

Are Sunspots Learnable? An Experimental Investigation in a Simple Macroeconomic Model

Jasmina Arifovic, George Evans and Olena Kostyshyna*

August 20, 2019

Abstract

We conduct experiments with human subjects in a model with a positive production externality in which productivity is a nondecreasing function of the average level of employment of other firms. The model has three steady states and a sunspot equilibrium that fluctuates between the high and low steady states. Steady states are payoff ranked: low values give lower profits than higher values. We investigate whether subjects can learn a sunspot equilibrium. We observe coordination on the extrinsic announcements in our experimental economies. Cases of apparent convergence to the low and high steady states are also observed.

JEL Categories: D83, G20

Keywords: sunspots, learning, experiments with human subjects

*Arifovic: Simon Fraser University, Department of Economics, 8888 University Drive, Burnaby, BC, V5A 1S6, Canada, arifovic@sfu.ca. Evans: University of Oregon, 1285 University of Oregon, Eugene, OR, USA and University of St. Andrews, gevas@uoregon.edu. Kostyshyna: Bank of Canada, 234 Wellington Street Ottawa, ON, K1A 0G9, Canada, kost@bankofcanada.ca. The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the Bank of Canada. We would like to thank Heng Sok and Brian Merlob for helpful research assistance.

1 Introduction

In this paper, we present an experimental study of a model with multiple payoff-rankable equilibria, including a sunspot equilibrium. The objective of this work is to explore whether subjects can coordinate on a sunspot equilibrium and under what circumstances such coordination arises.¹

The possibility of multiple rational expectations equilibria (REE) is a central issue in macroeconomics, raising the question of coordination and forcefully emphasizing the importance of expectations. In models with multiple REE the equilibria typically include stationary sunspot equilibria (SSE) solutions, in which agents' actions are conditioned on an extraneous random variable (Cass and Shell (1983)). This phenomenon can be found, for example, in overlapping-generations models of money, real business cycle-type models with non-convexities, New Keynesian models, endogenous growth models, and monetary search models. SSEs provide an interpretation of economic fluctuations as at least in part due to self-fulfilling changes in household and firm expectations.

A question of considerable interest in macroeconomics is whether agents can coordinate on SSEs. For example, Farmer (1999) and Clarida et al. (2000) have argued that SSEs may provide an explanation for business cycle fluctuations. Sunspot-driven fluctuations are more plausible in models in which SSEs are stable under adaptive learning. This possibility has been demonstrated by Woodford (1990), Evans and Honkapohja (1994), Evans et al. (1998), Honkapohja and Mitra (2004), Evans et al. (2007), Shea (2013), McGough, Meng and Xue (2013), as well as in numerous other papers.

The potential importance of SSEs continues to generate interest. For example, in the context of multiple steady states due to the interest-rate zero lower bound, stressed by Benhabib, Schmidt-Grohe and Uribe (2001), the possible relevance to the Great Recession of the existence of associated SSEs has been emphasized by Mertens and Ravn (2014). In a new generation

¹Experimental studies of models with multiple equilibria have been done in different environments, including overlapping generations models with money (Marimon and Sunder (1993, 1994, 1995), Lim et al (1994), Arifovic (1995, 1996), simultaneous games (Cooper et al. (1990, 1992))), effort coordination games (van Huyck et al. (1990)), optimal growth models with nonconvex production technology (Lei and Noussair (2007)), and bank runs (Garratt and Keister (2009), Schotter and Yorulmazer (2009)), Corbae and Duffy (2008), Arifovic et al. (2012)). Duffy (2016) and Arifovic and Duffy (2018) provide surveys of this literature.

of bubbles models, SSEs have recently been suggested as a source of asset price volatility – see Gali (2014), Farmer (2015) and Miao, Shen and Wang (2018).

Given this background, although there is of course important related literature discussed below, there has been relatively little research on SSEs in experimental macro settings. Our aim, therefore, is to determine whether it is possible for subjects in an experimental setting to coordinate on an SSE within a simple and transparent economy. The sunspot equilibria we consider are in the spirit of macroeconomic models that obtain cyclical fluctuations in settings in which multiple equilibria arise through strategic complementarities, externalities or monopolistic competition. Examples include the coordination failure models of Cooper and John (1988), the "animal spirits" model of Howitt and McAfee (1982) based on transactions externalities, and the multiplicities and sunspot equilibria found in non-convex real business cycle and endogenous growth models when positive production externalities or monopolistic competition are present, as in Benhabib and Farmer (1994) or Evans et al. (1998).

Because in experiments it is crucial to keep the setting simple, we develop a stripped-down, essentially static, framework in which a positive production externality leads to the existence of multiple steady states and sunspot equilibria. We also ensure that the decision the subjects need to make each period is simple and transparent: to forecast aggregate employment. Although the framework is very simple, this makes it possible to study whether agents can coordinate on a sunspot equilibrium in which there are economic fluctuations of the type found in the macroeconomic literature just cited.

The model has three steady states. In the language of the macro learning literature, the low-employment and high-employment steady states are E-stable (and thus stable under adaptive learning rules), while the middle steady state is not E-stable. There also exists an E-stable sunspot equilibrium on which we concentrate in our experiments. This sunspot equilibrium involves fluctuations between values near the two E-stable steady states, i.e. between low and high steady states. These two steady states are payoff ranked: the high-employment steady state has higher profits than the low-employment steady state does. This feature presents an additional challenge for coordination on a sunspot equilibrium because it implies switching from high-payoff to low-payoff outcomes.

The payoff rankability of the certainty equilibria motivates two different experimental treatments. In the first, the subjects' payoff is based on prof-

its, which presents the challenge discussed above. In the second treatment, the subjects' payoff is based on the forecasting accuracy of their forecasts (forecast squared error), and thus the two certainty equilibria are not payoff-ranked in this case. These two treatments allow us to address question of whether it is easier to coordinate on sunspot equilibria if they yield the same payoffs as certainty equilibria.

Our motivation for using two treatments with different payoffs based on profits and forecast squared error is also based on the observation about accuracy of inflation expectations of firms and consumers relative to the accuracy of inflation expectations of professional forecasters. Although these stylized facts relate to inflation expectations, we believe that the underlying link to the optimization function of these agents is relevant to our understanding of expectations about employment in our experiments. Based on the firm survey in New Zealand, Coibion et al. (2018) document that inflation expectations of firm managers are higher than those of professional forecasters and exhibit more disagreement than those of professional forecasters. Ehrmann et al. (2015) show that household inflation expectations in Michigan Survey of Consumers are above those of professional forecasters. Coibion et al. (2018) argue that inaccuracy in firms forecasts could be attributed to the inattention linked to the characteristics of the firms such as duration until the next price change, slope of the profit function and share of exports in their sales. In other words, firms do not pay attention to the dynamics in aggregate inflation because it is not directly linked to their decision making which determines their payoffs – profits. Similar reasons were suggested to explain bias in consumers inflation expectations – they care about inflation for their own consumption bundle, not aggregate level. Professional forecasters, on the other hand, care directly about the accuracy of their forecasts, and, therefore, their forecasts are more accurate than those of firms or consumers. In our experiments, we implement two treatments with different payoffs to capture different optimization functions – profits for firms and forecast squared error for professional forecasters. Our hypothesis is that forecasts are more accurate in the treatment with payoff based on the forecast squared error than forecasts in the treatment with payoff based on the profits. Our findings are consistent with this hypothesis.

The main result of our experiments is that subjects can coordinate on sunspot announcements in both treatments. We also observe coordination on high- and low-employment steady states where subjects disregarded sunspot announcements. While the experiments show frequent examples of coordi-

nation on the sunspot announcements in both treatments, in the treatments with forecasting accuracy, the subjects' forecasts and outcomes are closer to the equilibria corresponding to the announcement, i.e. the coordination is more accurate (consistent with the hypothesis outlined above). We should note that less accurate forecasts in our profit treatment can be attributed partly to the flatness of the profit function, consistent with the finding in Coibion et al. (2018) that firms with with steeper profit function make smaller forecast errors.

The paper is organized as follows. In section 2 we describe related literature. In section 3, we describe the model. In section 4, we present the design of the experiments. Section 5 describes the results of the experiments, followed by section 6 which presents the discussion of adaptive learning in the experiments. Section 7 concludes the paper.

2 Related literature

There are several related papers in the experimental literature that have looked for sunspot equilibria in different settings. Marimon et al. (1993) perform an experimental study based on the overlapping generations model with money. For an appropriate specification of preferences, this model can have multiple regular perfect foresight cycles and sunspots.² In their set-up the model has a unique steady-state equilibrium and a two-period cycle equilibrium (which can be viewed as a perfect foresight sunspot). Marimon et al. (1993) find that while the presence of extrinsic shocks (sunspots) is not sufficient in itself to generate cyclic patterns in behavior, cyclic behavior is observed when agents are trained to experience it together with a sunspot at the beginning of the experiment. During training periods, the cyclic behavior is achieved by a real shock to the number of agents in a generation that amounts to varying endowments; this shock is not observable by the subjects. The change in the number of agents is accompanied by a sunspot - a blinking square of a corresponding color on the computer screen. The number of agents in a generation is kept fixed after the training period, but the colored square continues to appear on the screens during the input stage and during the display of the results. The display of history is also color-coded. Marimon et al. (1993) find that the price fluctuations are smaller during the experiment than those during the training periods, but that the price fluctuations persist.

²See, for example Azariadis and Guesnerie (1986).

Thus there appears to be coordination on a cyclic equilibrium, though it is difficult to tell, given the length of the experiments, how long this cyclic behavior would continue. The cyclic behavior tends to trail off towards the end, and thus it is not clear that the SSEs are durable.

Duffy and Fisher (2005) study sunspot equilibria in a microeconomic setting in which heterogeneity of agents, both buyers and sellers, plays a central role, and motivates trades. They consider two mechanisms that have different information flows: the closed-book call market and double auction. In their set-up, the marginal valuations of buyers and marginal costs of sellers depend on the median price, and thus the payoffs of agents depend on the actual price realized in the market. This feature turns the set-up into a coordination game, with two equilibria, but by design the two equilibria are not Pareto rankable: some subjects are better off in one equilibrium, whereas other subjects are better off in the other equilibrium. Duffy and Fisher (2005) find that subjects can coordinate on a sunspot equilibrium based on a public announcement, though this result is sensitive to semantics, i.e. the wording of the announcements, and institutions: sunspot equilibria are observed in all sessions with call markets, but in less than half of the sessions with double auction.

Fehr et al. (2018) study a two-player coordination game with multiple equilibria: players pick a number between zero and one hundred, and the payoffs are determined as the squared deviations from the other player's choice. All equilibria have the same payoff, but choosing 50 is a risk-dominant equilibrium. Fehr et al. (2018) use public and/or private signals of different precision and study experimentally how they affect which equilibrium subjects coordinate on³. Their sunspots are semantically salient (they are the numbers just like the strategies) and they do not have training periods which distinguishes this paper from Marimon et al. (1993), Duffy and Fisher (2005) and ours. They find that sunspot equilibria arise endogenously in case of public signals that are easy to aggregate. They also observe sunspot-like behavior in case of highly correlated private signals. When both public and private signals are provided, full coordination on public signal is disturbed as some subjects condition their decisions on private signals. In our setup, we have

³The global games literature has shown that the existence of multiple equilibria, in coordination-type games, is sensitive to the presence of private and public signals (for example, Heinemann et al. (2004)). These signals can in effect operate as sunspots.

only public announcements, and our sunspot announcements are less semantically salient ("Low/high employment is forecast this period") than in Fehr et al. (2018) as the announcement does not specify the values of low/high employment. In our model sunspot equilibrium is payoff-dominated by the high-employment equilibrium, while in Fehr et al. (2018) sunspot equilibria do not present welfare losses.

