, where $\tan\psi^j = \alpha^j/g$ (which is held constant at the continuous-time limit). Substituting this into the first component gives: $R^j = \mathbb{E}_z[yz^1]\cos\psi^j$. In particular, we find that R^j is proportional to $\cos\psi^j$, with the constant of proportionality common to all agents: thus

$$\xi^j = R\cos\psi^j \begin{pmatrix}\cos\psi^j \\ \sin\psi^j\end{pmatrix}, \quad \tan\psi^j = \frac{\alpha^j}{g}, \quad \psi \in \left(-\frac{\pi}{2}, \frac{\pi}{2}\right). \tag{12}$$

All that remains is to find R. To do this, we substitute the above into (9), and calculate $\mathbb{E}_z[yz^1]$. The resulting equation is

$$R = \frac{1}{1-\theta}\mathbb{E}\left[z^1 \log \int_0^1 e^{(1-\theta)(\tilde{A}^j + \tilde{B}^j z^1)} dj\right]. \tag{13}$$

where

$$\tilde{A}^j = \frac{\sigma_\varepsilon^2}{2(1-\theta)}\left[\theta + \frac{1-\theta}{\theta}\lambda m(R^2\cos^2\psi^j)\right]^2 + \frac{1}{2\theta}\hat{\Sigma}(R^2\cos^2\psi^j),$$
$$\tilde{B}^j = \frac{(1-\lambda)}{\theta} m(R^2\cos^2\psi^j) R\langle\cos^2\psi^i\rangle.$$

The derivation is in the appendix (A.2), which also contains a worked out example for a simple case where (13) can be evaluated analytically. Typically, one must evaluate the right-hand side of (13) numerically, but this is still straightforward since the distributions of ψ^j and z_1 are known, and therefore the right-hand-side of (13) is purely a function of R.

4.3 Analysis of the Stationary Solutions and IRFs

Equation (13) always admits the solution $R = 0$: in this case $\tilde{B}^j = 0$, implying that the right-hand side of (13) also vanishes. In this solution $\xi^j = 0$ and $\phi^j = \phi^C$ for all j, and output is constant. Economically, this solution reflects a situation in which agents do not learn to coordinate on the sunspot, and end up playing the non-stochastic equilibrium of equation (5). This equilibrium is not stable for $\lambda < 1/2$ in the original model of Benhabib et al. (2015), but can be stable in our model when the variance of α^j is large enough.

Equation (13) may have additional solutions with $R > 0$. In these solutions, agents are imperfectly coordinated on the use of the sunspots. The beliefs ξ^j, as described by (12) are all located on a circle of radius $R/2$, and centered a distance of $R/2$ from the origin. This is depicted in figure 1 for

three uniform distributions $\alpha^j/g \sim \mathrm{U}[-\eta, \eta]$. The solutions with $R > 0$ are analogous to the stochastic solutions of Benhabib et al. (2015), but unlike in their model, where all agents share the same ξ^j with $\|\xi^j\| = \xi^S$ (see equation (6)), we have heterogeneous beliefs.

As explained in the previous subsection, the stationary solutions exist in the limit where learning is 'slow', in the sense that agents get to observe many realizations of z_t with beliefs held close to constant. This leads us to suggest a natural extension to the concept of impulse-response-functions (IRF): we start at an exact stationary solution, fix $g > 0$ and study how beliefs change in response to a particular realization of z_t. In figure 2 we draw two IRFs describing the reaction to realizations of z_t that are parallel and perpendicular to $\langle \xi^j \rangle$.

The initial belief distribution is represented by the dotted black line. As in (12), we choose a solution in which $\langle \xi^j \rangle$ is in the 1-direction, thus, a parallel (perpendicular) sunspot shock is proportional to $z_t = (1, 0)$ ($z_t = (0, 1)$). The perpendicular shock has no effect on output, because the action of each agent with $\alpha^j > 0$ is cancelled by an agent with $-\alpha^j$. A parallel shock will increase or decrease output depending on the sign of z_t^1. Therefore, a perpendicular shock causes agents to decrease the perpendicular component of their ξ^j, which squishes the distribution towards the 1-axis (see the dashed red line). Conversely, a parallel shock leads agents to increase the parallel component of their belief and causes the distribution to expand outward (the solid blue line).

The impulse responses demonstrate an important phenomenon: there is a positive correlation between uncertainty and disagreement. In periods where agents have more similar beliefs – such as following a perpendicular shock – the uncertainty that the average agent has about output ($\langle |\xi^j| \rangle$) is smaller. Similarly, following parallel shocks both disagreement and uncertainty are higher. So far we have demonstrated this with the impulse-response functions, but in the following it is argued that it holds more generally in the model.

4.4 Numerical Analysis

The general model defined by (7) and (9) is a stochastic nonlinear model, and as such not many analytical results can be obtained. However, numerical simulations suggest that the solutions described in the previous subsection capture many of its important characteristics.

To get a general impression, first let us choose the particular distribution $\alpha^j \sim \mathrm{N}(0, \eta^2)$. Figure 3 displays the results of three typical trial runs

with low, mid, and high values of η.[5] In the left column of the figure, we see that when η is small, the solution is a noisy version of the stochastic equilibria of (6): the average belief about ϕ is close to ϕ^S and the beliefs about ξ remain close to some point on a circle of radius ξ^S. We call this case 'synchronization'.

