



BANK OF CANADA
BANQUE DU CANADA

Working Paper/Document de travail
2013-14

Are Sunspots Learnable? An Experimental Investigation in a Simple General-Equilibrium Model

by Jasmina Arifovic, George Evans and Olena Kostyshyna

Bank of Canada Working Paper 2013-14

May 2013

Are Sunspots Learnable? An Experimental Investigation in a Simple General-Equilibrium Model

by

Jasmina Arifovic,¹ George Evans² and Olena Kostyshyna³

¹Simon Fraser University

²University of Oregon and University of St. Andrews

³Canadian Economic Analysis Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
kost@bankofcanada.ca

Bank of Canada working papers are theoretical or empirical works-in-progress on subjects in economics and finance. The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the Bank of Canada.

Acknowledgements

We would like to thank Heng Sok and Brian Merlob for helpful research assistance. We would also like to thank Luba Petersen and Gabriele Camera for useful feedback, and participants at the Experimental Macroeconomic Conference, Pompeu Fabra, May 2011, as well as seminar participants at the Bank of Canada, May 2012, University of California, Irvine, September, 2012, and ESA Meetings, Tucson, AZ, November 2012 for helpful comments.

Abstract

We conduct experiments with human subjects in a model with a positive production externality in which productivity is a non-decreasing function of the average level of employment of other firms. The model has three steady states: the low and high steady states are expectationally stable (E-stable), and thus locally stable under learning, while the middle steady state is not E-stable. There also exists a locally E-stable 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 in our experimental economies can learn a sunspot equilibrium. Our experimental design has two treatments: one in which payoff is based on the firm's profits, and the other in which payoff is based on the forecast squared error. We observe coordination on the extrinsic announcements in both treatments. In the treatments with forecast squared error, the average employment and average forecasts of subjects are closer to the equilibrium corresponding to the announcement. Cases of apparent convergence to the low and high steady states are also observed.

JEL classification: D83, G20

Bank classification: Economic models

Résumé

Dans le cadre d'expériences effectuées à l'aide de sujets humains, les auteurs étudient un modèle comportant une externalité de production positive, dans lequel la productivité d'une entreprise est une fonction non décroissante du niveau moyen de l'emploi dans les autres firmes. Le modèle distingue trois états stationnaires : deux (bas et haut) qui sont stables sous anticipations (A-stables) – et par conséquent localement stables sous apprentissage – et un état stationnaire intermédiaire qui n'est pas A-stable. Il existe également un équilibre à « taches solaires » présentant une A-stabilité locale, lequel fluctue entre les deux états stationnaires limites (bas et haut). Les états stationnaires sont ordonnés en fonction du gain obtenu, les profits étant faibles ou élevés selon que le niveau d'emploi est bas ou haut.

Les auteurs examinent si les sujets participants à l'expérience peuvent apprendre d'un équilibre à « taches solaires ». Deux situations sont envisagées : dans l'une, les gains dépendent des profits de l'entreprise; dans l'autre, de l'erreur quadratique de prévision. Les auteurs constatent que dans les deux cas, les sujets coordonnent leurs anticipations sur les annonces extrinsèques. Dans la situation où les gains sont liés à l'erreur quadratique de prévision, le niveau moyen de l'emploi et la moyenne des prévisions des sujets sont plus proches de l'équilibre correspondant à l'annonce. Les auteurs observent aussi des cas de convergence apparente vers les deux états stationnaires limites.

Classification JEL : D83, G20

Classification de la Banque : Modèles économiques

1 Introduction

In this paper, we describe 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. This is a first study of a coordination on sunspot equilibrium that involves switching between payoff-rankable equilibria in a macroeconomic environment. The switching between high and low equilibria can be interpreted as a change of outlook between optimism and pessimism leading to fluctuations in the experimental economy.

Experimental 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), simultaneous games (Cooper et al. 1992), effort coordination games (van Huyck et al. 1960), optimal growth models with non-convex production technology (Lei and Noussair 2007), and bank runs (Garratt and Keister 2005; Schotter and Yorulmazer 2003; Corbae and Duffy 2008). Ochs (1995) and Duffy (2008) provide surveys of this literature.

Models with multiple equilibria typically also have solutions called stationary sunspot equilibria (SSEs), in which agents' actions are conditioned on an extraneous random variable (Cass and Shell 1983). 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), and Evans et al. (2007), as well as in numerous other papers.

This paper examines whether SSEs can be observed in laboratory experiments in a simple macroeconomic model with positive production externalities generating multiple steady states that support sunspot equilibria. A central feature of our framework is that the steady states are Pareto ranked – high-employment steady states are superior to low-employment steady states – and that SSEs that fluctuate between a pair of steady states are inferior to the higher of the two steady states.