Beugnot et al. (2012) show that in a coordination game with two strategies with payoff-ranked equilibria the subjects coordinate on the payoff-dominant equilibrium in the treatment without sunspots. The introduction of sunspots disrupts the coordination on the payoff-dominant equilibrium and leads to off-equilibrium outcomes ('dis-coordination'). However, subjects never coordinate on the payoff-dominated equilibrium⁴. This study has one feature in common with our paper – payoff-rankability of certainty equilibria; however, this study is based on a normal form game with two strategies, while ours is based on a macroeconomic model with a continuum of strategies. In addition, coordination on a sunspot is not observed in their experiment.

There are also experimental studies of correlated equilibrium, a concept that is related to sunspot equilibrium. For example, Duffy and Feltovich (2010) study the game of Chicken with private third-party recommendations and show that subjects follow the third-party recommendations if they are derived from a correlated equilibrium which is payoff-enhancing relative to Nash equilibria. While related, the concept of private recommendations in the context of correlated equilibrium in their study is different from public sunspot announcements in ours. In our paper, the signal is publicly announced and informs all subjects about the forecast that is randomly generated. This is different from private recommendations of strategies to subjects. Furthermore, while in Duffy and Feltovich (2010) subjects only coordinate on correlated equilibrium that is payoff-enhancing relative to Nash equilibria, our subjects are able to coordinate on the sunspot equilibrium which is payoff-dominated by the high-employment equilibrium.

The focus of our paper is different, and motivated instead by the macro literature. We look at a simple macro set-up with a positive production externality and multiple steady states. In this setting there exist SSEs, and only

⁴Previous experimental work without sunspots has shown convergence to payoff-dominated equilibria in both normal form games (for example, Cooper et al. (1990), Van Huyck et al. (1990)) and optimal growth model with increasing returns (Lei and Noussair (2007)).

a subset of these SSEs can be locally stable under adaptive learning rules. We are interested in whether, in this context, adaptively stable SSEs can be reached and sustained experimentally in the lab. In line with the macro literature, in which it is plausible that agents may not have complete knowledge of the full economic structure, we provide qualitative but incomplete quantitative information to subjects of the specification of the economy.

Our experiments are also related to the experimental studies of expectations formation. In these studies, subjects also do not know the underlying structure of the economy and need to form expectations of endogenous variables using observed past realizations in the economy. For the surveys of this macro experimental literature see Duffy (2016) and Arifovic and Duffy (2018).

Within the general literature, our work is closest to Marimon et al. (1993) and Duffy and Fisher (2005), both of which study experimentally whether the equilibrium selection can be driven by extraneous public announcements. Like Marimon et al. (1993), we use a macroeconomic setting to generate sunspots, but in contrast to their framework, which looks at SSEs near cycles in a neighborhood of the indeterminate steady state, in particular at SSEs near a 2-period cycle, we look at SSEs near a pair of distinct steady states. Duffy and Fisher (2005) also look at SSEs near distinct equilibria, but in our setting the steady states are Pareto rankable. Also we do not require heterogeneity of agent types, and thus the sunspot equilibria that we examine have the interpretation of switches between high and low levels of aggregate output, resulting from waves of optimism or pessimism driven by extraneous public announcements.

3 Model

3.1 Description of the economy

We use a simple, stylized macroeconomic set-up in which production externalities can generate multiple steady states and SSEs. Our framework is characterized by the contemporaneous production externality in which productivity of a firm is increased, over a range, by higher activity in other firms.⁵

⁵Our set-up is closely related to the “Increasing Social Returns” overlapping generations model described on pp. 72-81 in Evans and Honkapohja (1995). To keep the framework

In period t , each firm hires workers, n_t , to produce output, y_t using the production function

$$(1) \quad y_t = \psi_t \sqrt{n_t},$$

where ψ_t indexes productivity. Profit for the firm is computed as output minus labor costs. The cost of a unit of labor is wage w , and thus the firm maximizes profit

$$(2) \quad \Pi_t = \psi_t \sqrt{n_t} - wn_t,$$

The level of productivity ψ_t depends on the average level of employment *across all other firms* (not including firm's own employment)⁶. We will call average employment of other firms \bar{N}_t . The firm decides on employment, n_t , before knowing productivity, ψ_t , because it does not know the average employment of other firms, \bar{N}_t , when its decision is made. A firm is more productive when other firms are operating at a high level of employment.⁷ Specifically, productivity, ψ_t , depends on the average employment of other firms, \bar{N}_t , as follows⁸:

$$(3) \quad \begin{aligned} \psi_t &= 2.5 && \text{when } \bar{N}_t \leq 11.5 \\ \psi_t &= 2.5 + (\bar{N}_t - 11.5) && \text{when } 11.5 < \bar{N}_t < 13 \\ \psi_t &= 4 && \text{when } 13 \leq \bar{N}_t \end{aligned}$$

This model can be thought of as a very simple and stylized macroeconomic model, with a single consumption good, no capital or other means of saving,

as simple as possible, for laboratory experiments, we use a version that eliminates the dynamic optimization problem required in overlapping generations set-ups, and instead focuses entirely on the contemporaneous production externality.

⁶Another alternative would be to make ψ_t depend on the average level of employment of all firms. Neither implementation matters in a rational expectations equilibrium (REE) (under competitive assumptions) but they can affect the behavior of the experimental economy as will be discussed later.

⁷As emphasized by Farmer (1999) in the context of RBC models, perfect competition with positive production externalities is formally equivalent to increasing returns to scale production combined with monopolistic competition; for the laboratory model we employ a simple version of the production externality used in Evans and Honkapohja (1995).

⁸Lei and Noussair (2007) study an environment similar to ours. The productivity in their model depends on the aggregate level of capital: if aggregate capital is above the threshold, productivity is high; if aggregate capital is below the threshold, the productivity is low. They find that experimental economies can get into poverty traps with low levels of capital and output.

and in which household utility is such that labor supply is infinitely elastic at wage w .⁹ Firms are owned by households and profits are distributed as dividends back to the households each period. Note that, because we have a repeated static equilibrium model, the household problem is trivial: supply the labor demanded by firms at wage w and consume all income, generated by wages and dividends. In the experiments we therefore focus solely on the firm problem within this setting.

3.2 Equilibria

For this economy, profits are maximized when firms choose:

$$(4) \quad n = \left(\frac{\psi}{2w} \right)^2$$

Depending on the parameters there are (generically) one or three perfect foresight steady states. Within each of the three steady states, all firms hire the same quantity of labor and produce the same level of output. For productivity function (3) and with wages $w = 0.5$, this model has 3 steady states: $n_L = 6.25$ (“low-level”), $n_M = 12.54$ and $n_H = 16$ (“high-level”).

When there are three steady states, stationary sunspot equilibria (SSE) exist between any pair of steady states. For example, in the experiments we randomly generate announcements of “high” and “low” forecasted employment. Letting $A_t \in \{L, H\}$ denote the announcement at time t , where L represents the announcement “Low employment is forecasted this period” and H denotes the announcement “High employment is forecasted this period”, there exists an SSE $n_t = n_L$ if $A_t = L$ and $n_t = n_H$ if $A_t = H$. Other SSEs also exist, including those switching between other pairs of steady states or between all three steady states, as well as the three steady states themselves in which employment is independent of the announcement.¹⁰

Changes in the wage w , as well as changes in employment subsidies or taxes that alter the “effective” wage rate, can bifurcate the system. (Wage

⁹It would be straightforward to generalize the model to allow for less than fully elastic labor supply, as is done in Evans and Honkapohja (1995).

¹⁰Equilibria or SSEs can be constructed that depend on any observable, e.g. equilibria can be constructed that switch between steady states values depending on calendar time or on the past history of aggregate employment. These equilibria can be viewed as limiting SSE. There also exists an SSE in which $n_t = n_L$ if $A_t = H$ and $n_t = n_H$ if $A_t = L$.

subsidies or taxes are assumed offset by lump-sum taxes or subsidies, respectively so that the combined effect is revenue neutral).

When $w = 1$ only the “low-level” steady state exists, and when $w = 0.2$ only the “high-level” steady-state exists. There do not exist SSEs when the effective wage is such that only a single interior steady state exists. In this study, we concentrate on the issue of coordination on SSE, and so we use $w = 0.5$.

We now take up the issue of which equilibria are stable under simple adaptive learning rules that have been widely studied in the macro learning literature.

3.3 Temporary equilibrium framework and E-stability

Sunspot-driven fluctuations are more plausible in models in which SSEs are stable under adaptive learning (Woodford (1990), Evans and Honkapohja (1994), Evans et al. (1998), and Evans et al. (2007)). In our setting only a subset of the existing SSEs can be locally stable under adaptive learning rules. In this section, we establish which SSE is adaptively stable, and we concentrate on it in our experiments.

The optimal choice of n in equation (4) depends on the firm’s expectations of ψ . As ψ depends on the average employment of other firms, \overline{N}_t (equation 3), the optimal choice of n equivalently depends on the firm’s expectation of \overline{N}_t .

If we now drop rational expectations and also the assumption of homogeneous expectations, then the model equations are as follows. In period t , firm i chooses its employment level as: $n_t^i = (\frac{\psi_t^{e,i}}{2w})^2$ where the superscripts e, i denote the expectations of agent i .

Average employment of other firms for firm i is given by

$$(5) \quad \overline{N}_t^i = \frac{\sum_{j \neq i} n_t^j}{K - 1}$$

where there are K firms, and the actual current productivity level of firm i is given by $\psi(\overline{N}_t^i)$ according to equation (3). The output of firm i is $y_t^i = \psi(\overline{N}_t^i)\sqrt{n_t^i}$ and aggregate output is therefore given by

$$Y_t = \sum_i y_t^i$$

Thus, given the profile of time t expectations $\{\psi_t^{e,i}\}_{i=1}^K$, the above equations determine n_t^i, \bar{N}_t^i and Y_t . The profit of firm i at time t is given by

$$(6) \quad \Pi_t^i = \psi(\bar{N}_t^i) \sqrt{n_t^i} - wn_t^i$$

We have so far assumed that the expectations of agents are specified in terms of $\psi_t^{e,i}$. However, since ψ_t is a monotonic function of \bar{N}_t (equation 3), it is equivalent to specify expectations in terms of $\bar{N}_t^{e,i}$. That is, given the profile of time t expectations $\{\bar{N}_t^{e,i}\}_{i=1}^K$, employment levels are given by

$$(7) \quad n_t^i = \left(\frac{\psi(\bar{N}_t^{e,i})}{2w} \right)^2$$

Figure 1 shows a firm's optimal choice of employment as a function of the firm's forecast, as given by equation (7). The above equations then determine \bar{N}_t^i, Y_t and profits Π_t^i .

Because, in our set-up, there are multiple equilibria, including steady states and SSEs, a natural question is: which equilibria are stable under learning? We now briefly examine the stability properties under simple adaptive learning schemes.¹¹ For convenience (this is not essential) assume that firms have homogeneous expectations $\bar{N}_t^{e,i} = \bar{N}_t^f$ concerning the average level of employment of other firms. Their corresponding forecast of their own productivity is then $\psi(\bar{N}_t^f)$ and the optimal choice of employment for each firm is therefore

$$T(\bar{N}_t^f) = \left(\frac{\psi(\bar{N}_t^f)}{2w} \right)^2$$

This is the map illustrated for our numerical example in Figure 1. The fixed points of this map correspond to the perfect-foresight steady states.

