Synchronization in Sunspot Models

Patir, Assaf

The Hebrew University of Jerusalem

25 August 2019

Synchronization in Sunspot Models

Assaf Patir*

August 25, 2019

Abstract

This note illustrates how agents' beliefs about economic outcomes can dynamically synchronize and de-synchronize to produce business-cycle-like fluctuations in a simple macroeconomic model. I consider a simple macroeconomic model with multiple equilibria, which are different ways that sunspots can forecast future output in a self-fulfilling manner. Agents are assumed to learn to use the sunspot variable through econometric learning. I show that if different agents have different interpretations of the sunspot, this leads to a complex nonlinear dynamic of synchronization of beliefs about the equilibrium being played. Depending on the extent of disagreement on the interpretation of the sunspot, the economy will be more or less volatile. The dispersion of the 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 the output is stable.

The equation describing the evolution of the economy can be interpreted as a nonlinear-stochastic version of the Kuramoto model, a prototypical model of synchronization phenomena, and simulations confirm that the qualitative features of the model are in agreement with results from the Kuramoto literature.

1 Introduction

This paper has two goals: first, to analyze, in a model with strategic uncertainty, how using a sunspot can emerge as a coordination mechanism even when agents don't agree on what the sunspot is; and, second, to show that in such a situation the perpetual learning about the equilibrium leads to fluctuations in output.

Macroeconomic models with strategic uncertainty frequently use sunspots to facilitate coordination between agents (Cass and Shell, 1983). The literal interpretation is that the agents incorporate some extrinsic source of randomness into their decision making process in a manner that generates an actual law of motion that is 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

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

that a certain real-life system is well described by a sunspot model, how come they can't tell us what the sunspot is? How could it be that the lay agents are all coordinating on a sunspot, yet the economists are not clued in on what the sunspot it?

Of course, one shouldn't be so literal when interpreting a model: a more plausible interpretation is that the agents are observing various processes that are influenced by a random process that is not related to the fundamentals, and 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 on using them, therefore, unlike signals about fundamentals, it is far less clear if agents can learn its relationship to economic outcomes unless the other agents are already coordinated. Previous literature has assumed that agents know that some process, $\{z_t\}$, is potentially relevant to outcomes, and analyzed the circumstances under which they can learn to use it,[1] but this might be considerably more difficult when they don't 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 preformed a fixed amount of time after sunrise. Both stochastic series would be iid with the same distribution as would infinitely 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 won't have such different interpretations. Thus, the problem of learning to coordinate on sunspots is before anything else, the problem of learning to construct a stochastic process out of many sources of randomness.

In this paper I demonstrate 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, I consider a system with a continuum of sunspot equilibria that are distinguished by different choices of the sunspot, and show that agents that are using a simple learning rule can converge on playing an outcome similar to an equilibrium, even while not agreeing about the identity of the sunspot. As the level of disagreement increases up to some critical level, the outcome is that 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.

Learning models have typically been used in economics to study stability properties of equilibria. If the learning process leads to convergence to the equilibrium, then it is more plausible that it will be observed in reality, and the converse, if the process diverges, then the equilibrium would not be played by agents whom are not perfectly rational and knowledgable. In contrast, the analysis here shows that the learning process can lead to complicated non-linear dynamics in the agents' belief-space, in which their beliefs

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

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 the 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 the agents' actions canceling out and output being roughly constant).

As we shall see, the learning process does not converge to a point in the belief space. Instead, the agents' beliefs about the relationship between the sunspot and output fluctuate around a "steady-state". This volatility in the belief space is highest for some medial level of disagreement, and drops toward higher and lower levels.

While the results of the paper can easily be adapted to any sunspot model, I focus for concreteness 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 two publicly observed stochastic processes ($z_t^i, i = 1, 2$) 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 note is as follows: in the next section I review some of the relevant economic and mathematical literature. Section 3 describes the model, which is a modified version of the model of Benhabib, Wang, and Wen (2015). Section 4 includes the main analysis: describing the rational-expectations equilibria of the model and their stability properties, the results of numerical simulations for the full model, as well as some analytical results based on simplifications of the model. Finally, some concluding remarks are left to section 5.

2 Related Literature

Learning has a long history in macroeconomics, but the stochastic recursive description in this paper originates with Marcet and Sargent (1989). For a comprehensive account of the state of this field see Evans and Honkapohja (2012). Some papers that study learning in 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 for the agents to coordinate their actions on (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 the *sentiment* to choose their actions. While similar in spirit, these are formally different situations. This note 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.