Turning to the right-most column, we see that when η is high, coordination fails and players end up playing the non-stochastic equilibrium. This equilibrium would not be stable with rational agents (for $\lambda < 1/2$), but the disagreement over choice of a sunspot precludes coordination with our non-rational agents.

Finally, the central column illustrates the case for an intermediate value of η. Here we see an example of a learning process that does not converge, and yet does not diverge either. Over time beliefs about the use of the sunspot ξ_t^j converge and increase the average $\|\langle \xi_t^j \rangle\|$, but then rapidly diverge as coordination fails. Consequently, output becomes more or less volatile, and average output fluctuates between ϕ^S and ϕ^C.

Conjecture 1. *Consider the main model with $\alpha^j \sim N(0, \eta^2)$, and $g_t = (1-q)/(1-q^t)$. If equation (13) has no solution with $R > 0$, then for all sufficiently small $|1-q|$ as $t \to 0$ the beliefs weakly converge to the non-stochastic equilibria, i.e. $\phi_t^j \to \phi^C$ and $\xi^j \to 0$.*

We test this conjecture by fixing all the model parameters except η. We first find the bifurcation point of (13) and denote it by η^*. It is defined by the condition that there exist non-zero solutions if and only if $\eta < \eta^*$. Next, we run numerical simulations of the full model for $\eta > \eta^*$ and test for convergence to the non-stochastic equilibria. This was performed for many values of $\theta \in (0,1)$, $\lambda \in (0, 1/2)$ and $q \in (0, 0.1)$ and we have not found a counter-example.

We can now describe the full behaviour of the model. Starting at $\eta = 0$, the model converges to one of the stochastic equilibria where $\phi_t^j \to \phi^S$ and ξ_t^j all converge on a point on a circle of radius ξ^S. As η is increased from zero to η^*, two things happen: first, the radius of the associated stationary solution decreases from ξ^S to zero, and, second, the learning process becomes more 'noisy'. Thus, for small η the model can roughly be described as small

[5] The technical details of the simulations: we fix all parameter values, discretize to $J = 800$ agents, and set $\alpha^j = \eta \Phi^{-1}(j/J)$, $g_t = (1-q)/(1-q^t)$. Initial conditions are chosen at random. The simulation is run for $T = 10^5$ periods with different initial conditions. The results reported in the graph are typical for many values of the parameters that have been checked (the only requirements are $\lambda < 1/2$ and $1-q$ small enough to avoid immediate divergence).

fluctuations around the stationary solution with R slightly smaller than ξ^S, and for large η (but below η^*) the motion in the belief space appears to be very chaotic. Finally, once η is increased beyond η^*, we are in the regime of the above conjecture and the model converges to the non-stochastic equilibrium.[6] To better understand the above interpolation, figure 4 displays summary results of simulations for different values of η.[7] The four subplots display the values of: (a) the average ϕ_t^j across agents, (b) the dispersion of ϕ^j across agents: $\left\langle (\phi_t^j)^2 \right\rangle - \left\langle \phi_t^j \right\rangle^2$, (c) the norm of the average ξ_t^j, (d) the dispersion of the last quantity across agents. Note that these quantities change over time, so every point in the plot is an average both across time and across simulations. The dashed lines are standard deviations which give a sense of the ergodic distribution. These results suggest that the distributions are continuous in η. They also show the bifurcation point: above some η^* the system is exactly at the deterministic equilibrium, and the statistics (b)-(d) appear to have a discontinuous first derivative at the critical point.

5 Economic Implications and Empirical Evidence

The first conclusion of this paper is that synchronization on sunspot models does not require agents to be able to agree explicitly on what the sunspot is. As long as the different interpretations are not too far apart, agents can spontaneously learn to synchronize on the use of random noise, and while they will differ in their use of the noise, in the aggregate there will be a sunspot that is correlated with output.

The second conclusion is that, unless the differences in interpretation are minimal, the learning dynamics will not lead to a settled use of the sunspot. Rather, the process leads to a constant flow in the belief-space that generates periods of higher and lower levels of coordination. In periods of high coordination, the sunspot has a larger impact on agents' actions, which results in output being more volatile. We can, in fact, translate the results for low, high, and mid dispersions in α^j into three corresponding macroeconomic scenarios: (a) output is constantly volatile, (b) output is

[6] The motion of ξ_t^j through the belief space for different values of η is best illustrated in a video that we have made available at https://youtu.be/Xn2DR-CmWTg.

[7] Here we run the simulation a number of times with different initial conditions. We drop the first 1000 periods and calculate the statistics for the remaining periods. The averages are over time and runs, and the dashed lines represent a one-standard-error interval around the means, i.e. it is a statistic of the ergodic distribution of the variable.

constantly non-volatile ,(c) the volatility of output is itself volatile. This will also correspond to low (a,b) or high (c) levels of dispersion of beliefs of forecasters.

Thus, this model is a theory of how coordination emerges, and also of endogenously generated stochastic volatility. Furthermore, the model explains why we should expect agents' agreement on the use of any specific sunspot to be stable over a long period of time.

In the following subsection we describe some more implications of the model which are supported by empirical findings from surveys of professional forecasters.