Under adaptive learning, consider first the case in which announcements are not present and agents believe they are in a (possibly noisy) steady state in which the average employment of other firms is $\bar{N}^f = \bar{N}^f + \eta_t$, where η_t is an independent zero mean random variable. Each period t they revise their forecasts of average employment of other firms, which they use to determine their employment in period t , according to the adaptive rule

$$\bar{N}_t^f = \bar{N}_{t-1}^f + \gamma_t(\bar{N}_{t-1} - \bar{N}_{t-1}^f),$$

¹¹For details and further discussion of adaptive learning see Evans and Honkapohja (2001).

where γ_t are the “gain” parameters, which might, for example, be fixed at a number γ such that $0 < \gamma_t = \gamma < 1$.¹² This learning rule is a recursive update of expectations based on the past average of average employment of other firms, \bar{N} . This type of beliefs and their recursive representation is frequently used in the learning literature. It can be shown that a steady state $\bar{n} = \bar{N}_t^f = \bar{N}_t^{e,i} = n_t^i$, for all i, t , is locally stable under learning if and only if the derivative $T'(\bar{n}) < 1$. This is known as the E-stability condition.¹³

Thus when there are three steady states $n_L < n_M < n_H$, steady states $\bar{n} = n_L, n_H$ are locally stable, while $\bar{n} = n_M$ is not locally stable under learning. Here “local” means that initial expectations are sufficiently close and “stable” means that $\bar{N}_t^f \rightarrow \bar{n}$ as $t \rightarrow \infty$. While the stated result is asymptotic, the tendency toward convergence should be visible in finite time, in particular for experiments.

The learning rule just described assumes that agents do not condition on announcements. We now turn to that possibility. The adaptive learning rule then is as follows. Let \bar{N}_t^{Hf} denote the time t forecast of the average employment of other firms if A_t , the announcement at t , is H and let \bar{N}_t^{Lf} denote the time t forecast of the average employment of other firms if A_t , the announcement at t , is L . Forecasts over time are revised according to the rule

$$\begin{aligned}\bar{N}_t^{Hf} &= \bar{N}_{t-1}^{Hf} + \gamma_t(\bar{N}_{t-1} - \bar{N}_{t-1}^{Hf}) \text{ if } A_t = H \text{ and } \bar{N}_t^{Hf} = \bar{N}_{t-1}^{Hf} \text{ if } A_t = L, \\ \bar{N}_t^{Lf} &= \bar{N}_{t-1}^{Lf} + \gamma_t(\bar{N}_{t-1} - \bar{N}_{t-1}^{Lf}) \text{ if } A_t = L \text{ and } \bar{N}_t^{Lf} = \bar{N}_{t-1}^{Lf} \text{ if } A_t = H,\end{aligned}$$

where again, for convenience, we are here assuming homogeneous expectations.

The optimal choice of employment, given these expectations is

$$n_t = T(\bar{N}_t^{Hf}) \text{ if } A_t = H \text{ and } n_t = T(\bar{N}_t^{Lf}) \text{ if } A_t = L.$$

It can be shown that an SSE between two steady states is E-stable, and hence locally stable under learning, if and only if both steady states are themselves E-stable. Thus an SSE fluctuating between n_L and n_H is E-stable, since $T'(n_L) < 1$ and $T'(n_H) < 1$, while sunspots fluctuating between n_L and n_M or between n_H and n_M are not stable under learning. Sunspot fluctuations

¹²We need to assume that $\sum_{t=1}^{\infty} \gamma_t = +\infty$.

¹³We exclude from consideration nongeneric cases in which $T'(\bar{n}) = 1$ for a steady state \bar{n} .

between high and low steady states are locally asymptotically stable in the sense that $\bar{N}_t^{Hf} \rightarrow n_H$ and $\bar{N}_t^{Lf} \rightarrow n_L$ as $t \rightarrow \infty$, provided initial conditional expectations for $A_t = L, H$ are sufficiently close to the two steady state values, and provided both announcements are generated infinitely often over time.¹⁴ In our experiments, we will concentrate on this E-stable SSE.

It can *also* be shown that, even when agents allow for announcements in their learning rule, the steady states n_L and n_H are *also* locally stable under learning, and are thus possible outcomes. That is, if initially expectations \bar{N}^{Lf} and \bar{N}^{Hf} are both close to one of the two steady states n_L or n_H , then convergence will be to that steady state, rather than to an SSE. Put differently, an SSE is an outcome of the learning rules given above only if the initial beliefs of agents exhibit an appropriately large difference between \bar{N}^{Lf} and \bar{N}^{Hf} . We will discuss this further when discussing the results below.

4 Design of experiments

As described earlier, for wages $w = 0.5$ there are two stable steady states at $n_L = 6.25$ and $n_H = 16$ as well as an unstable steady state at $n = 12.54$. In the experiments, announcements are generated using Markov transition probabilities chosen so that ‘high’ forecasts are followed next period by ‘high’ forecasts with probability $\pi_{11} = 0.8$ and ‘low’ forecasts are followed by ‘low’ forecasts with probability $\pi_{22} = 0.7$. There is thus an adaptively stable sunspot equilibrium in which employment switches between n_L and n_H depending on the value of A_t . The objective of the experiments is to see whether subjects can coordinate on the sunspot announcements.

In this model, the firm’s profit in the high steady state is 8 which is higher than its profit in the low steady state (3.125). Therefore, it might seem likely that subjects would coordinate on the high steady state. The payoff dominance of the high steady state makes coordination on the sunspot equilibrium challenging. To investigate this point, we also have a treatment in which the payoffs are based on the forecast squared error. When payoffs are based on the forecasting accuracy, the steady states are no longer payoff ranked.¹⁵

¹⁴For additional discussion and details of adaptive learning of SSEs, see Evans and Honkapohja (1994) and Ch. 4.6 and 12 of Evans and Honkapohja (2001).

¹⁵The payoff dominance of one of the steady states in our model is a key difference

Information about the economy We provide descriptive information about the economy without technical details and equations¹⁶. For example, the instructions provide the following information (instructions are available in the online appendix). "The producer hires labor and produces output... The productivity of each producer depends on the average labor hired (employed) by other producers in the market. The average employment of other producers is equal to the sum of the labor hired by each producer in the market divided by the total number of producers. The higher the average labor hired in the market, the higher the productivity of each individual producer." Thus, the subjects know the qualitative relationships between the variables, but not the quantitative ones.

Decision making In each period t , subjects make forecasts of average employment of the other firms $\bar{N}_t^{e,i}$. Their own optimal choice for hiring is determined using (7).

After all subjects submit their forecasts and their employment is determined according to (7), the actual average employment of other firms \bar{N}_t is computed according to equation (5). The level of productivity, ψ_t , is determined based on the average employment of other firms according to equation (3). Each session lasts 50 periods and subjects have this information.

Payoffs We conduct two treatments in which the payoffs of the subjects are evaluated in two different ways. In the first treatment, the payoff is based on the firm's profit computed according to equation (6). We refer to this treatment as the 'Profits' treatment.

In the second treatment, the payoff is based on the forecasting accuracy of the subjects' forecasts. Forecasting accuracy is evaluated as forecast squared error:

$$(8) \quad FSE_t^i = (\bar{N}_t^{e,i} - \bar{N}_t^i)^2$$

And forecasting payoff is computed as:

$$(9) \quad FP_t^i = \max(8 - FSE_t^i, 0)$$

between our experiments and those in Duffy and Fisher (2005).

¹⁶This has become the standard practice in the experimental literature, for example, Lei and Noussair (2002), Hommes et al. (2005a, 2005b), Lei and Noussair (2007), Hommes et al. (2008), Heemeijer et al. (2009), Capra et al. (2009), Bao et al. (2011)

where δ is the maximum payoff when forecast squared error, FSE_t^i , is zero. This value was chosen to match the maximum profit of the firm in the high steady state of the model.

In this treatment, the subjects are rewarded for their forecasting accuracy only: as long as the subjects' forecasts are close to the actual outcomes, they can get the maximum payoff, and so the steady states are not payoff ranked. We refer to this treatment as the 'FSE' treatment.

Announcements The sunspot announcements "Low employment is forecasted in this period" or "High employment is forecasted in this period" appear on the subjects' screens during the input stage of the decisions. The following information is provided to the subjects in the instructions. "At the beginning of each period, you will see an announcement on your computer screen. The announcement will be either "Low employment is forecasted this period" or "High employment is forecasted this period". The announcements are randomly generated. There is a possibility of seeing either announcement, but the chance of seeing the same message that you saw in the previous period is higher than the chance of seeing a different announcement. These announcements are forecasts, which can be right or wrong. The experimenter does not know better than you what employment is going to result in each period. The employment in each period is based on the decisions of all subjects."

The sequence of announcements is randomly generated by the experimenters before the experiment.

Practice periods Each experimental session includes 6 practice periods during which subjects can familiarize themselves with the environment. We also use practice periods for 'training' (conditioning) subjects to experience different equilibria and their payoffs and introduce the sunspot announcements (as is done in Marimon et al. (1993) and Duffy and Fisher (2005)). The training periods are set up such that subjects experience 3 periods of high employment and then 3 periods of low employment with corresponding announcements in each period. The average employment of other firms is predetermined by the experimenters such that the resulting employment in the economy is high or low. The low values are generated as 6.8 plus a random number from a uniform distribution with support $[0, 1]$. The high values are generated as 14.8 plus a random number from a uniform distribution with

support $[0, 1]$. Subjects are not aware that the average employment of other firms is predetermined by the experimenters. After practice periods are over, the first announcement of the experiment is about low employment.

Information on the computer screen The forecast announcement is given at the beginning of each period. At the end of each period, data from past periods, including the last one is presented in the table as well as in the graph that is updated with new observation in each time period. Screenshots are provided in Figures 2 and 3. In the Profits treatment, the table presents the announcement, subject's forecast and actual average employment of other firms, subject's output and labor costs, as well as subject's payoff (Figure 2). In the FSE treatment, the table presents the announcement, subject's forecast and actual average employment of other firms, and subject's forecast squared error and payoff (Figure 3). The graph presents subject's forecast and actual average employment of other firms.

The experimental software was programmed using z-Tree (Fischbacher, 2007). Each subject sat at a personal computer station, and was not able to observe the decisions of other subjects or interact with them. The experiments were conducted at the Economic Science Institute, Chapman University. The participants were recruited using the Chapman online recruiting software.

We ran two treatments - one is with payoff based on the firm's profits (Profits treatment), and the other is with payoff based on the forecast squared error (FSE treatment). We ran 6 sessions of each treatment, with 6 subjects participating in each session (total of 72 subjects). Each subject participated only in one of 12 sessions. At the beginning of each session, subjects were seated at computer stations in random order. The instructions were distributed and read out loud and if a subject had any questions, these were answered in private. Each session lasted on average 70 minutes, including time spent on instructions. In addition to the show-up fee of \$7, subjects were paid based on their payoffs accumulated over 50 periods. The payoffs were expressed in terms of the experimental currency with the exchange rate of 30 experimental currency units per \$1. The average payoff in profits treatment was \$15.62, and the average payoff in the FSE treatment was \$18.50.

5 Results of the experiment

We observe coordination on announcements (sunspots) in both of our treatments. However, in both treatments, there are instances within a session, or the entire session, where we observe a failure of coordination on a sunspot. We first present the results observed in individual sessions for each treatment, and then analyze the data in terms of deviations from the announcements. We also compare the results observed in the two treatments.

5.1 Profits treatment

Figures 4 - 9 present the results of the Profits treatment for each individual session.¹⁷ Each figure consists of two panels. The first panel presents average employment, average forecast and the equilibrium employment corresponding to the announcement.¹⁸ (Note that participants were not given the history of equilibrium employment corresponding to the announcement on their screens. We present these series in our figures for ease of comparison with the actual data.) The second panel presents the percentage deviations of average employment and average forecast from equilibrium employment corresponding to the announcement.

In session 1, the economy follows the announcements closely as can be seen from Figure 4 where average employment is the same as the equilibrium corresponding to announcement in most of the periods except five instances. Figure 4 shows that the percentage deviations from the announced equilibrium are 0 for all periods, except for 7 periods. This means that subjects have coordinated on the announcements. However, there is also evidence of learning during early periods of high-employment announcements. In the first three stretches with high-employment announcements, it takes the economy two to three periods to reach the equilibrium values.

We observe coordination on the announcements in sessions 2, 3 and 4 as illustrated on Figures 5 - 7 but with some departures from the equilibria corresponding to the announcements and somewhat larger percentage deviations than in session 1. We can also see again that learning/adaptation takes

¹⁷We report data for each session to illustrate how close the coordination is or is not, which would be obscured by reporting average values for the treatment because of the variation across sessions. We provide a comparison of the two treatments in Section 5.3.