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 my model links the agents' beliefs about the equilibrium being played. This is, to my 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 amplitudes 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

$$C_t = \frac{1}{\psi} \cdot \frac{W_t}{P_t}, \tag{1}$$

$$N_t = \frac{1}{\psi} - \frac{\Pi_t}{W_t}. \tag{2}$$

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

$$Y_t = \left[\int_0^1 \epsilon_{jt}^{\theta} Y_{jt}^{1-\theta} dj \right]^{\frac{1}{1-\theta}}, \quad \theta > 0 \tag{3}$$

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 which 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 firms 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, I choose units of output 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

$$y_{jt} = \frac{1}{\theta} \log \mathbb{E}_{jt} \left[e^{\theta \epsilon_{jt} - (1-\theta) y_t} \right]. \tag{4}$$

3.2 Information

There is a large number of 'forecasters' who observe both the 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, i.e. 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 the forecasters to estimate their demand curves. However, the firm is limited in its ability to conduct market research, so that it eventually obtains a signal that mixes the information that the forecasters have:

$$s_{jt} = \lambda\varepsilon_{jt} + (1-\lambda)\mathbb{E}_t^f y_t, \quad \lambda \in (0,1), \tag{5}$$

where $\mathbb{E}_t^f y_t$ is the forecasters average estimate 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

$$y_t = \theta\sigma_\varepsilon^2/(2(1-\theta)) \equiv \phi^C. \tag{6}$$

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 my notation, the sunspot equilibrium is obtained when all agents assume that output follows $y_t = \phi^S + \xi^S \cdot z_t$ with

$$\|\xi^S\|^2 = \frac{\theta\lambda(1-2\lambda)}{(1-\lambda)^2}\sigma_\varepsilon^2, \qquad \phi^S = \phi^C\left(1 - \frac{(1-\theta)(1-2\lambda)}{1-\lambda}\right). \tag{7}$$

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 if 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.

In this paper I focus on the case of $\lambda > 1/2$, when both types of equilibria exist, but, as shown in the appendix, only the sunspot equilibria are stable under the learning scheme. Therefore, the agents are trying to learn about the use of z_t.

3.3 Learning

The sunspot-process takes values in $\otimes_{t=0}^{\infty}(\mathbb{R}^k)$, and I will use a particular basis in this space to describe z_t. However, as explained in the introduction, I do not want to assume that this basis is special in some way, or that agents should give that choice any precedence. Indeed, for any arbitrary sequence of orthogonal matrices $\{M_t\}_{t=0}^{\infty}$, with $M_t' = M_t^{-1}$, the series $\tilde{z}_t = M_t z_t$ has the same statistical properties as z_t. an arbitrary sequence of orthogonal matrices $(M_t' = M_t^{-1})$. The choice of a series is a matter of how agents understand what the sunspot is, and since I assume that the sunspot is an amorphous source of randomness, I allow different agents to have different understandings of what the sunspot process is.

Specifically, I assume that each agents is considering a sunspot series of the form $\tilde{z}_t^j = (M^j)^t z_t$, i.e. agents j believes that the sunspot $\tilde{z}_t^j$ is conceptually related to $M^j \tilde{z}_{t+1}^j$. The model of Benhabib et al. (2015) is simply the case when all agents use the same M^j.

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.[2]

The updating process can be written recursively:

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

where g_t is the gain sequence and Υ_t^j is the estimated variance-covariance martix. 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, I assume 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

$$\phi_{t+1}^j = \phi_t^j + g_t(y_t - \phi_t^j - \xi_t^j \cdot z_t^j), \tag{8a}$$

$$\xi_{t+1}^j = {M^j}' \cdot (\xi_t^j + g_t z_t (y_t - \phi_t^j - \xi_t^j \cdot z_t)). \tag{8b}$$

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.

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

4 Analysis

4.1 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 (4)

$$y_{jt} = \theta^{-1} \log \mathbb{E}_{jt}\left[e^{(\theta-1)y_t+\theta\epsilon_{jt}}|s_{jt}\right] = (1-\theta^{-1})\phi^j + \theta^{-1}\log \mathbb{E}_{jt}\left[e^{x_{jt}}|s_{jt}\right]. \tag{9}$$

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

$$m(\|\xi^j\|^2) = \frac{\theta\lambda\sigma_\varepsilon^2 - (1-\lambda)(1-\theta)\|\xi^j\|^2}{\lambda^2\sigma_\varepsilon^2 + (1-\lambda)^2\|\xi^j\|^2}, \tag{10a}$$

$$\hat{\Sigma}(\|\xi^j\|^2) = \frac{(\theta+\lambda-2\theta\lambda)^2\|\xi^j\|^2\sigma_\varepsilon^2}{\lambda^2\sigma_\varepsilon^2 + (1-\lambda)^2\|\xi^j\|^2}. \tag{10b}$$

Therefore, from (9)

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

Using this in (3),

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

where (5) is used in the last step. Since ε_{jt} is independent of beliefs, we can integrate

$$\begin{aligned}(1-\theta)y_t = \log\int_0^1 & e^{\frac{\sigma_\varepsilon^2}{2}[\theta+(\theta^{-1}-1)\lambda m(\|\xi^j\|^2)]^2}\times \\ & \times e^{(1-\theta)\{(1-\theta^{-1})\phi^j+\theta^{-1}[(1-\lambda)m(\|\xi^j\|^2)\langle\xi^i\rangle\cdot z_t+\frac{1}{2}\hat{\Sigma}(\|\xi^j\|^2)]\}}dj.\end{aligned} \tag{11}$$

Equation (11) describes the mapping from the full belief space to actual output.