5.1 Comparison to Surveys of Forecasters

Since this paper is all about beliefs, it makes sense to compare our findings to those from the literature analyzing surveys of professional forecasters. To do so, we first need to define how we relate the quantities in our model to empirical variables. The following table describes the most natural relationship:

Model Concept	Variable	Survey Concept
Individual point belief about future output (before observing the sunspot)	ϕ_t^j	Individual forecast for next period
Average belief about future output	$\langle \phi_t^j \rangle$	Average forecast for next period
Disagreement about future output	$V(\phi_t^j)$,[8]	Average forecast for next period
Individual uncertainty about future output	$\|\xi_t^j\|$	Individual forecast uncertainty for next period
Average uncertainty about future output	$\langle \|\xi_t^j\| \rangle$	Average uncertainty for next period

It is generally accepted that individual uncertainty about the future is counter-cyclical (Bloom, 2014; Bloom et al., 2018; Fajgelbaum et al., 2017; Ludvigson, 2016). This also holds in our model, and is a direct result of the fact that $\phi^C > \phi^S$. Generally speaking, agents are more uncertain about the future when they believe that the sunspot is more important ($\|\xi^j\|$ is larger), and $\mathbb{E}_z[y_t]$ can be shown to be lower when this happens. This perception was verified in simulations.

[8] The variance is over j at a given t.

While previous literature offers models that demonstrate how an exogenous increase in uncertainty can generate a recession, or, conversely, how an exogenous negative shock to output causes an increase to uncertainty, in this paper both are a result of a third mechanism: the build-up of coordination.

Belief dispersion is also counter-cyclical and positively correlated with uncertainty, but the relationship appears to be weaker (Jurado et al., 2015; Rich and Tracy, 2018, 2010; Mankiw et al., 2004). There are two ways to measure disagreement in our model, depending on whether we consider the forecast that agents make before or after observing the sunspot. In either case, our model predicts the empirical finding, but the reasoning is slightly different: if we consider the forecast after observing the sunspot, then it is enough to note that since the ξ^j fluctuate around the stationary solution, which is a circle of radius R, they are larger when they are also further away from each other. Therefore, larger average uncertainty ($\langle\|\xi^j\|\rangle$)) is associated with more dispersion in the forecast (dispersion of $\phi^j + \xi^j z_t$). If we consider the forecast before observing the sunspot, then it is an indirect result of the fact that more dispersion in ξ^j is associated with more dispersion in ϕ^j as a result of the learning process. Since agents are observing the same (z_t, y_t), estimating different ξ^j can only be consistent with also estimating different ϕ^j.

In numerical simulations we find the correlation between belief dispersion and output to be negative, but smaller than the correlation between uncertainty and output. We also find the correlation between belief dispersion and uncertainty to be positive.[9]

Another important finding of Rich and Tracy (2018) is that much of the variation in forecasters' uncertainty is explained by a fixed effect: that is, some forecasters just tend to be more uncertain about the economy. This also holds in our model: those agents with larger α^j will consistently have lower $\|\xi^j\|$. Intuitively, these agents' understanding of what the sunspot is is persistently farther away from the average understanding, and therefore, the learning process leads them to downplay the importance of the sunspot relative to agents with lower α^j.

[9]To the best of our knowledge the only paper other than this that endogenously generates a positive correlation between belief dispersion and uncertainty is Zohar (2020).

6 Additional Comments

6.1 Will they ever learn?

The agents in our model are using a misspecified learning model: they are not considering a law of motion with time-independent parameters, which is not the actual law of motion when other agents are also learning. Clearly, an agent who understands this could profit by making superior forecasts. However, this would require a very clever agent, who is also able to understand how other agents perceive the sunspot. Since the premise of the model was that the sunspot is not easily definable, that seems unlikely.

Still, one can try to address this concern by considering a modification of the model in which the forecasters who consistently make bad predictions fall out of the profession over time. Recall that the deterministic solution (12) describes a circle, and note that the forecast error is on average simply the difference between the belief of the agent and the average belief. If the average belief was the center of the circle, then there would be no difference in the average forecast error. Thus, in the deterministic model one can figure out which agents will get eliminated by considering the difference between the average belief and the center of the circle. One can construct examples where the worst forecasters are the ones with extreme values of α^j, and examples where they are the ones close to the mean. Therefore, dropping the worst forecasters might increase or decrease dispersion.

More specifically, when dispersion is low, the agents with the α^j that is farthest from the mean are the ones making the worst predictions. If they fall out, then dispersion becomes even lower until eventually we reach the S equilibrium of the rational expectations model. When dispersion is high, all the agents are virtually ignoring the sunspot, so no elimination would happen at all. For medium levels of dispersion, the relationship is not so clear: even if it is true that on average and over long periods of time the agents with extreme α^j are the worst forecasters, over short periods of time the differences are small compared to the variance, and the stochastic nature of the sunspot combined with the seemingly chaotic flows in the belief space introduce uncertainty into who will drop out. In numerical simulations we find that even over samples of 30-40 periods, sometimes it is actually the agents close to the mean who are performing worst.

To conclude, close to the extreme levels of dispersion, termination of 'bad' forecasters will reinforce the results we already have. For medium levels, the details of the model in combination with the rate of attrition will determine in which direction things go.

6.2 The Kuramoto Connection and the Matthews-Strogatz Model

Finally, we would like to add a comment about a similarity between our model and a version of the Kuramoto model (Kuramoto, 1975) that is due to Matthews and Strogatz (1990) (hence, MS).