¹⁸By the “equilibrium employment corresponding to the announcement” we mean n_H if $A_t = H$ and n_L if $A_t = L$.

place. It appears that it is harder to switch from the low to the high steady state, and all the figures show that it takes a bit of time for the economies to reach the high steady state values.

Sessions 5 and 6 have instances of a lack of coordination on the announcements. In session 5 (Figure 8), between periods 12 and 31 there are both high and low announcement stretches in which average employment does not correspond to the announcement. However, after period 32 average employment becomes close to the high equilibrium. Although it is not clear what would have happened if session 5 had continued for more than 50 periods, it appears possible that there would have been convergence to the high steady state.

In session 6 (Figure 9), average employment is again initially in line with announcements, but after period 12 employment during periods of high announcements begins to fall short of the high equilibrium and then eventually, after period 32, average employment becomes close to the low steady state. Again, we do not know what would have happened if session 6 had continued for more than 50 periods, but there might plausibly have been eventual convergence to the low steady state. Other experimental studies without sunspots have shown convergence to payoff-dominated equilibria in normal form games (for example, Cooper et al. (1990), Van Huyck (1990)) and optimal growth model with increasing returns (Lei and Noussair (2007)), but subjects in Beugnot et al. (2012) failed to coordinate on payoff-dominated equilibrium in the presence of a sunspot. As we will discuss in section 6.1, in our setup it is harder to switch from the low to the high steady state than vice versa. This difficulty could be an explanation of why subjects coordinate on low steady state.

In summary, we observe close coordination on the extrinsic announcements in sessions 1-4. However, we also observe a lack of coordination on an SSE in sessions 5 and 6, with apparent convergence to the high steady state in one case and to the low steady state in the other case.

5.2 FSE treatment

Figures 10 - 15 present the results of the FSE treatment. Again, the data for each session is presented in a figure that consists of two panels. The first presents average employment, average forecast and equilibrium employment corresponding to the announcement; and the second presents percentage deviations of average employment and average forecasts from equilibria

corresponding to the announcement.

In sessions presented in Figures 10-12, 14, and 15, the experimental economies exhibit close coordination on the announcements, and the percentage deviations from the equilibrium corresponding to the announcement are zero during almost all periods. In these sessions, the subjects are rewarded for their forecasting accuracy only: as long as the subjects' forecasts are close to the actual outcomes, they can get the maximum payoff, and it does not matter which steady state is the outcome as the steady states are not payoff ranked. We can see better coordination on the announcements and smaller deviations from the equilibrium employment than in the treatment with payoff based on profits. The formal test results are presented in section 5.3.

Figure 13 illustrates the results of session 6 in which we observe that the subjects coordinated on the low-employment steady state by the end of the session. During periods 14-17, 23-25, 32-41 and 47-49 the subjects ignored high-employment announcements and remained in the low-employment steady state. The lack of coordination on the high-employment announcement does not cost these subjects lower payoffs because they are rewarded for their forecasting accuracy only. Therefore, it is less of a puzzle in comparison to the sessions in which the payoffs are based on profits.

In summary, in the sessions with payoff based on FSE we observe both coordination on the extrinsic announcements and coordination on the low-employment equilibrium. It is interesting to observe coordination on announcements in this treatment because the subjects could have ignored the announcements and stayed in one of the two equilibria, and they still would have achieved the maximum payoff. It is a matter of coordination in this game, and the subjects coordinated on the announcements in many sessions.

5.3 Comparison of the two treatments

Next, we analyze the data and compare the two treatments. We want to evaluate how closely the experimental economies coordinate on the announcements and whether there is a difference between the two treatments.

5.3.1 Employment and forecasts

We pool the data on individual employment in periods with the low-employment announcements and in periods with the high-employment announcements

separately, and pool these data for all experimental sessions for each treatment. Table 1 presents the fractions of observations in the ranges containing two equilibria; this table corresponds to the histograms presented in Figures 16 and 17.

The top left panel of Figure 16 represents the histogram of individual employment decisions during periods with high-employment announcements in the FSE treatment and shows that employment is concentrated on the high-employment equilibrium of 16 (83.54% of employment outcomes according to Table 1). The top right panel of Figure 16 presents the histogram of individual employment decisions during periods with low-employment announcements in the FSE treatment and illustrates that the values of employment are heavily concentrated on the low-employment equilibrium of 6.25 (98.23% of outcomes according to Table 1). The bottom left and right panels of Figure 16 present the histograms of individual employment decisions during periods with high-employment (bottom left) and low-employment announcements (bottom right) in the Profits treatment. These histograms also illustrate that the values of employment are very close to the equilibrium values corresponding to the announcements. During periods with high-employment announcements, 76.03% of the employment outcomes are close to the high-equilibrium employment of 16; and during periods with low-employment announcements, 88.76% of the employment outcomes are close to the low-equilibrium employment of 6.25.

We also pool the data on the individual forecasts made in periods with low-employment and in periods with high-employment announcements separately. The top left and right panels of Figure 17 present the histograms of individual forecasts made in periods with high- and low-employment announcements in the FSE treatment and show that the forecasts are heavily concentrated on the respective equilibrium values corresponding to the announcements: 61.83% of forecasts are in the range containing the high-equilibrium employment of 16, and 70.20% of forecasts correspond to the low-equilibrium employment of 6.25. The bottom left and right panels of Figure 17 present the histograms of individual forecasts made in periods with high- and low-employment announcements in the Profits treatment and show that the forecasts are centered around the equilibrium values, but the fractions of forecasts in the ranges containing equilibria are much lower than in the FSE treatment: 28.60% of forecasts are in the range containing the high-equilibrium employment of 16, and 33.46% of the forecasts are in the range containing the low-equilibrium employment of 6.25. In the Profits

treatment, the subjects' performance is not evaluated based on the accuracy of their forecasts, therefore, we observe very high variability in the forecasts. We will explore this in more detail in the next section 5.3.2.

Table 2 presents the data on average, median and standard deviations of employment and forecasts in both treatments.

5.3.2 Deviations from the equilibrium employment

We would like to test how closely subjects coordinate on the announcements. We compute the percentage deviations of employment and forecasts from the equilibrium corresponding to the announcement for all periods and pool the data over all sessions for each treatment. Then we test whether the two treatments are different.

The left panel of Figure 18 presents the cumulative density function (CDF) of the percentage deviations of individual forecasts from equilibrium employment corresponding to the announcements in both treatments. The CDF for the FSE treatment is larger than the CDF for the Profits treatment, which is statistically significant using the Kolmogorov-Smirnov test (with a p-value of 0, and test statistic of 0.3827). This implies that forecasts are closer to the equilibrium values in the FSE treatment than in the Profits treatment. As the subjects are rewarded based on the accuracy of their forecasts in the FSE treatment, their forecasts are closer to the equilibrium values than those in the Profits treatment. As illustrated by Figure 1, when forecasts are below 11.5, employment is constant at 6.25; and when forecasts are above 13, employment is constant at 16. Thus, even if the subjects' forecasts are not equal to equilibrium employment, but are in the appropriate range, their employment outcomes and profits take equilibrium values corresponding to the low or high equilibrium. Thus the subjects in the Profits treatment do not have to make very accurate forecasts to arrive at equilibrium employment and profits. The shape of the employment function explains why forecasts are less accurate in the Profits treatment than in the FSE treatment.

The right panel of Figure 18 presents the CDF of the percentage deviations of individual employment outcomes from the equilibrium employment corresponding to the announcements in both treatments (Figure 1 explains why the lowest value of employment is 6.25 and the highest value is 16). The CDF for the FSE treatment is larger than the CDF for the Profits treatment, which is statistically significant using Kolmogorov-Smirnov test (with p-value

of 0, and test statistic of 0.0839). This implies that employment outcomes are closer to the equilibrium values in the FSE treatment than in the Profits treatment. Forecast decisions are closer to the equilibrium values in the FSE treatment than in the Profits treatment. Because employment outcomes are based on forecasts, employment is also closer to the equilibrium employment in the FSE treatment than in the Profits treatment.

6 Further discussion of adaptive learning

The results of our experiments exhibit a large degree of consistency with the adaptive learning theory results described in Section 3.3, which is a theory that motivated our experimental study. There it was shown that an SSE fluctuating between n_H and n_L is locally stable under learning, as are the steady states n_H and n_L themselves. In our practice periods we ensured that subjects saw a strong correlation between the announcement A_t and reported average employment of others, \bar{N}_t . In many of the experimental sessions this was sufficient to generate convergence or approximate convergence to the SSE throughout the experimental session. In this section, we provide descriptive evidence about behavior of subjects suggestive of learning during the experiment.¹⁹

For example in the Profits treatment, sessions 1 to 3, we see initial deviations from the SSE in early periods, with a process of learning in which subjects eventually closely approximate the SSE. This is seen also in session 4 of Profits treatment, but with larger initial errors: the extended sequence of nine high announcements between periods 32 and 41, followed by five low announcements between periods 42 and 46, appear to have been helpful in inducing apparent eventual convergence to the SSE.

Even the cases in which there were substantial deviations from the announcement-based SSE are illuminating in terms of adaptive learning. In session 5 of the Profits treatment, subjects appear less certain about the relevance of the announcement. During the sequence of high announcements between periods 14 – 17 and 23 – 25, forecasts are significantly below $n_M = 12.5$, which implies that actual observations of average employment are less than the average forecast, which under adaptive learning pushes agents towards n_L rather than n_H . However, during the extended sequence of high announce-

¹⁹This descriptive evidence can be useful as a motivation for a formal testing of learning mechanisms by subjects which is outside the scope of this study.

ments in periods 32-41, subjects relearn the high equilibrium and continue to make high forecasts during the subsequent low announcements. At the end of session 5 it appears possible that subjects have converged on the high steady state. These results might be consistent with some of the subjects conditioning their learning rule on the announcements, with other subjects disregarding the announcements and instead using a simple non-conditional adaptive learning rule. In session 6 we see a similar pattern, except that in the middle part of the high announcement periods 32-41, expectations are slightly lower and this means that the even lower observed \bar{N}_t pushes forecasts down toward the low steady state. At the end of session 6 it appears possible that subjects have converged on the low steady state.

Similar interpretation can be given to the FSE sessions. Evidence of adaptive learning is seen in several of them, particularly of the forecast of average employment during periods of high announcements early in the experimental session. Where there is apparent convergence to the SSE, the convergence is quite close in several of the FSE sessions. In session 4 of the FSE treatment, however, there appears clearly to be eventual convergence to the low steady state. This again might be consistent with a substantial proportion of the subjects using non-conditional adaptive learning rules.

The adaptive learning framework described and discussed in Section 3.3 can be extended in various ways. For example, one can allow for heterogeneous priors of subjects, i.e. allow for different subjects to have different initial expectations and degrees of subjective uncertainty about their forecasts. Furthermore, more general adaptive learning rules along the line of Evans, Honkapohja and Marimon (2001), allow for heterogeneity in gains, inertia and experimentation. This can greatly increase the variety of possible paths under adaptive learning, and lead to more subtle learning dynamics in which heterogeneous expectations can emerge. Our results appear to be consistent with such generalized adaptive learning rules.

6.1 The role of heterogeneity in learning

Continuing with this last point, we note that when agents have heterogeneous expectations, learning dynamics can depend on the dispersion of expectations as well as on the average forecast. For example, even if, when the announcement is high, most agents have expectations near the high steady state, if there are several agents that have sufficiently low expectations, this

can be enough to destabilize coordination on an SSE.²⁰ Furthermore, under the Profits treatment, it is possible that subjects understand that their loss function is not symmetric around a given equilibrium, and they may take this feature into account when making their forecasts.²¹

Thus, how quickly subjects learn to coordinate on each announcement may also be influenced by the fact that switching to high employment from low employment more quickly than other subjects is costlier in terms of lower profits than switching to low employment while other subjects still choose high employment. How profitable choosing high or low employment is depends on how many subjects choose high and low values.