4.2 Rational Expectations Equilibria

We can recover the equilibria found in Benhabib et al. (2015) by considering equation (11) in the case that all agents have common beliefs $\phi^j = \phi$, $\xi^j = \xi$. We have

$$(1-\theta)y_t = -\frac{(1-\theta)^2}{\theta}\phi + \frac{1}{2}[\theta + (\theta^{-1}-1)m(\|\xi\|^2)\lambda]^2\sigma_\varepsilon^2 \\ + (\theta^{-1}-1)\left[m(\|\xi\|^2)(1-\lambda)\xi \cdot z_t + \frac{1}{2}\hat{\Sigma}(\|\xi\|^2)\right].$$

Since y_t is a linear function of z_t, the mapping from the commonly perceived law of motion $y_t = \phi + \xi \cdot z_t$ to the actual law of motion is:

$$\phi \to -\frac{(1-\theta)}{\theta}\phi + \frac{1}{2\theta}\hat{\Sigma}(\|\xi\|^2) + \frac{[\theta + (\theta^{-1}-1)m(\|\xi\|^2)\lambda]^2\sigma_\varepsilon^2}{2(1-\theta)}, \tag{12a}$$

$$\xi \to \frac{1}{\theta}m(\|\xi\|^2)(1-\lambda)\xi. \tag{12b}$$

Rational expectations equilibria (REE) are fixed points of the above mapping. One such fixed point is:

$$\xi^C = 0, \qquad \phi^C = \frac{\theta\sigma_\varepsilon^2}{2(1-\theta)},$$

which is the equilibrium described above in (6). The superscript C stands for 'certainty', since output is constant in this fixed point. The S equilibria (for 'stochastic') of (7) are the fixed points that exist when $0 < \lambda < 1/2$:

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

As noted in Benhabib et al. (2015), average output is lower in the stochastic equilibrium ($\phi^C > \phi^S$), and it is also straightforward to show that the welfare of the representative consumer is lower.

It is worth noting that the C equilibrium is only stable-under-learning for $\lambda \in (1/2, 1)$. For $\lambda \in (0, 1/2)$, when both equilibria exist, only the S equilibria are stable. The remainder of the paper focuses on the case $\lambda \in (0, 1/2)$.

4.3 Numerical Analysis

Before going into further analysis, it is useful to get a general idea of the dynamics of this system through numerical simulations. Consider first 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 $\alpha \sim \mathrm{N}(0, \eta^2)$, i.e. I include all the special-orthogonal matrices and choose some distribution around the unit matrix.

To get a general impression, figure 1 displays the results of three typical trial runs with low, mid, and high values of η^2.[3] Recall that for the parameter values chosen, the stochastic equilibria defined by $\phi = \phi^S$, $|\xi| = \xi^S$ are stable, i.e. for $\eta = 0$ the learning process will converge to one of these equilibria. The specific equilibrium chosen in this case depends on the initial beliefs and the realization of the stochastic process. The left-most column in figure 1 demonstrates that the dynamics seem to be continuous: for low levels of disagreement (η^2) the economy settles into a noisy version of the stochastic equilibria. We see that the average belief about ϕ is close to ϕ^S and the belief about ξ remains on a circle of radius ξ^S. Further investigation shows that dispersion is also low, i.e. each individual belief remains close to the average belief, and the errors that agents are making are roughly averaged out of the aggregate outcome.

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

Finally, the central column illustrates the case for an intermediate value of η^2. 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.

To better understand the above interpolation, figure 2 displays summary results of simulations for different values of η^2.[4] The four subplots display the values of: (a) the average ϕ_t^j across agents, (b) the dispersion of ϕ^j across agents: $\left\langle \phi_t^j \right\rangle^2 - \left\langle \phi_t^j \right\rangle^2$, (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 suggests that the distributions are continuous in η^2. It also shows that the system has a bifurcation point: above some critical level of η^2 the system is exactly at the deterministic equilibrium, and the statistics (b)-(d) appear to have a discontinuous first derivative at the critical point.

[3]The technical details of the simulations: I fix all parameter values, discretize to $J = 800$ agents, and set $\alpha^j = \eta\Phi^{-1}(j/J)$. 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 q small enough to avoid immediate divergence).

[4]Here I allow the simulation to run for 1000 periods and calculate the statistics for the remaining periods, and over a number of simulations with different initial conditions. The averages are over time, 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.

4.4 A Non-Stochastic Simplification

The full model discussed is both stochastic and highly nonlinear. However, one can gain much insight by considering a simplification of the model that does away with the stochasticity and eliminates some of the non-linearity. In this subsection I shall discuss this simplified model and relate the results back to the full model. The inspiration for this simplification is the Kuramoto model (Kuramoto, 1975), which describes synchronization phenomena in general, and has been used in various fields with considerable success (Strogatz, 2000).