In section 4.2 I derived equation (10). Using equation (9) to substitute out y_t in the right-hand side of (10) and expanding to second order in ξ_t^j, one can derive:

$$\dot{\xi}_t^j = \frac{\alpha^j}{g_t}\begin{pmatrix}0 & -1\\ 1 & 0\end{pmatrix}\cdot\xi_t^j + g\left(\left\langle\xi_t^i\right\rangle - \xi_t^j\right) - \frac{g}{\theta}\left\langle 1 + \frac{\lambda(\theta(1-2\lambda)+\lambda)\sigma_\epsilon^2}{(1-\lambda)^2(\|\xi_t^i\|^2 - \xi^{S2})}\right\rangle^{-1}\left\langle\xi_t^i\right\rangle. \tag{14}$$

This equation bears similarity to an equation studied by MS, which can be written as:[10]

$$\dot{\xi}_t^j = \begin{pmatrix}0 & -\omega^j\\ \omega^j & 0\end{pmatrix}\cdot\xi_t^j + g\left(\left\langle\xi_t^i\right\rangle - \xi_t^j\right) - (\|\xi_t^j\|^2 - \xi^{S2})\xi_t^j. \tag{15}$$

The first two terms are identical. To first order, the last term of (14) is proportional to $(\langle\|\xi_t^i\|^2\rangle - \xi^{S2})\langle\xi^i\rangle$, which is the same as (15) except that the latter has ξ_t^j instead of $\langle\xi_t^i\rangle$.

Both equations (14) and (15) lead to steady-state solutions where agents are synchronized and a trivial $\xi_t^j = 0$ solution. In addition, both systems have regions of the parameter space where each of the steady-state solutions is a global attractor. MS find additional regions where neither of the above are attractors, which is not the case in our model. The reason for the difference is due to the last term of the equation. In equation (15) each ξ^j has its own natural frequency ω^j, an interaction with the average $\langle\xi^i\rangle$, and a nonlinear interaction with itself. Without the interaction between ξ^j and the average, all agents would end up moving in a circle of radius ξ^S, with constant angular velocity ω^j. In (14) there is no limit cycle. Lacking the interaction with the average, each agent would spiral down to $\xi^j = 0$.

The phenomena that MS find (e.g. periodic fluctuations, non-periodic fluctuations, chaos) are fascinating and should be of interest to economists. Models that combine the methods of this paper with limit-cycle models would naturally lead to these phenomena.

[10] The notation was changed to conform to the notation of this paper

Acknowledgements

For their comments and suggestions, I would like to thank: Gadi Barlevi, Jess Benhabib, Victoria Consolvo, Alex Cukierman, George Evans, Roger Farmer, Jesus Fernandez-Villaverde, Oded Galor, Luca Guerrieri, Joao Guerreiro, Zhen Huo, Mohammad Jahan-Parvar, Guido Lorenzoni, Bruce McGough, Guido Menzio, Karel Mertens, Kristoffer Nimark, Robert Rich, Karl Shell, and Mirko Wiederholt. I am grateful to Doron Zamir for his assistance in research.

This research was supported in part by the Israel Science Foundation (grant 2048/16).

References

J. Abel, R. Rich, J. Song, and J. Tracy. The Measurement and Behavior of Uncertainty: Evidence from the ECB Survey of Professional Forecasters. *Journal of Applied Econometrics*, 31(3):533–550, 2016.

J. A. Acebrón, L. L. Bonilla, C. J. Pérez Vicente, F. Ritort, and R. Spigler. The Kuramoto model: A simple paradigm for synchronization phenomena. *Reviews of Modern Physics*, 77(1):137–185, Apr. 2005.

G. Angeletos and J. La'O. Sentiments. *Econometrica*, 81(2):739–779, 2013.

R. Aumann, J. Peck, and K. Shell. Asymmetric information and sunspot equilibria: A family of simple examples. Working Paper 88-34, Center for Analytic Economics, Cornell University, Ithaca, Oct 1988.

J. Benhabib and R. E. A. Farmer. Indeterminacy and Increasing Returns. *Journal of Economic Theory*, 63(1):19–41, June 1994.

J. Benhabib and M. M. Spiegel. Sentiments and Economic Activity: Evidence from US States. *Economic Journal*, 129(618):715–733, 2019.

J. Benhabib, P. Wang, and Y. Wen. Sentiments and Aggregate Demand Fluctuations. *Econometrica*, 83:549–585, March 2015.

N. Bloom. Fluctuations in Uncertainty. *Journal of Economic Perspectives*, 28(2):153–176, 2014.

N. Bloom, P. Bunn, S. Chen, P. Mizen, P. Smietanka, G. Thwaites, and G. Young. Brexit and Uncertainty: Insights from the Decision Maker Panel. *Fiscal Studies*, 39(4):555–580, 2018.

D. Cass and K. Shell. Do Sunspots Matter? *Journal of Political Economy*, 91(2):193–227, 1983.

L. J. Christiano and S. G. Harrison. Chaos, sunspots, and automatic stabilizers. Staff Report 214, Federal Reserve Bank of Minneapolis, 1996.

H. Daido. A solvable model of coupled limit-cycle oscillators exhibiting partial perfect synchrony and novel frequency spectra. *Physica D: Nonlinear Phenomena*, 69(3):394–403, Dec. 1993.

H. Daido. Generic scaling at the onset of macroscopic mutual entrainment in limit-cycle oscillators with uniform all-to-all coupling. *Phys. Rev. Lett.*, 73(5):760–763, Aug. 1994.