With our parametrization, choosing high employment is relatively more profitable than choosing low employment when five out of six subjects choose high values (see Table 3). In contrast, if two or more subjects out of six choose low employment, they get higher profit than those who choose high employment. Thus, the coordination on the announcement about high employment is more demanding in terms of how many subjects need to coordinate (five out of six) to make coordination on high employment profitable. And the coordination on low employment is relatively simpler: it requires only two subjects following the announcement about low employment to make choosing low employment more profitable. The technical reason for this asymmetry is the functional form of the production function and the calibration of positive production externality. The square root production function was chosen as a simple production function with diminishing returns, and form and calibration of production externality and wage rate were chosen to be as simple as possible and calibrated to yield three well-spaced Pareto-ranked steady states. Well-spaced steady states were needed for easy identification of the correspondence of experimental results with the equilibria and for easy understanding by the subjects. Given these design constraints, the asymmetry of the basins of attraction was difficult to avoid. With heterogeneous forecasts, the "basin of attraction" is a complicated multidimensional object, and can be a subject for future research.

Let us take a closer look at how learning happens during an experimental

²⁰More formally, the basin of attraction of an SSE, in terms of initial expectations, depends on the dispersion of these expectations as well as on the mean.

²¹"Direct criterion" versions of the adaptive learning rules, described in Section 3.3, can be developed in which decisions or forecasts are adjusted in the direction of the decision or forecast that would have been most profitable in the preceding period. For an example, see Woodford (1990).

session and at the role of heterogeneity. For example, in session 4 of the Profits treatment (Figure 7) during periods 23-26 with high announcements, the subjects fail to coordinate on high employment as only two or three subjects choose high values. Next in periods 32-41, four, then five and eventually all six subjects choose high values. During the final sequence of high announcements in periods 47-50, all subjects choose high values after one period of high announcements. Similarly, subjects learn in periods with low announcements. During periods 12-13, three and then four subjects choose low values. In periods 18-21, three subjects choose low values in period 18, and then all the subjects choose low values. Next in periods 26-31, five subjects choose low values immediately, and then all subjects choose low values. Thus, we can see that as the session proceeds it takes fewer periods for subjects to coordinate on the announcements, i.e. the subjects learn during the experiment.

However, coordination on the announcements does not always happen. In session 6 (Figure 13) subjects coordinate on the low value by the end of the session. During periods 32-41 with high announcements, only two or three subjects choose high values. This makes choosing the low value more profitable, and eventually all subjects choose low values. When not enough subjects choose high values, lower profits drive them towards the low equilibrium.

Similar dynamics are present in the FSE treatment. Coordination on the high-employment equilibrium is more demanding because it requires five out of six subjects to choose high values for the system dynamics to be driven towards the high equilibrium. When less than five subjects choose high values, average employment of others is below their forecasts which under adaptive learning pushes them towards low equilibrium.

In session 2 of the FSE treatment (Figure 11) during periods 5-11 and 14-17 with high announcements, the subjects learn to forecast high values after two periods during which four and then five subjects choose high values. In periods 23-25, all subjects choose high values, and in period 25 all subjects learn the exact high-equilibrium value. In periods 32-41, everybody chooses high-equilibrium value after one period of high announcements. During the final sequence of high announcements, all subjects choose the high-equilibrium values. Similarly, the subjects learn to choose low-equilibrium values during periods with low announcements. During periods 12-13, all subjects choose low forecasts which are quite heterogeneous. In periods 18-22, four and then five subjects choose the low-equilibrium value of 6.25 while

one chooses 7. This behavior continues during the remaining periods with low announcements (26-31 and 42-46).

In session 4 of the FSE treatment (Figure 13) subjects ignore the announcements and coordinate on the low equilibrium. It is interesting that at the beginning of this session during periods 5-11 with high announcements, all subjects choose high values. However, their forecasts are heterogeneous (14.25, 14.35, 15, 17, 13.8, 15). In periods 14-17, only one subject tries the high value for two periods and then switches to the low value. Next in periods 23-25, all the subjects choose low, heterogeneous forecasts resulting in low employment, and in the subsequent periods with high announcements, the forecasts are equal to the low-equilibrium value or very close to it.

7 Conclusion

We have conducted experiments in a simple, stylized macroeconomic model with a production externality that generates multiple equilibria. The equilibria are payoff-ranked – the low-employment equilibrium has lower profit than the medium or high-employment equilibria do – which adds to the challenge for coordination and switching between them. We observe that subjects can indeed coordinate on extraneous announcements (a ‘sunspot’ equilibrium), with switching between low- and high-employment states, in treatments with two different payoff structures. When subjects payoffs are evaluated based on forecast squared error (FSE treatment), their forecasts and employment outcomes are closer to the equilibrium corresponding to the announcement than they are in the treatment based on the profits. This is explained by the functional form of the employment and a reward based on the accuracy of the forecasts.

In our set-up coordination on the sunspot equilibrium is Pareto ranked superior to coordination on the low equilibrium but inferior to coordination on the high equilibrium. It is striking that in our set-up we appear able to induce subjects to coordinate on sunspot equilibria in a high proportion of the sessions, and that this occurs even when there exists an equilibrium steady state that would provide higher payoffs to all agents. However, the stability of the sunspot equilibria under adaptive learning is local, and we also see experiments in which subjects appear to eventually coordinate on the low or high steady state. Our results raise a number of important questions. What would happen if the initial experience obtained in training session were

different? Will the results be robust to the form of the externality? How would agents react if there were a regime change in which the number of steady states were reduced to one? Can our results be extended to dynamic versions of the model in which agents need to forecast both the level of current employment and the average level of employment next period? We reserve these questions and other extensions to future research.

References

- [1] Adam, K. 2007. Experimental Evidence on the Persistence of Output and Inflation. *Economic Journal*. 117, 603-636.
- [2] Arifovic, J. 1995. Genetic Algorithms and Inflationary Economies. *Journal of Monetary Economics*. 36, 219-243.
- [3] Arifovic, J. 1996. The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies. *Journal of Political Economy*. 104, 510-541.
- [4] Arifovic, J., Duffy, J. 2018. Heterogeneous Agent Modeling: Experimental Evidence, in: Hommes, C., LeBaron, B. (Eds.), *Handbook of Computational Economics*, Volume 4, Amsterdam: North-Holland, pp. 491-540.
- [5] Azariadis, C., Guesnerie, R. 1986. Sunspots and Cycles. *Review of Economic Studies*. 53, 725-737.
- [6] Bao, T., C. Hommes, Duffy, J. 2013. Learning, Forecasting and Optimizing: an Experimental Study. *European Economic Review*. 61, 186-204.
- [7] Bao, Te, Hommes, C., Sonnemans, J., Tuinstra, J. 2012. Individual expectations, limited rationality and aggregate outcomes. *Journal of Economic Dynamics and Control*. 36, 1101-1120.
- [8] Benhabib, J. and R. Farmer. 1994. Indeterminacy and Increasing Returns. *Journal of Economic Theory*. 63, 19-41.
- [9] Benhabib, J., Schmidt-Grohe, S., Uribe, M. 2001. The perils of Taylor-rules. *Journal of Economic Theory*. 96, 40-69.

- [10] Beugnot, J., Gurguc, Z., Ovlisen, F. Roose, and M. Roos. 2012. Coordination Failure Caused by Sunspots. *Economic Bulletin*. 32 (4), 2860-2869.
- [11] Capra, M., Tanaka, T., Camerer, C., Feiler, L., Sovero, V., Noursair, C.N. 2009. The Impact of Simple Institutions in Experimental Economies with Poverty Traps. *Economic Journal*. 119(539), 977-1009
- [12] Cass, D., Shell, K. 1983. Do Sunspots Matter? *Journal of Political Economy*. 91, 193-227.
- [13] Clarida, R., Gali, J., Gertler, M. 2000. Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory. *Quarterly Journal of Economics*. 115, 147-180.
- [14] Coibion, O., Gorodnichenko, Y., Kumar, S. 2018. "How do firms form their expectations? New survey evidence. *American Economic Review*. 108, 2671-2713.
- [15] Cooper, R., John, A. 1988. Coordinating Coordination Failures in Keynesian Models. *Quarterly Journal of Economics*. 113, 441-464.
- [16] Cooper, R., DeJong, D., Forsythe, R., Ross, T. 1990. Selection Criteria in Coordination Games: Some Experimental Results. *American Economic Review*. 80 (1), 218-233.
- [17] Cooper, R. W., DeJong, D. V., Forsythe, R., Ross, T.W. 1992. Communication in Coordination Games. *Quarterly Journal of Economics*. 107, 739-771.
- [18] Corbae, D., Duffy, J.. 2008. Experiments with Network Formation. *Games and Economic Behavior*. 64, 81-120.
- [19] Duffy, J. 2016. Macroeconomics: A Survey of Laboratory Research, in: J.H. Kagel and A.E. Roth (Eds.), *Handbook of Experimental Economics*. Volume 2, Princeton University Press, Princeton, pp. 1-90.
- [20] Duffy, J., Feltovich, N. 2010. Correlated Equilibria, Good and Bad: An Experimental Study. *International Economic Review*. 51, 701-721.
- [21] Duffy, J., Fisher, E. 2005. Sunspots in the Laboratory. *American Economic Review*. 95, 510-529.

- [22] Ehrmann, M., Pfajfar, D., Santoro, E. 2017. Consumer attitudes and their inflation expectations. *Journal of Central Banking*. 13, 225-259.
- [23] Evans, G., Honkapohja, S. 1995. Increasing Social Return, Learning and Bifurcation Phenomena, in: Kirman, A., Salmon, M. (Eds.), *Learning and Rationality in Economics*, Blackwell, Oxford, pp. 216-235.
- [24] Evans, G., Honkapohja, S. 1994. On the Local Stability of Sunspot Equilibria under Adaptive Learning. *Journal of Economic Theory*. 64, 142-161.
- [25] Evans, G., Honkapohja, S.. 2001. *Learning and Expectations in Macroeconomics*, Princeton: Princeton University Press.
- [26] Evans, G., Honkapohja, S., Marimon, R. 2001. Convergence in Monetary Inflation Models with Heterogeneous Learning Rules. *Macroeconomic Dynamics*. 5, 1-31.
- [27] Evans, G., Honkapohja, S., Marimon, R. 2007. Stable Sunspot Equilibria in a Cash-in-Advance Economy. *The B.E. Journal of Macroeconomics, Advances*. 7 (1), Article 3.
- [28] Evans, G., Honkapohja, S., Romer, P. 1998. Growth Cycles. *American Economic Review* 88, 495-515.
- [29] Farmer, Roger. 1999. *The Macroeconomics of Self-Fulfilling Prophecies*, second ed. MIT Press, Cambridge.
- [30] Farmer, 2015. Global sunspots and asset prices in a monetary economy. NBER working paper No 20831.
- [31] Fehr, D., Heinemann, F., Llorente-Saguer, A. 2018. The Power of Sunspots: An Experimental Analysis. *Journal of Monetary Economics*.
- [32] Fischbacher, U. 2007. Z-Tree: Zurich Toolbox for Ready-made Economic Experiments. *Experimental Economics*. 10 (2), 171-178.
- [33] Gali, J., (2014). Monetary policy and rational asset price bubbles. *American Economic Review*. 104, 721-752.