Consider a modification to our model: rather than playing the game once within every period $t = 0, 1, \ldots,$ imagine that the players play the game *many* times within each period while holding their beliefs fixed. At the end of each period the agents use all of their observations to update their beliefs as before. The point is that for enough repetitions within each period, the stochastic nature of the learning process (8) is eliminated and the correction to the players' beliefs at the end of every period does not depend on the random realization of z_t. This simplification is certainly not "legal" in any mathematical sense, but is useful for the purpose of gaining intuition on the underlying mechanism.

The learning rule for ϕ (8a) is simplified to

$$\phi^j_{t+1} = \phi^j_t + g_t(\hat{y}_t - \phi^t_j),$$

where $\hat{y}_t$ is the within period average of y_t (over all realizations of z_t). Notice that this rule also implies that for any j, j',

$$|\phi^j_{t+1} - \phi^{j'}_{t+1}| = (1 - g_t)|\phi^j_t - \phi^{j'}_t| = |\phi^j_0 - \phi^{j'}_0| \prod_{s=0}^{t}(1 - g_s) \xrightarrow{t\to\infty} 0,$$

i.e. after enough time the simplified learning process leads to a common belief about ϕ.

Next, consider equation (11). It gives a formula for y_t that is the logarithm of an average of exponents. Schematically, this can be written as:

$$(1 - \theta)y = \log \int_0^1 e^{A(\phi^j,\xi^j)+B(\xi^j,\langle\xi^i\rangle)\cdot z} dj = \log \int_0^1 e^{A^j+B^j\cdot z} dj.$$

The right-hand-side is the log of the moment-generating function of the argument of the exponent. Assuming that the moments are finite, we can expand:

$$(1 - \theta)y = \left\langle A^j \right\rangle + \left\langle B^j \right\rangle \cdot z + \frac{1}{2}\,\mathrm{Var}[A^j + B^j \cdot z] + \text{higher moments}.$$

Combining this with the previous simplification and (8b), we get the learning dynamics:

$$M^j \cdot \xi^j_{t+1} = \xi^j_t + \frac{g_t}{1 - \theta}\left(\left\langle B^i_t \right\rangle - (1 - \theta)\xi^j_t + \mathrm{Cov}[A^i_t, B^i_t] + \cdots\right).$$

We shall return to the above setup, but for now, let us ignore drop all the higher-order terms. Replacing the actual expression for $\langle B_t^i \rangle$ in the above, we get

$$M^j \cdot \xi_{t+1}^j = \xi_t^j + g_t \left(\frac{1-\lambda}{\theta} \left\langle m(\|\xi_t^i\|^2) \right\rangle \left\langle \xi_t^i \right\rangle - \xi_t^j \right).$$

Using the definition (10a), and after some manipulation,

$$\frac{1-\lambda}{\theta} m(\|\xi_t^i\|^2) = 1 - \frac{1}{\theta} \left[1 + \frac{\lambda(\theta(1-2\lambda)+\lambda)\sigma_\epsilon^2}{(1-\lambda)^2(\|\xi^i\|^2 - \xi^{S2})} \right]^{-1},$$

which allows us to rewrite

$$M^j \cdot \xi_{t+1}^j = \xi_t^j + g_t \left(\left\langle \xi_t^i \right\rangle - \xi_t^j \right) - \frac{g_t}{\theta} \left\langle 1 + \frac{\lambda(\theta(1-2\lambda)+\lambda)\sigma_\epsilon^2}{(1-\lambda)^2(\|\xi^i\|^2 - \xi^{S2})} \right\rangle^{-1} \left\langle \xi_t^i \right\rangle.$$

Next I would like to consider the continuous-time limit of this equation. To do this, we take the limit $\Delta t \to 0$, and also $g_t \to 0$ (that determines how fast agents learn), and $\alpha^j \to 0$ (that determines how agents' interpretation of the sunspot changes in time), keeping the ratios constant. In a slight abuse of notation, I call the limit $g_t/\Delta t \equiv g$ and $\alpha^j/\Delta t \equiv \omega^j$, to arrive at

$$\dot{\xi}_t^j = \begin{pmatrix} 0 & -\omega^j \\ \omega^j & 0 \end{pmatrix} \cdot \xi_t^j + g \left(\left\langle \xi_t^i \right\rangle - \xi_t^j \right) - \frac{g}{\theta} \left\langle 1 + \frac{\lambda(\theta(1-2\lambda)+\lambda)\sigma_\epsilon^2}{(1-\lambda)^2(\|\xi_t^i\|^2 - \xi^{S2})} \right\rangle^{-1} \left\langle \xi_t^i \right\rangle. \quad (13)$$