H. Daido. Multibranch Entrainment and Scaling in Large Populations of Coupled Oscillators. *Phys. Rev. Lett.*, 77(7):1406–1409, Aug. 1996a.

H. Daido. Onset of cooperative entrainment in limit-cycle oscillators with uniform all-to-all interactions. *Physica D: Nonlinear Phenomena*, 91(1): 24–66, Mar. 1996b.

J. Dovern, U. Fritsche, and J. Slacalek. Disagreement Among Forecasters in G7 Countries. *Review of Economics and Statistics*, 94(4):1081–1096, May 2011.

G. B. Ermentrout. Oscillator death in populations of "all to al" coupled nonlinear oscillators. *Physica D: Nonlinear Phenomena*, 41(2):219–231, Mar. 1990.

G. W. Evans and S. Honkapohja. Existence of adaptively stable sunspot equilibria near an indeterminate steady state. *Journal of Economic Theory*, 111(1):125–134, 2003a.

G. W. Evans and S. Honkapohja. Expectational stability of stationary sunspot equilibria in a forward-looking linear model. *Journal of Economic Dynamics and Control*, 28(1):171–181, 2003b.

G. W. Evans and S. Honkapohja. *Learning and Expectations in Macroeconomics*. Princeton University Press, 2012.

G. W. Evans, S. Honkapohja, and S. Honkapohja. Learning, convergence, and stability with multiple rational expectations equilibria. *European Economic Review*, 38(5):1071–1098, 1994.

P. D. Fajgelbaum, E. Schaal, and M. Taschereau-Dumouchel. Uncertainty Traps. *The Quarterly Journal of Economics*, 132(4):1641–1692, 2017.

D. Fehr, F. Heinemann, and A. Llorente-Saguer. The power of sunspots: An experimental analysis. *Journal of Monetary Economics*, 103(C):123–136, 2019.

J. Fernández-Villaverde and P. A. Guerrón-Quintana. Uncertainty Shocks and Business Cycle Research. Technical Report 26768, National Bureau of Economic Research, Inc, Feb. 2020. Publication Title: NBER Working Papers.

R. Guesnerie and M. Woodford. Stability of Cycles with Adaptive Learning Rules. DELTA Working Paper 90-25, DELTA (Ecole normale supérieure), 1990.

S. Honkapohja and K. Mitra. Are non-fundamental equilibria learnable in models of monetary policy? *Journal of Monetary Economics*, 51(8):1743–1770, 2004.

K. Jurado, S. C. Ludvigson, and S. Ng. Measuring Uncertainty. *American Economic Review*, 105(3):1177–1216, 2015.

Y. Kuramoto. Self-entrainment of a population of coupled non-linear oscillators. In P. H. Araki, editor, *International Symposium on Mathematical Problems in Theoretical Physics*, number 39 in Lecture Notes in Physics, pages 420–422. Springer Berlin Heidelberg, 1975.

S. Ludvigson. Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response. Technical Report 183, Society for Economic Dynamics, 2016.

N. G. Mankiw, R. Reis, and J. Wolfers. Disagreement about Inflation Expectations. NBER Chapters, National Bureau of Economic Research, Inc, 2004.

A. Marcet and T. J. Sargent. Convergence of least squares learning mechanisms in self-referential linear stochastic models. *Journal of Economic Theory*, 48(2):337–368, Aug. 1989.

P. C. Matthews and S. H. Strogatz. Phase diagram for the collective behavior of limit-cycle oscillators. *Phys. Rev. Lett.*, 65(14):1701–1704, Oct. 1990.

P. C. Matthews, R. E. Mirollo, and S. H. Strogatz. Dynamics of a large system of coupled nonlinear oscillators. *Physica D: Nonlinear Phenomena*, 52(2):293–331, Sept. 1991.

N. P. Nayar and A. Levchenko. TFP, News and "Sentiments:" The International Transmission of Business Cycles. Technical Report 1076, Society for Economic Dynamics, 2017.

S. H. Park and S. Kim. Noise-induced phase transitions in globally coupled active rotators. *Phys. Rev. E*, 53(4):3425–3430, Apr. 1996.

A. J. Patton and A. Timmermann. Why do forecasters disagree? Lessons from the term structure of cross-sectional dispersion. *Journal of Monetary Economics*, 57(7):803–820, Oct. 2010.

R. Rich and J. Tracy. The Relationships Among Expected Inflation, Disagreement, and Uncertainty: Evidence from Matched Point and Density Forecasts. *The Review of Economics and Statistics*, 92(1):200–207, 2010.

R. W. Rich and J. Tracy. A Closer Look at the Behavior of Uncertainty and Disagreement: Micro Evidence from the Euro Area. Working paper 1811, FRB of Dallas, July 2018.

S. H. Strogatz. From Kuramoto to Crawford: exploring the onset of synchronization in populations of coupled oscillators. *Physica D: Nonlinear Phenomena*, 143(1–4):1–20, Sept. 2000.