- [34] Garratt, R., Keister, T. 2009. Bank runs as coordination failures: an experimental study. *Journal of Economic Behavior and Organization*. 71, 300-317.
- [35] Honkapohja, S., Mitra, K. 2004. Are non-fundamental equilibria learnable in models of monetary policy. *Journal of Monetary Economics*. 51, 1743-1770.
- [36] Heinemann, F., Nagel, R., Ockenfels, P. 2004. The Theory of Global Games on Test: Experimental Analysis of Coordination Games with Public and Private Information. *Econometrica*. 72, 1583-1599.
- [37] Howitt, P., McAfee, R.P. 1992. Animal Spirits. *American Economic Review*. 82, 493-507.
- [38] Lei, V., Noussair, C.N. 2002. An Experimental Test of an Optimal Growth Model. *American Economic Review*. 92(3), 549-570.
- [39] Lei, V., Noussair, C.N. 2007. Equilibrium Selection in an Experimental Macroeconomy. *Southern Economic Journal*. 74, 448-482.
- [40] Lim, S.S., E. Prescott, E., S. Sunder, S. 1994. Stationary Solution to the Overlapping Generations Model of Fiat Money: Experimental Evidence. *Empirical Economics*. 19, 255-77.
- [41] Marimon, R., Sunder, S. 1993. Indeterminacy of Equilibria in a Hyperinflationary World: Experimental Evidence. *Econometrica*. 61, 1073-1107.
- [42] Marimon, R., Sunder, S. 1994. Expectations and Learning under Alternative Monetary Regimes: An Experimental Approach. *Economic Theory*. 4, 131-62.
- [43] Marimon, R., Sunder, S. 1995. Does a Constant Money Growth Rule Help Stabilize Inflation? *Carnegie-Rochester Conference Series on Public Policy*. 43, 111-156.
- [44] Marimon, R., Spear, S.E., Sunder, S. 1993. Expectationally Driven Market Volatility: An Experimental Study. *Journal of Economic Theory*. 61, 74-103.

- [45] McGough, B., Q. Meng, Q., Xue, J. 2013. Expectational stability of sunspot equilibria in non-convex economies. *Journal of Economic Dynamics and Control*. 37, 1126-1141.
- [46] Mertens, M., Ravn, M.O. 2014. Fiscal Policy in an Expectations-Driven Liquidity Trap. *Review of Economic Studies*. 81, 1637–1667
- [47] Miao, J., Shen, Z., Wang,P. 2019. Monetary Policy and Rational Asset Bubbles: Comments. *American Economic Review*. 109(5), 1969-1990.
- [48] Shea, P., 2013. Learning By Doing, Short-Sightedness, and Indeterminacy. *Economic Journal*. 123, 738-763.
- [49] Schotter, A., Yorulmazer, T. 2009. On the Severity of Bank Runs: An Experimental Study. *Journal of Financial Intermediation*. 18, 217-241.
- [50] Van Huyck, J. B., Battalio, R. C., Beil, R.O. 1990. Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure. *American Economic Review*. 80(1), 234-48.
- [51] Woodford, M. 1990. Learning to Believe in Sunspots. *Econometrica*. 58, 277-307.

Table 1: Percentage of observations in each range of values.

range	FSE treatment				Profits treatment			
	E, H	E, L	F, H	F, L	E, H	E, L	F, H	F, L
5.5-6.5	14.71	98.23	10.19	70.20	17.28	88.76	2.78	33.46
15.5-16.5	83.54	1.64	61.83	1.14	76.03	8.84	28.60	1.77

Note: E, H = employment during periods with high-employment announcements; E, L = employment during low announcements; F, H = forecasts during high announcements; F, L = forecasts during low announcements.

Table 2: Descriptive statistics of the data.

	High-employment announcements				Low-employment announcements			
	average	median	std	skewness	average	median	std	skewness
Employment								
FSE	14.46	16.00	3.51	-1.86	6.42	6.25	1.27	7.34
Profits	13.91	16.00	3.83	-1.36	7.20	6.25	2.81	2.73
Forecasts								
FSE	14.35	16.00	3.21	-1.84	6.66	6.25	1.37	4.78
Profits	14.21	15.00	3.30	-1.06	7.85	7.00	3.14	1.51

Table 3: Productivity and profits of subjects forecasting low and high values.

Number of forecasters		Productivity of forecasters		Profit of forecasters	
low	high	low	high	low	high
6	0	2.5	-	3.125	-
5	1	2.5	2.5	3.125	2
4	2	2.5	2.5	3.125	2
3	3	3.1	2.5	4.65	2
2	4	4.0	3.1	6.87	4.4
1	5	4.0	4.0	6.87	8
0	6	-	4.0	-	8

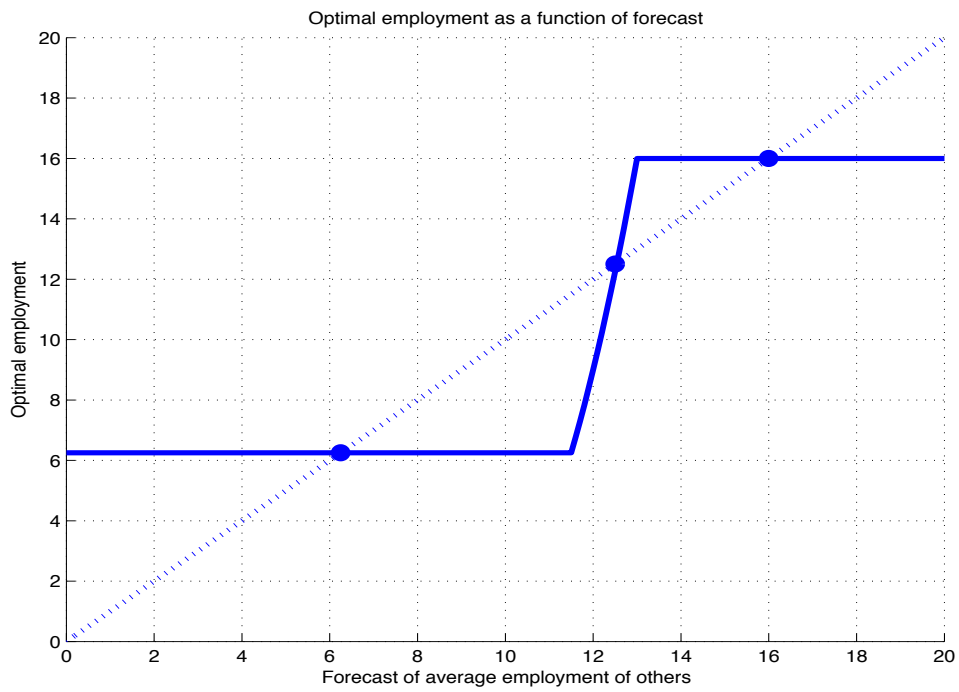


Figure 1: This figure presents employment as a function of forecast of average employment of others according to equation 1.

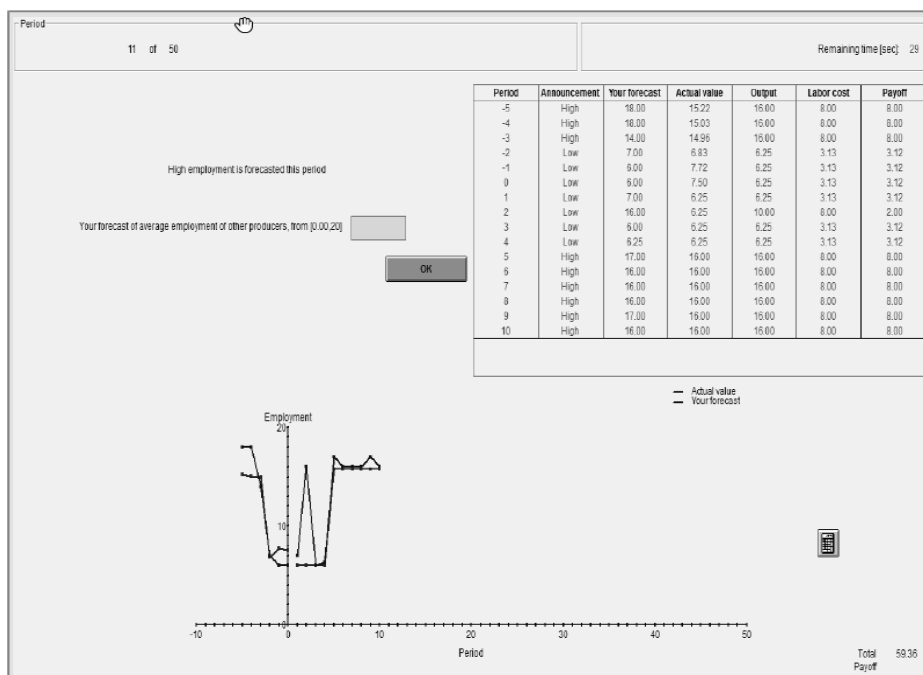


Figure 2: Screenshot for Profit treatment.

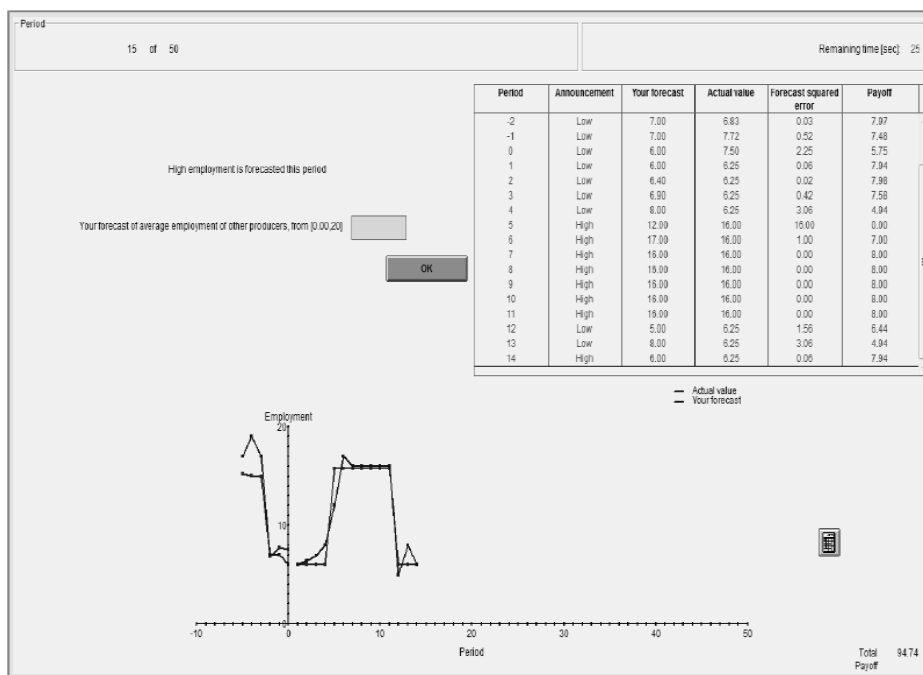


Figure 3: Screenshot for FSE treatment.

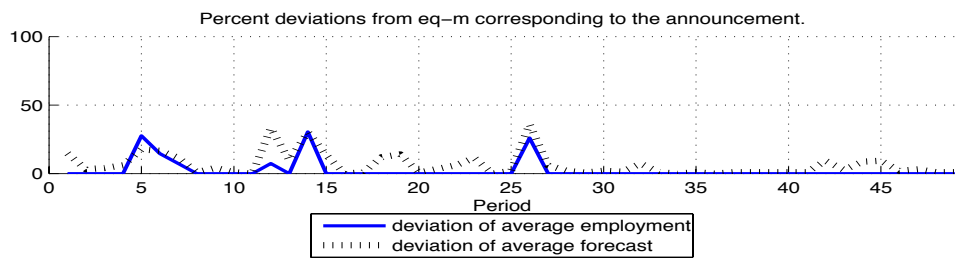
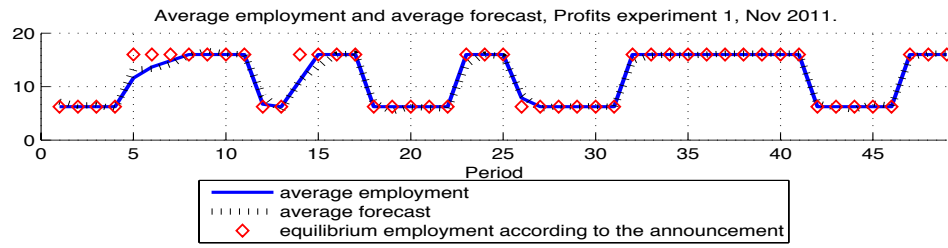


Figure 4: Session 1 of profits treatment.