4.4.1 Solving the nonstochastic equation

The system (13) has two important steady-state solutions. The first is the trivial $\xi_t^j = 0$. To obtain the second solution, let us assume without loss of generality that $\langle \xi_t^i \rangle$ is entirely in the 1-direction. The 2-component of (13) gives $\xi_t^{2j}/\xi_t^{1j} = \omega^j/g$. Therefore, define $\tan\psi^j = \omega^j/g$, and consider the solution $\xi_t^j = R^j(\cos\psi^j, \sin\psi^j)'$. The remaining equation is:

$$R^j = \cos\psi^j \left\langle R^i \cos\psi^i \right\rangle \left[1 - \frac{1}{\theta} \left\langle 1 + \frac{\lambda(\theta(1-2\lambda)+\lambda)\sigma_\epsilon^2}{(1-\lambda)^2(R^{i2} - \xi^{S2})} \right\rangle^{-1} \right],$$

i.e. R^j must be proportional to $\cos\psi^j$ (which is exogenously determined), therefore let $R^j = C\cos\psi^j$, and we are left with

$$1 = \left\langle \cos^2\psi^i \right\rangle \left[1 - \frac{1}{\theta} \left\langle 1 + \frac{\lambda(\theta(1-2\lambda)+\lambda)\sigma_\epsilon^2}{(1-\lambda)^2(C^2\cos^2\psi^i - \xi^{S2})} \right\rangle^{-1} \right]. \quad (14)$$

For a given distribution of ω^j (and therefore ψ^j) equation (14) determines C. Furthermore, since the right-hand-side is decreasing in C and becomes negative for large enough C (since $\theta < 1$), it is enough to evaluate the right-hand-side at $C = 0$ to determine if a solution exists. After some algebraic manipulation one arrives at the condition

$$\left\langle \cos^2 \psi^i \right\rangle = \left\langle \frac{1}{1 + (g/\omega^i)^2} \right\rangle > \frac{\lambda}{1 - \lambda}. \tag{15}$$

Thus,

Theorem 1. *The simplified model equation* (13) *admits a trivial steady-state solution,* $\xi_t^j = 0$*. If the distribution of* ω^j *is that* (15) *is satisfied, then there also exists a family of solutions:*

$$\xi_t^j = C \cos \psi^j \begin{pmatrix} \cos(\psi^j + \psi^0) \\ \sin(\psi^j + \psi^0) \end{pmatrix}, \tag{16}$$

for any ψ_0 *and* C *is the unique solution to* (14)*. Furthermore, both solutions are locally (Lyapunov) stable.*

The stability property is trivial: since a change in a single ψ^j (a small mass) has a negligible effect on the terms inside the $\langle \cdots \rangle$ operators, the equation for a small deviation is simply

$$\dot{\xi}_t^j = -\begin{pmatrix} g & \omega^j \\ -\omega^j & g \end{pmatrix} \cdot \xi_t^j + \text{constant}.$$

Since the eigenvalues for this equation, $-g \pm i\omega^j$, have negative real part, local stability is assured.

Note that (16) describes a circle of radius $C/2$ centered at $v = (C/2)(\cos \psi^0, \sin \psi^0)$. The agents will always end up distributed on such a circle, with the exact distribution determined by the distribution of ω^j. This also means that $v \cdot \xi^j > 0$ for all j, so that they all agree that $v \cdot z_t > 0$ implies higher output this period. That is, they agree on the sign, but not the magnitude, of the impact of a sunspot realization parallel to v. They will not agree on the relevance of the perpendicular component: some will believe that $(I - vv')z_t > 0$ implies higher than average output and some lower.

In numerical simulations I find that the non-trivial is also a global attractor if it exists, and that otherwise the trivial solution is an attractor. Indeed, the system (13) tends to one of the two solutions for all the specifications that I have tried.

4.5 From the Deterministic to the Stochastic Model

In order to arrive at (13) we ignored the higher moments of the belief distribution. This, however, is only needed in order to find the explicit condition (15) that determines when the nontrivial solution exists. In fact, in the fully nonlinear version the equation in the direction perpendicular to $\left\langle \psi^i \right\rangle$ still requires that $\xi_t^{2j}/\xi_t^{1j} = \omega^j/g$, and therefore leads to a solution of the form $\xi_t^j = C \cos \psi^j (\cos \psi^j, \sin \psi^j)'$. The equation for C is now more complicated, but for a given distribution of ω^j, one can still determine if a solution exists numerically and find it. The local stability properties also work in exactly the same way: the nontrivial solution is stable when it exists, and appears to be a global attractor; and when it does not, the trivial solution is the attractor.

Figure 3 shows examples of the stead-state solutions of the fully nonlinear but non-stochastic model for different distributions of ω^j. In all examples $\omega^j \sim \mathrm{U}(-\eta, \eta)$, which makes it possible to plot the support of ξ^j. We can see that for small η, the stead-state solution resembles a rational-expectations solution of type S. Instead of the agents agreeing on some point on the ξ^S-circle they are dispersed on a small arc close to the circle. For η large, but below the critical point, the beliefs form a small loop close to the origin. In both cases the belief-dispersion is low and the average belief is close to the S or C REE. For intermediate values, we get a much larger dispersion of beliefs, and the average belief takes a value within the circle. This is all in agreement with the properties of the full stochastic model (as described in figure 2).