M. Woodford. Learning to Believe in Sunspots. *Econometrica*, 58(2):277–307, Mar. 1990.

O. Zohar. Cyclicality of Uncertainty and Disagreement. Feb. 2020.

A Mathematical Derivations

A.1 The Firm's Problem

First, consider a firm whose beliefs are given by (ϕ^j, ξ^j). Defining $x_{jt} = (\theta-1)(y_t - \phi^j) + \theta\varepsilon_{jt}$, we have from (3)

$$y_{jt} = \theta^{-1} \log \mathbb{E}_{jt}\left[e^{(\theta-1)y_t + \theta\epsilon_{jt}} | s_{jt}\right] = (1-\theta^{-1})\phi^j + \theta^{-1} \log \mathbb{E}_{jt}\left[e^{x_{jt}} | s_{jt}\right]. \quad (16)$$

Since $\log \epsilon_{j,t} \sim \mathrm{N}(0, \sigma_\varepsilon^2)$, we have $\mathbb{E}_{jt}[x_{jt}] = \mathbb{E}_{jt}[s_{jt}] = 0$, and the subjective variance-covariance matrix of (x_{jt}, s_{jt}) is

$$\Sigma = \begin{pmatrix} \theta^2\sigma_\varepsilon^2 + (1-\theta)^2\|\xi^j\|^2 & \theta\lambda\sigma_\varepsilon^2 - (1-\lambda)(1-\theta)\|\xi^j\|^2 \\ \text{sym.} & \lambda^2\sigma_\varepsilon^2 + (1-\lambda)^2\|\xi^j\|^2 \end{pmatrix}.$$

Therefore, $x_{jt}|s_{jt} \sim \mathrm{N}(m(\|\xi^j\|^2)s_{jt}, \hat{\Sigma}(\|\xi^j\|^2))$, where

$$m(\|\xi^j\|^2) = \frac{\theta\lambda\sigma_\varepsilon^2 - (1-\lambda)(1-\theta)\|\xi^j\|^2}{\lambda^2\sigma_\varepsilon^2 + (1-\lambda)^2\|\xi^j\|^2},$$
$$\hat{\Sigma}(\|\xi^j\|^2) = \frac{(\theta+\lambda-2\theta\lambda)^2\|\xi^j\|^2\sigma_\varepsilon^2}{\lambda^2\sigma_\varepsilon^2 + (1-\lambda)^2\|\xi^j\|^2}.$$

Therefore, from (16)

$$y_{jt} = (1-\theta^{-1})\phi^j + \theta^{-1}\left[m(\|\xi^j\|^2)s_{jt} + \frac{1}{2}\hat{\Sigma}(\|\xi^j\|^2)\right].$$

Using this in (2),

$$(1-\theta)y_t = \log\int_0^1 e^{\theta\varepsilon_{jt}+(1-\theta)y_{jt}}dj =$$
$$= \log\int_0^1 e^{\theta\varepsilon_{jt}+(1-\theta)\{(1-\theta^{-1})\phi^j+\theta^{-1}[m(\|\xi^j\|^2)s_{jt}+\frac{1}{2}\hat{\Sigma}(\|\xi^j\|^2)]\}}dj =$$
$$= \log\int_0^1 e^{\theta\varepsilon_{jt}+(1-\theta)\{(1-\theta^{-1})\phi^j+\theta^{-1}[m(\|\xi^j\|^2)(\lambda\varepsilon_{jt}+(1-\lambda)\langle\xi^i\rangle\cdot z_t)+\frac{1}{2}\hat{\Sigma}(\|\xi^j\|^2)]\}}dj,$$

where (4) is used in the last step. Since ε_{jt} is independent of beliefs, we can integrate

$$y_t = \frac{1}{1-\theta}\log\int_0^1 e^{(1-\theta)(A^j+B^j\cdot z_t)}dj,$$

where

$$A^j = \frac{\sigma_\varepsilon^2}{2(1-\theta)}\left[\theta + \frac{1-\theta}{\theta}\lambda m(\|\xi^j\|^2)\right]^2 + \frac{1}{2\theta}\hat{\Sigma}(\|\xi^j\|^2) - \frac{(1-\theta)}{\theta}\phi^j,$$
$$B^j = \frac{(1-\lambda)}{\theta}m(\|\xi^j\|^2)\langle\xi^i\rangle.$$

A.2 Solving for R

Starting with (9), we substitute $\phi^j = \phi$ and

$$\xi^j = R\cos\psi^j \begin{pmatrix}\cos\psi^j \\ \sin\psi^j\end{pmatrix}, \quad \tan\psi^j = \frac{\alpha^j}{g},$$

to find

$$y = \frac{1}{1-\theta}\log\int_0^1 e^{(1-\theta)(\tilde{A}^j + \tilde{B}^j z^1)}dj - \frac{1-\theta}{\theta}\phi,$$

where

$$\tilde{A}^j = \frac{\sigma_\varepsilon^2}{2(1-\theta)}\left[\theta + \frac{1-\theta}{\theta}\lambda m(R^2\cos^2\psi^j)\right]^2 + \frac{1}{2\theta}\hat{\Sigma}(R^2\cos^2\psi^j),$$
$$\tilde{B}^j = \frac{(1-\lambda)}{\theta}m(R^2\cos^2\psi^j)R\langle\cos^2\psi^i\rangle.$$

Since $R = \mathbb{E}[yz^1]$, we find

$$R = \frac{1}{1-\theta}\mathbb{E}\left[z^1\log\int_0^1 e^{(1-\theta)(\tilde{A}^j + \tilde{B}^j z^1)}dj\right]. \tag{17}$$

Recall that the ψ^j's are all known, so while the right-hand side of (17) is complicated, it is always straightforward to evaluate numerically, and the result is only a function of R (and the model parameters θ, λ and σ_ϵ).