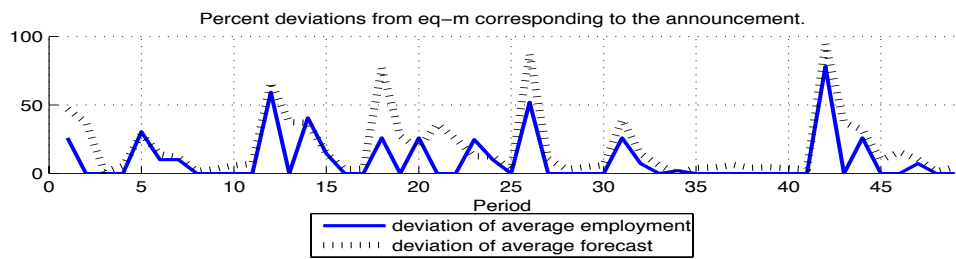
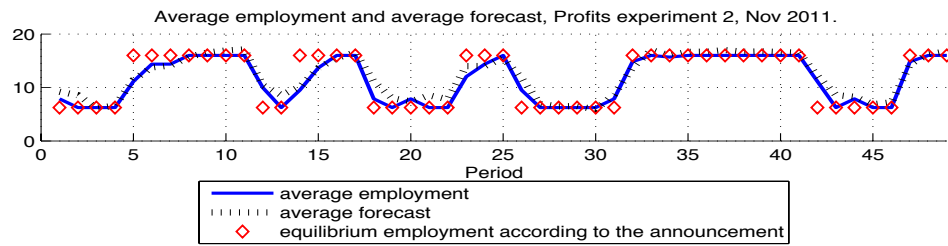


Figure 5: Session 2 of Profits treatment.

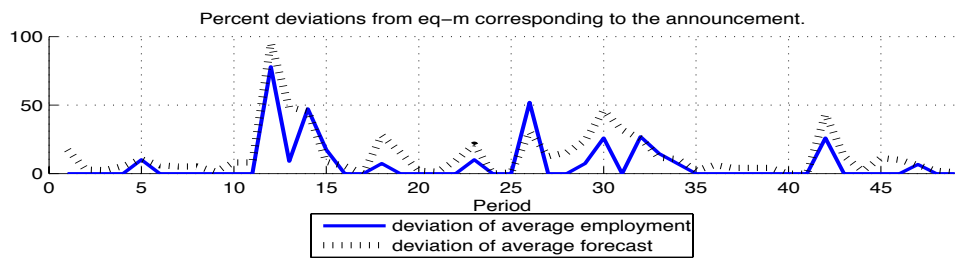
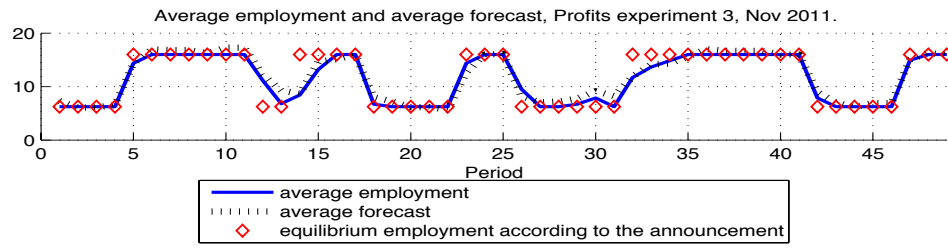


Figure 6: Session 3 of Profits treatment.

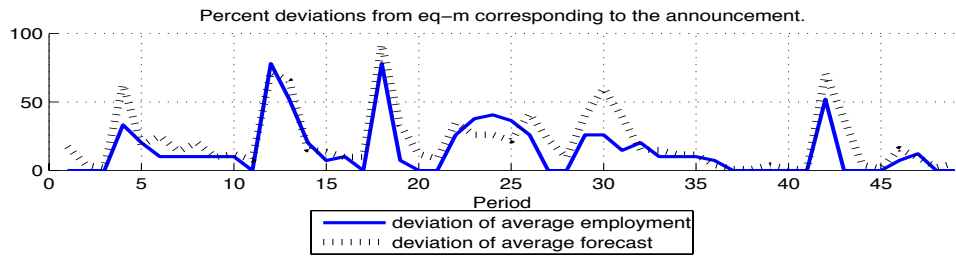
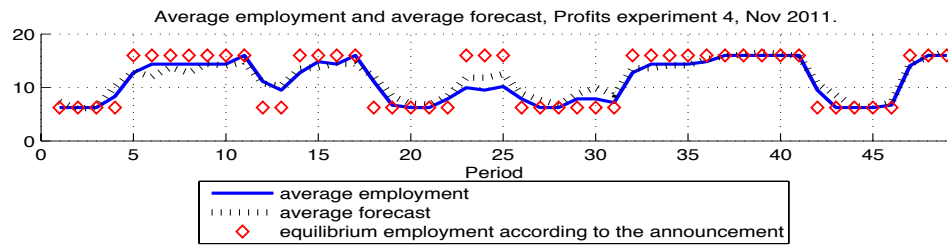


Figure 7: Session 4 of Profits treatment.

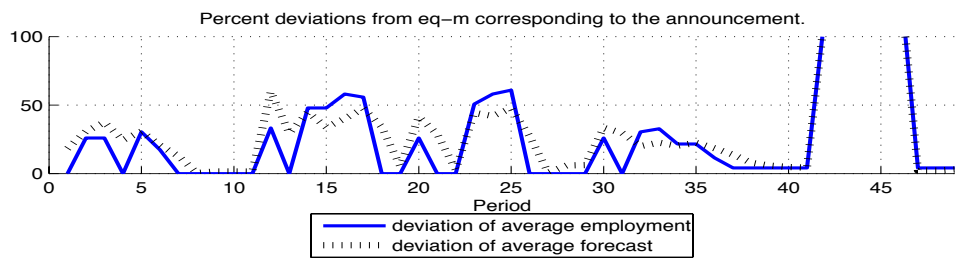
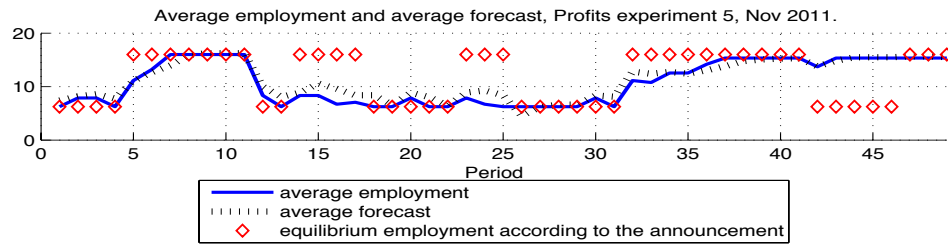


Figure 8: Session 5 of Profits treatment.

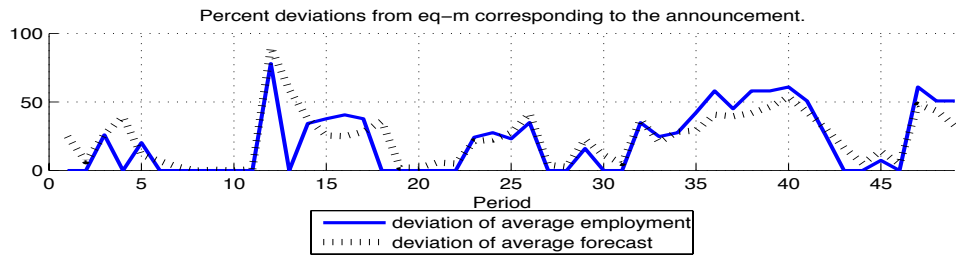
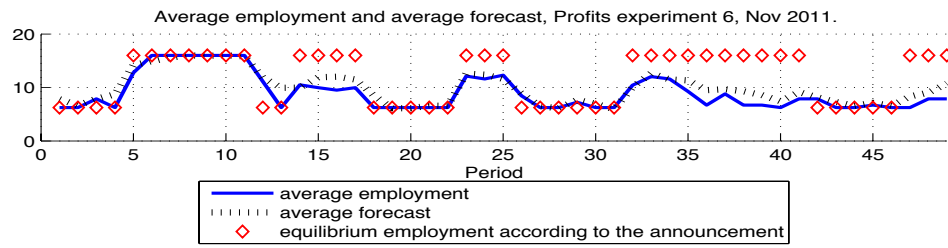


Figure 9: Session 6 of Profits treatment.

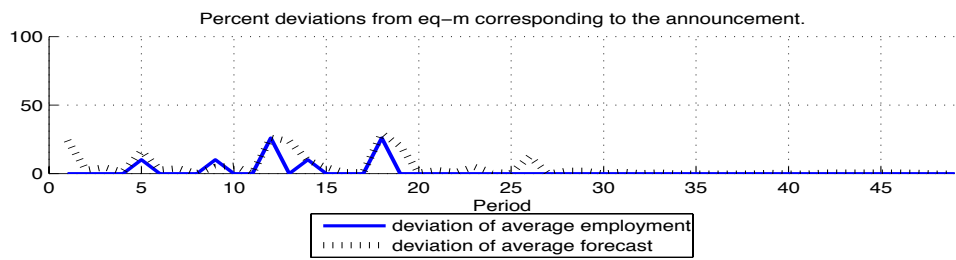
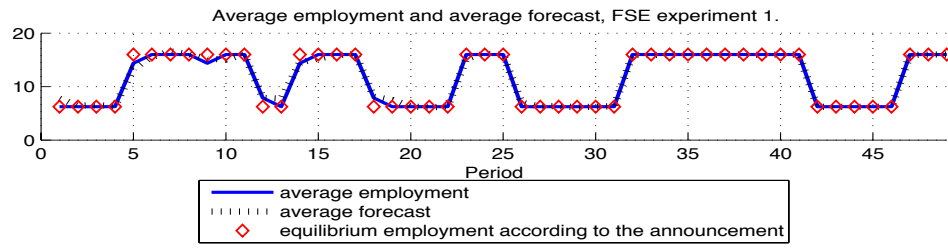


Figure 10: Session 1 of FSE treatment.

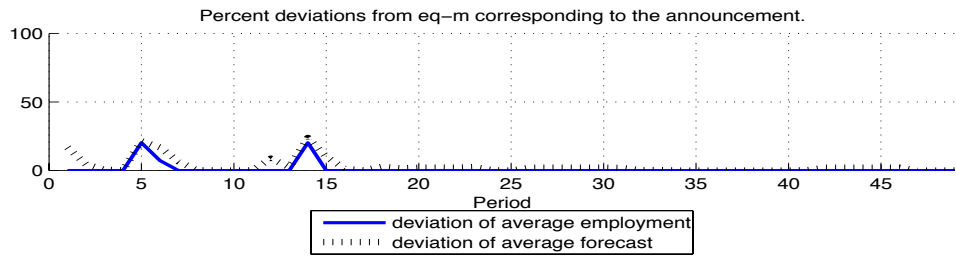
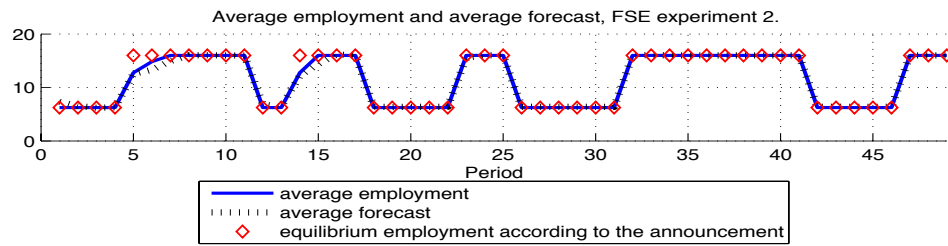


Figure 11: Session 2 of FSE treatment.