The difference between this model and the full stochastic one is in the information based on which learning happens: in the deterministic model the learning is based on the ergodic distribution of z_t, while in the full model it is based on a single realization. We can interpolate between the full model and the deterministic one by modifying the number of realizations of z_t between each learning episode (in the deterministic model it is infinite). Changing the distribution of z_t changes the steady-state in the deterministic model, so we cannot draw impulse-response-functions (IRFs) in the standard sense, but we can use the idea of interpolation from slow to fast learning to define a similar construct: begin at $t = 0$ with the steady-state solution of the deterministic model and play one period ($t = 1$) of learning as in the original model. After that, for all periods $t > 1$, the learning is again as in the deterministic model. This approach can be thought of as an impulse to the learning speed, or more accurately to the stochasticity that is introduced by the difference between the realization of z_t and its ergodic distribution.

Figure (4) shows the single-period responses of beliefs, ξ^j, to such a shock when the realization is in the direction parallel and perpendicular to the average belief $\langle \xi^i \rangle$. The reader is invited to view the IRF over multiple periods in a video attachment to this paper[5], but it can also simply be summarized by saying that the beliefs simply decay back to the steady-state. Notice that perpendicular shocks make the distribution of beliefs narrower, and parallel shocks make it wider regardless of whether the projection of z_t on $\langle \xi^i \rangle$ is positive or negative.

The IRFs help us understand another feature of the full model. As mentioned above, the steady-state distribution is more dispersed for intermediate values of η than for low and high values. Therefore, each 'shock' generates a larger movement in beliefs for intermediate values of η, resulting in the type of stochastic behaviour seen in the middle column of figure 1.

[5]The video is available at `https://youtu.be/Xn2DR-CmWTg`

4.6 Interim Summary

To summarize the results so far, I postulated a model in which agents face multiple sources of amorphous non-fundamental random processes, and the nature of these processes precludes explicit agreement on how to use them in the decision making process. The extent of how amorphous the processes are is captured in the dispersion of ω^j. I find that:[6]

1. Without any dispersion in ω^j (the distribution is a singleton), the agents converge over time to the S or C rational-expectations-equilibria found in Benhabib et al. (2015).

2. When the dispersion is beyond some critical value (which can be calculated for any particular family of distributions), the agents learn to play the C equilibrium.

3. Below this critical value, in the deterministic model, agents converge on a solution of the form (16), where the value of C can be determined by solving equation (14) for the particular distribution (or an equivalent equation in the nonlinear model).

4. In the full model, when dispersion is just below criticality or just above zero, the system tends toward a solution that is similar to the one mentioned in the pervious item. The beliefs fluctuate due to the random realizations of the sunspot, but overall the belief dispersion is low. Output in such an economy will be close to $y_t = \phi^C$ for high dispersion, and distributed approximately $y_t \sim \mathrm{N}(\phi^S, (\xi^S)^2)$ for low dispersion.

5. For intermediate values, the above solution is not a good approximation since realizations of the sunspot shift beliefs to create large fluctuations in the average belief. The economy will fluctuate between periods of high output and low volatility (high $\langle \phi_t^i \rangle$, low $\langle \xi_t^i \rangle$), and periods of low average output and high volatility.

5 Discussion

The first conclusion of this paper is that synchronization on sunspot models does not require agents to be able to explicitly agree 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 the agents' actions, which results in output being more volatile. We can, in fact,

[6] Based on numerical simulations for the full model and analytical results for the simplified deterministic version.

translate the results for low-, high-, and mid-dispersion in ω^j into three corresponding macroeconomic scenarios: (a) output is constantly volatile, (b) output is 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.

My model is clearly too stylized to be compared to data as is – in particular, it lacks capital that generates more persistence in the real economy – but these are all in principle predictions that can be tested empirically. At the very least, it demonstrates the sorts of phenomena that learning can create when the information is not very easily and clearly defined.

5.1 Will they ever learn?

The agents in our model are using a misspecified learning model: they are not considering a linear 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 can profit by making superior forecasts. However, this would require a very clever agents, 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 like a tall order.

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 (13) 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, i 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 ω^j, and examples where they are the ones close to the mean. Therefore, it is not clear if dropping the worst forecasters leads to higher or lower eventual dispersion.

More specifically, when dispersion is low, the agents with the ω^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 ω^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 I find that even over samples of 30-40 periods, sometimes it is actually the agents close to the mean that 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.

5.2 The Kuramoto Connection and the Matthews-Strogatz Model

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

In section 4.4 I derived equation (13)

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

This equation bears similarity to an equation studied by MS, which can be written as:[7]

$$\dot{\xi}_t^j = \begin{pmatrix} 0 & -\omega^j \\ \omega^j & 0 \end{pmatrix} \cdot \xi_t^j + g\left(\langle \xi_t^i \rangle - \xi_t^j\right) - (\|\xi_t^j\|^2 - {\xi^S}^2)\xi_t^j. \tag{17}$$

The first two terms are identical. To first order, the last term of (13) is proportional to $(\langle \|\xi_t^i\|^2 \rangle - {\xi^S}^2) \langle \xi^i \rangle$, which is the same as (17) except that the latter has ξ_t^j instead of $\langle \xi_t^i \rangle$.