For example, consider the simple case where

$$\alpha^j = \begin{cases}+\bar{\alpha} & j \in [0, 1/2] \\ -\bar{\alpha} & j \in (1/2, 1]\end{cases}$$

Note that in this example, ψ^j is equal to $\bar{\psi}$ for one half of the population and $-\bar{\psi}$ for the other, therefore $\tilde{A}^j$ and $\tilde{B}^j$ are independent of j, therefore (17) simplifies to

$$R = \tilde{B}^j = \frac{(1-\lambda)}{\theta}m(R^2\cos^2\bar{\psi})R\cos^2\bar{\psi}.$$

This equation has the trivial solution $R = 0$ and another nontrivial solution $R > 0$ if

$$\left(\frac{\lambda}{1-\lambda} - \cos^2\bar{\psi}\right)\left(\frac{\theta}{1-\theta} - \cos^2\bar{\psi}\right) > 0.$$

Figures

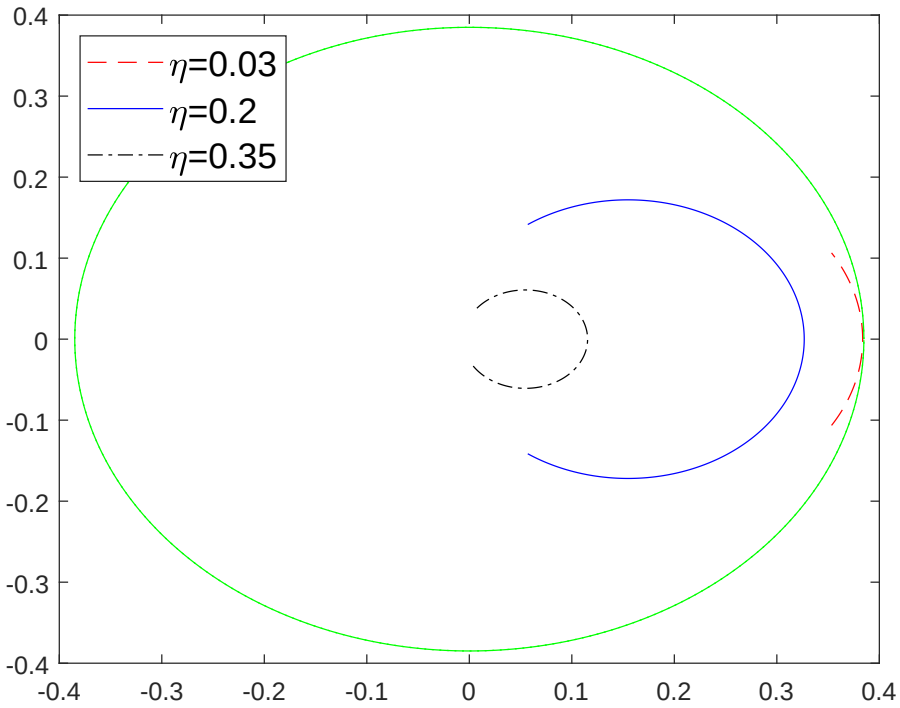

Figure 1: The steady state solution of the steady-state solution to the (fully nonlinear) non-stochastic model for different values of η. The outer circle is of radius ξ^S.

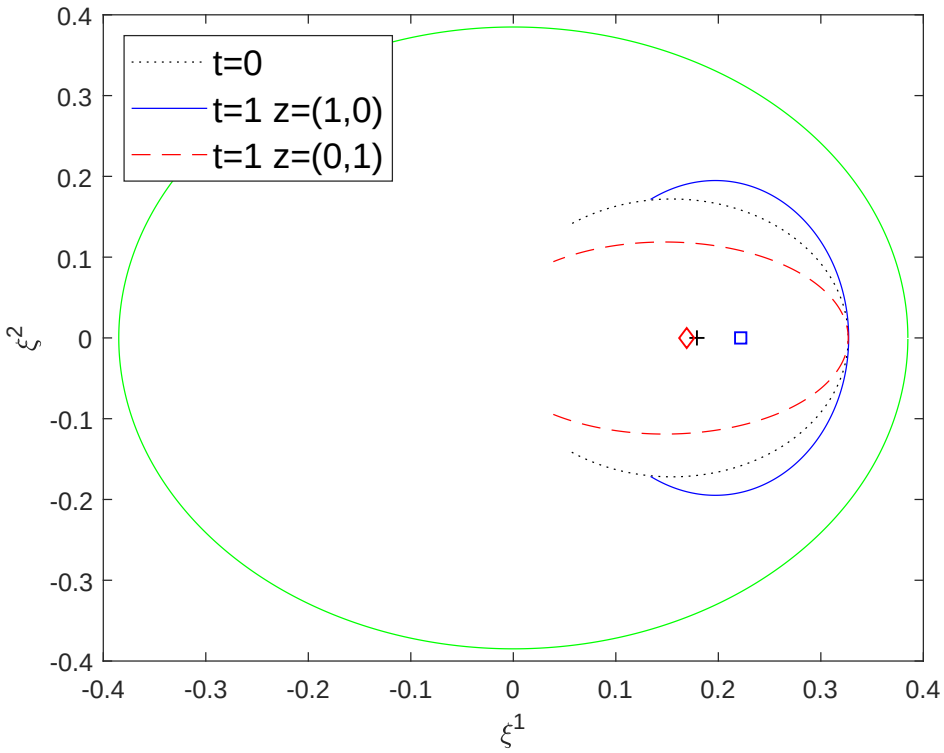

Figure 2: The 'Impulse Response Functions' described in subsection 4.3. The dotted black line is the steady-state solution to the (fully nonlinear) non-stochastic model. The solid and dashed lines show how the ξ^j distributions react to a shocks parallel and perpendicular to $\langle \xi^i \rangle$ respectively. The average belief at $t = 0$ and after the parallel and perpendicular shocks are denoted by a plus, square and diamond respectively.