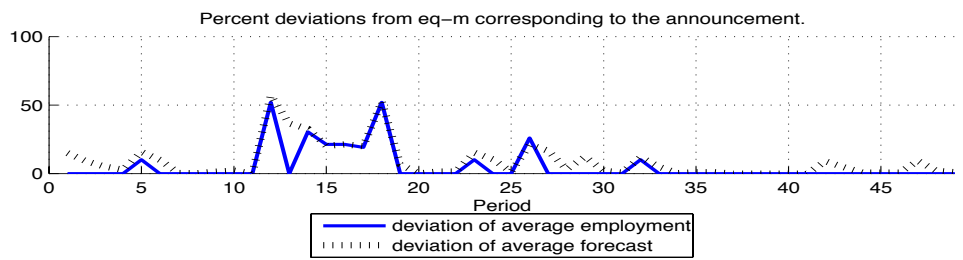
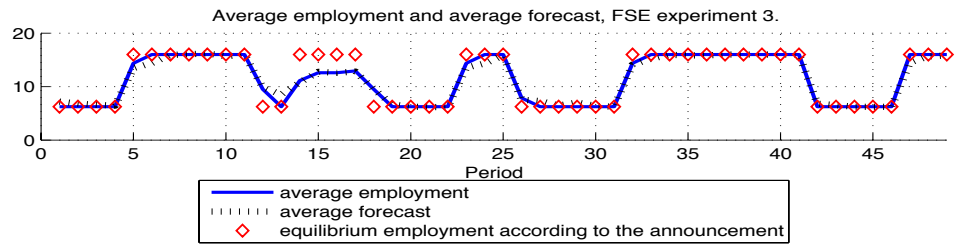


Figure 12: Session 3 of FSE treatment.

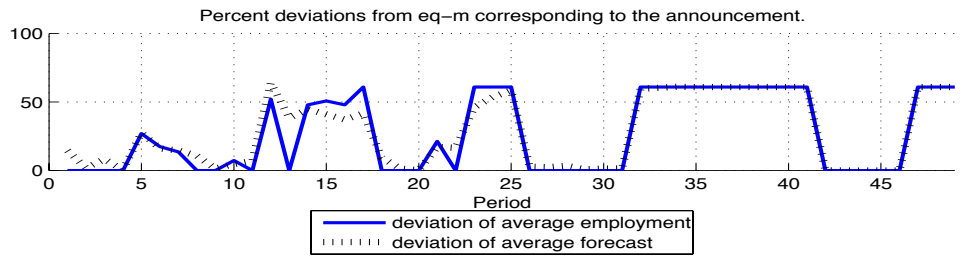
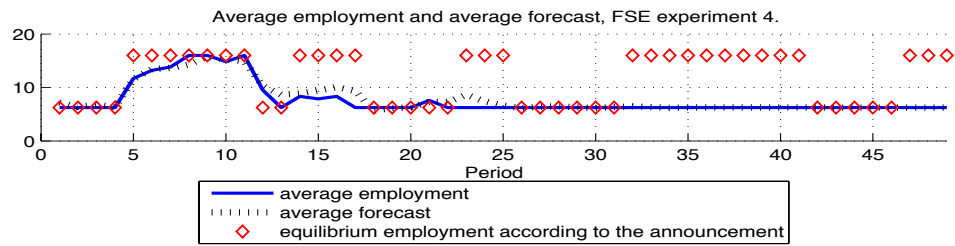


Figure 13: Session 4 of FSE treatment.

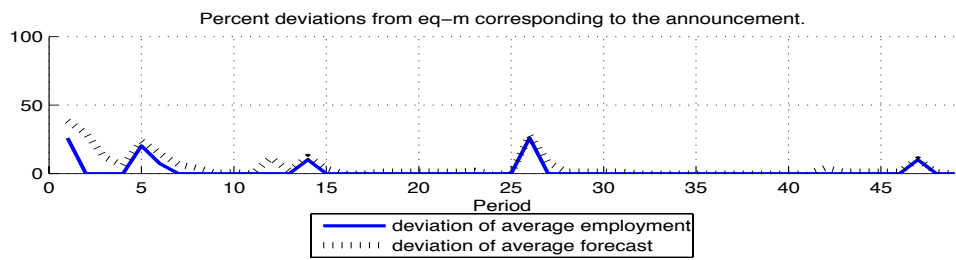
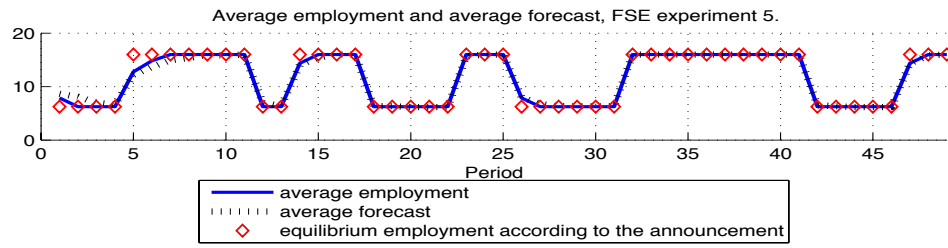


Figure 14: Session 5 of FSE treatment.

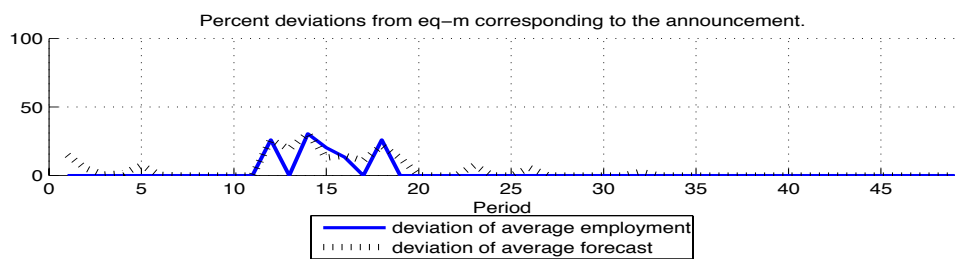
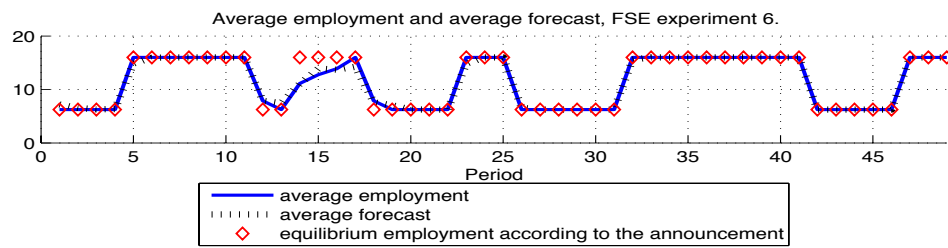


Figure 15: Session 6 of FSE treatment.

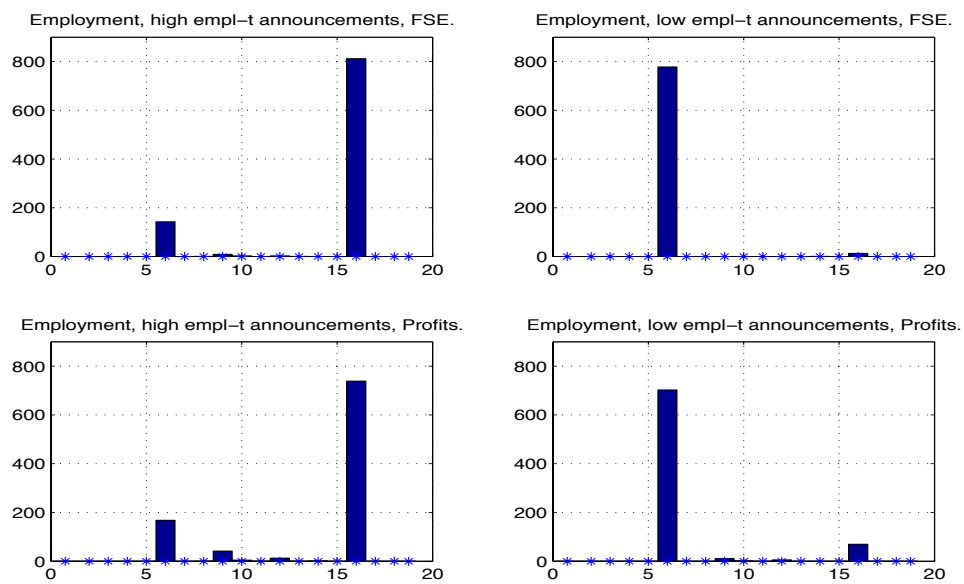


Figure 16: Histograms of employment during periods with high- and low-employment announcements in FSE and Profits treatments.

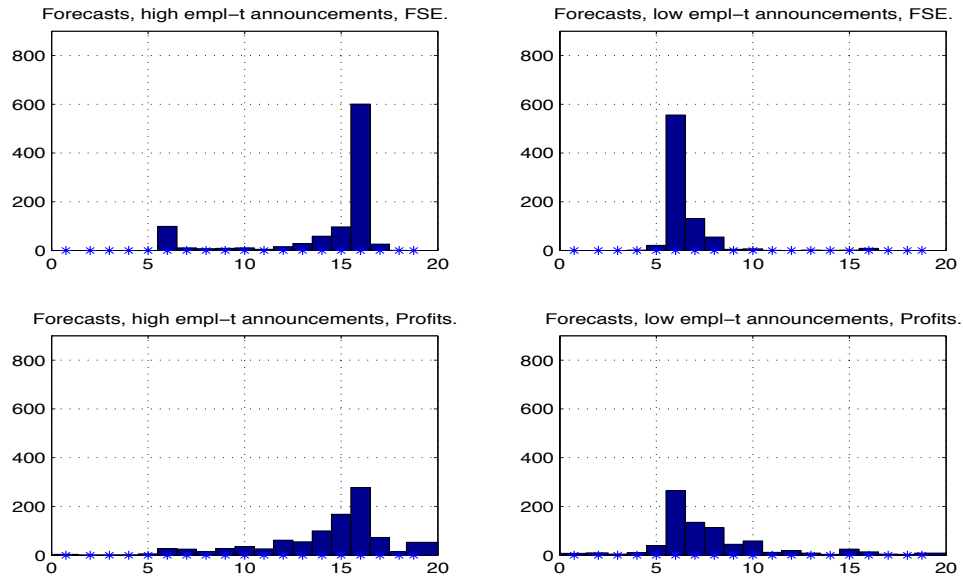


Figure 17: Histograms of forecasts during periods with high- and low-employment announcements in FSE and profits treatments.

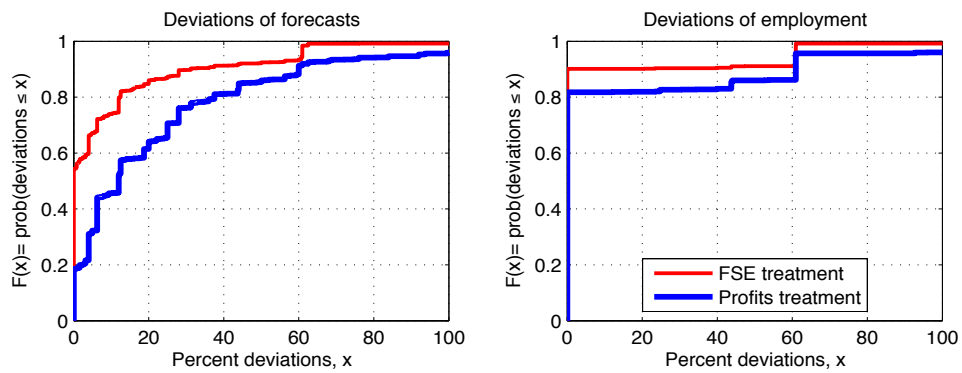


Figure 18: Cumulative density functions of deviations of forecasts and employment from equilibrium employment corresponding to the announcements.