Both my equation and (17) lead to a steady-state solutions where agents are synchronized and a trivial $\xi_t^j = 0$ solution. Additionally, 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 my model. The reason for the difference is due to the last term of the equation. In equation (17) 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 in a moving in a circle of radius ξ^S, with constant angular velocity ω^j. In (13) 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.

References

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.

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

[7] The notation was changed to conform to the notation my paper

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

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.

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.

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.

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.

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.

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

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.

Figures

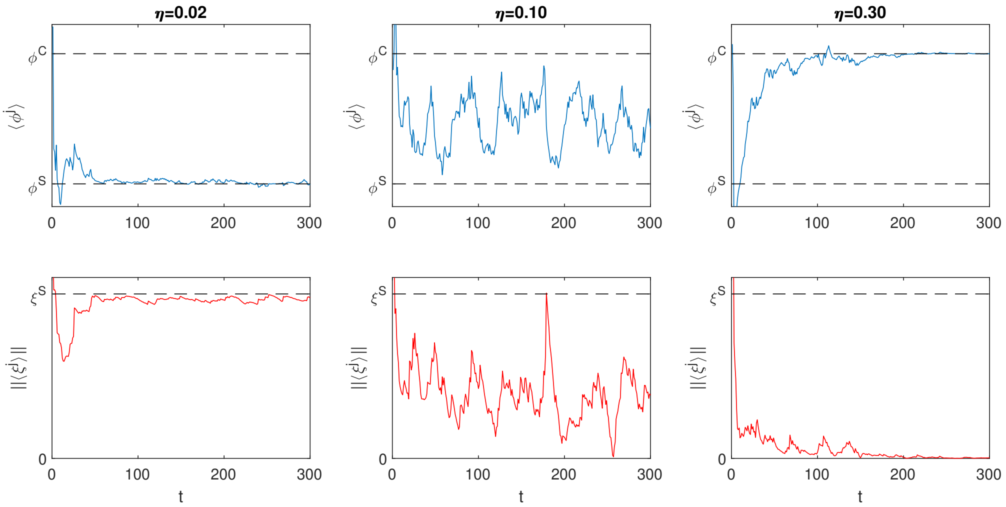

Figure 1: The left, center, and right columns, are each an example of a simulation of the full system for low, mid and high values of η^2 respectively. In each column the top graph displays the evolution of the average belief $\left\langle \phi_t^j \right\rangle$. The middle graph shows the norm of the 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`

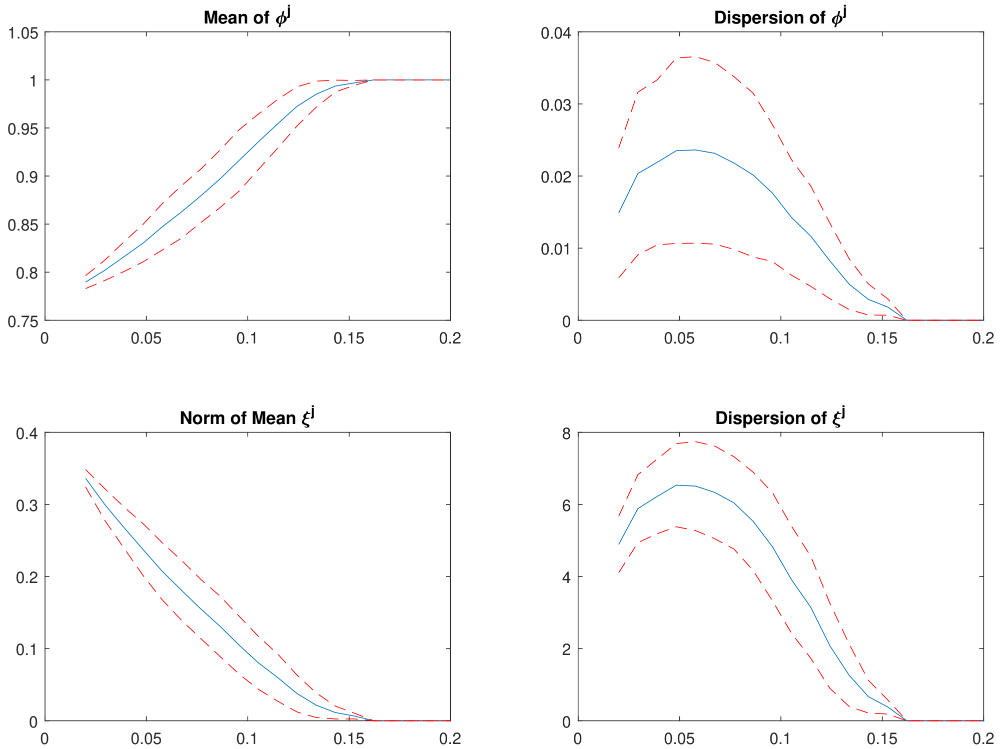

Figure 2: Summary statistics for simulations of the main model. The horizontal axes display values of η, and all other parameters are kept constant as in the previous figure.