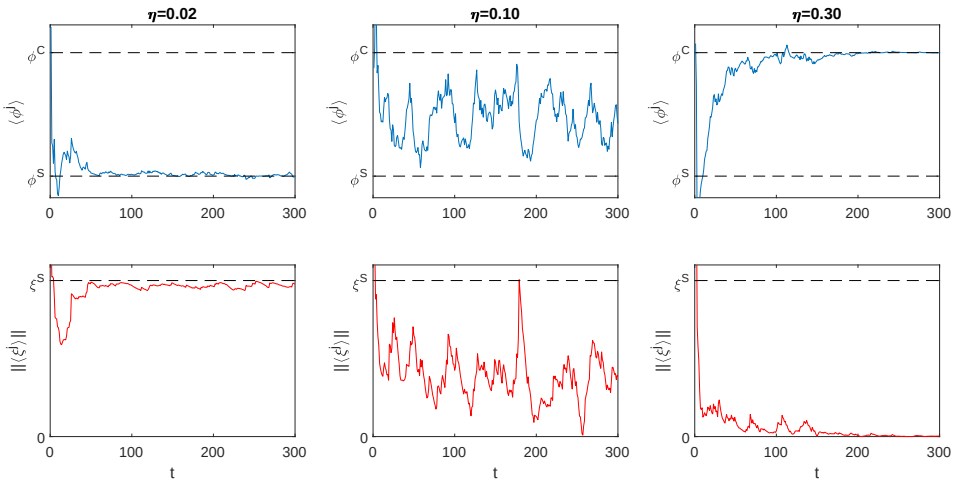

Figure 3: The left, center, and right columns, are each an example of a simulation of the full system for low, mid and high values of η respectively. In each column the top graph displays the evolution of the average belief $\left\langle \phi_t^j \right\rangle$. The middle graph shows the norm of the average belief on ξ, i.e. $\| \left\langle \xi_t^j \right\rangle \|$. The remaining parameters are: $\theta = 2/3, \lambda = 1/4, \sigma_\epsilon = 1, q = 0.9$. A video of these simulations is available at https://youtu.be/Xn2DR-CmWTg

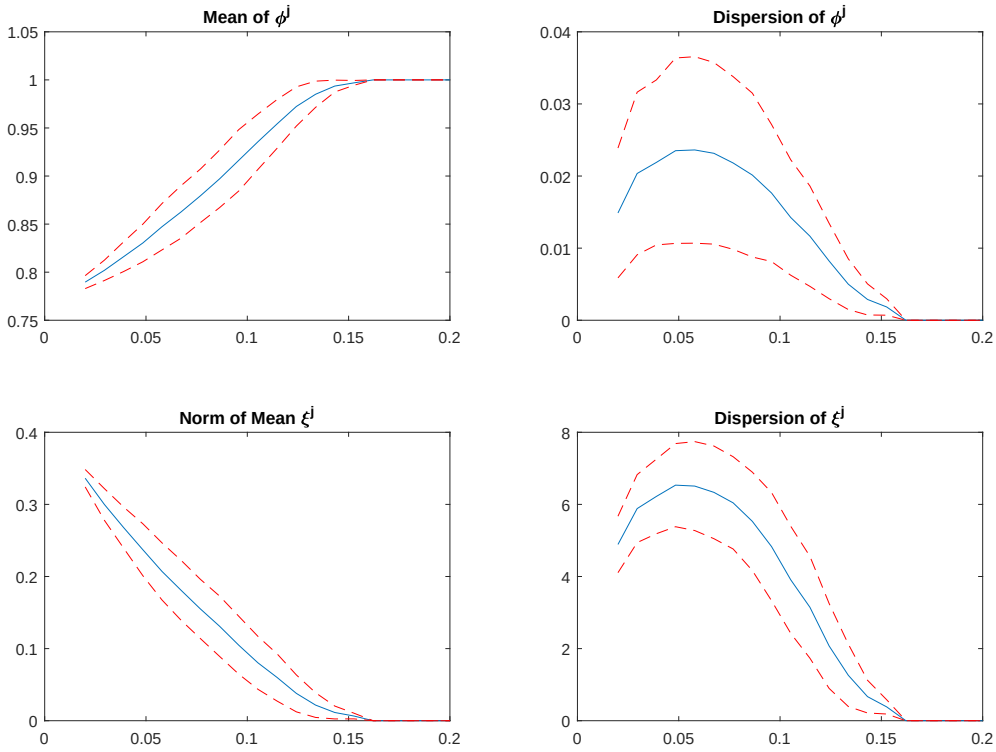

Figure 4: Summary statistics for simulations of the main model. The horizontal axes display values of η, and all other parameters are kept constant as in the previous figure.