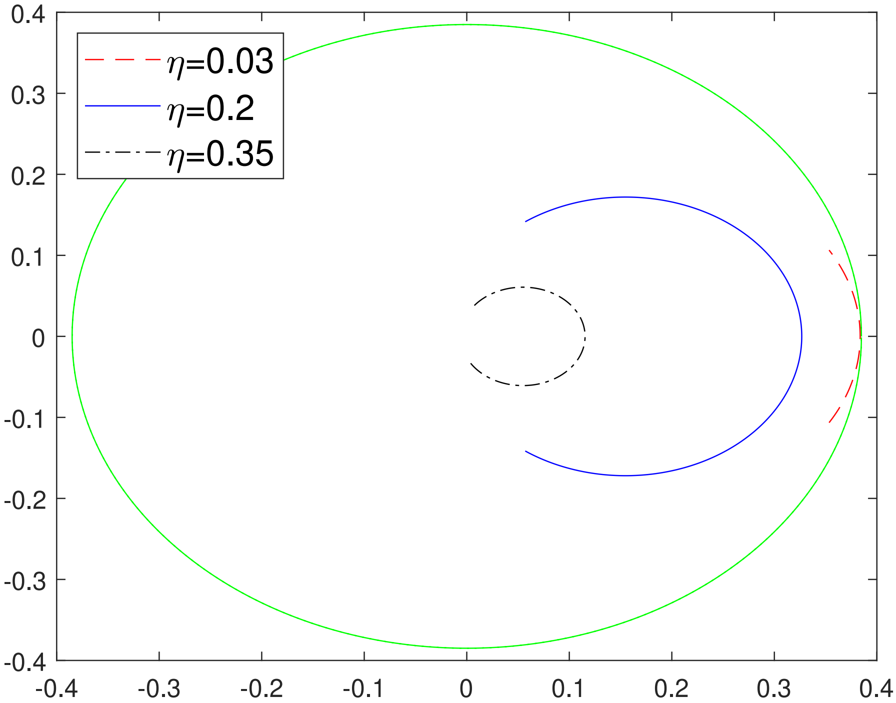

Figure 3: The steady state solution of the steady-state solution to the (fully nonlinear) non-stochastic model for different values of η. The outer circle is of radius ξ^S.

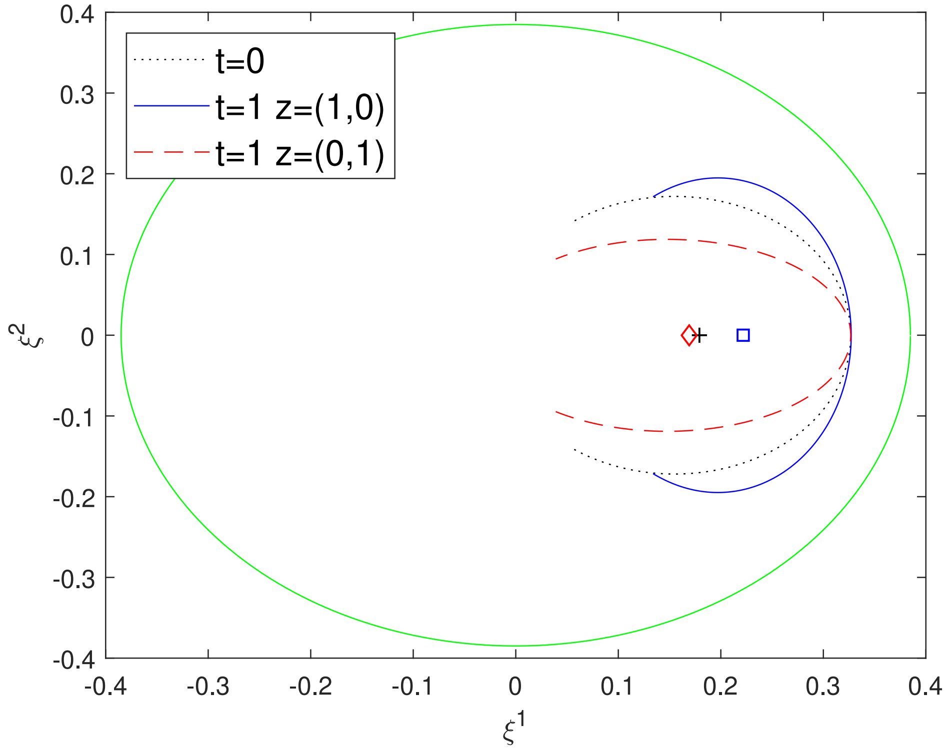

Figure 4: The 'Impulse Response Functions' described in subsection 4.5. The dotted black line is the steady-state solution to the (fully nonlinear) non-stochastic model. The solid and dashed lines show how the ξ^j distributions react to a shocks parallel and perpendicular to $\langle \xi^i \rangle$ respectively. The average belief at $t = 0$ and after the parallel and perpendicular shocks are denoted by a plus, square and diamond respectively.