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.¹ In their set-up, the model has a unique steady-state equilibrium and a two-period cyclic 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 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. Thus there appears to be coordination on a cyclical equilibrium, though it is difficult to tell, given the length of the experiments, how long this cyclic behavior would

¹See, for example, Azariadis and Guesnerie (1986).

continue. The cyclic behavior tends to trail off toward 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 both the closed-book call market and double auction implementations, two mechanisms that have different information flows. In their set-up, the marginal valuations of buyers and the 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 ranked: 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 (the wording of the announcements) and institutions: sunspot equilibria are observed in all sessions with call markets, but in less than half of sessions with double auction markets.

Fehr et al. (2011) 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. The global games literature has shown that the existence of multiple equilibria, in coordination-type games, is sensitive to private and public signals (for example, Heinemann et al. 2004). These signals can in effect operate as sunspots. Motivated by the global games results, Fehr et al. (2011) use public and/or private signals of different precision and study experimentally how they determine on which equilibrium subjects coordinate.

The focus of our paper is different, and motivated instead by the macro literature. We look at a simple macro set-up with positive production externalities and multiple steady states. We do not have any private signals and the public signals have zero information content. In this setting there exist SSEs, and only 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 (Hommes et al. 2005; Hommes et al. 2008; Heemeijer et al. 2009; and Adam 2007). 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. Hommes et al. (2008) and Heemeijer et al. (2009) find that price bubbles are possible when feedback from expectations to the realized variables is positive. Adam (2007) finds evidence of a restricted perceptions equilibrium. While these experiments are clearly related to our work, they differ from the current study in that we focus on the question of whether agents can reach a sunspot equilibrium.

Within this general literature, our work is closest to Marimon et al. (1993) and Duffy and Fisher (2005), both of whom study experimentally whether the equilibrium can be driven by extraneous public announcements. Like Marimon, Spear and Sunder (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 an indeterminate steady state, in particular at SSEs near a two-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 ranked, driven by an aggregate production externality. We do not require heterogeneity of agents, and thus the SSEs we examine can be interpreted as switches between high and low levels of aggregate output, resulting from waves

of optimism or pessimism driven by extraneous public announcements.

Our experimental environment is characterized by a positive production externality in which each firm’s productivity is a non-decreasing function of the average employment of other firms. In the theoretical model, each firm chooses employment to maximize profit, and its decision about employment depends on its productivity, while its productivity depends on the average employment of other firms. The decision about employment must be made before employment decisions of other firms are known, and so it depends on the firm’s forecast of the average employment of other firms. Therefore, in our experiments the subjects make forecasts of the average employment of other firms, and the employment of their firms is determined optimally based on this forecast.

The model has three steady states. In the language of the macro learning literature, the low-employment and high-employment steady states are expectationally stable (E-stable) (and thus are stable under adaptive learning rules), while the middle steady state is not E-stable. There also exists an E-stable sunspot equilibrium that 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 poses an additional challenge for coordination on SSEs, 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 profits, which poses 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. The experiments show frequent examples of coordination 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. The experiments also show examples of coordination on both low and high steady states.

The paper is organized as follows. In section 2, we describe the model. In section 3, we describe the design of the experiments. Section 4 reports the results of the experiments, and section 5 discusses adaptive learning in the experiments. Section 6 concludes the paper.

2 Model

2.1 Description of the economy

We use a 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.²

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

$$y_t = \psi_t \sqrt{n_t}, \tag{1}$$

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

$$\Pi_t = \psi_t \sqrt{n_t} - wn_t. \tag{2}$$

²Our set-up is closely related to the “Increasing Social Returns” overlapping-generations model described in Evans and Honkapohja (1995, 72–81). To keep the framework 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.

The level of productivity ψ_t depends on the average level of employment *across all other firms* (not including a firm’s own employment).³ We will call average employment of other firms \overline{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, \overline{N}_t , when its decision is made. A firm is more productive when other firms are operating at a high level of employment. Specifically, productivity, ψ_t , depends on the average employment of other firms, \overline{N}_t , as follows⁴:

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

This model can be thought of as a very simple and stylized general-equilibrium model, with a single consumption good, no capital or other means of saving, and in which household utility is such that labor supply is infinitely elastic at wage w .⁵ Firms are owned by households and profits are distributed as dividends back to the households each period. Note that 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.

2.2 Equilibria

For this economy, profits are maximized when firms choose:

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

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 three steady states: $n_L = 6.25$ (“low-level”), $n_M = 12.5$ and $n_H = 16$ (“high-level”).

When there are three steady states, SSEs exist between any pair of steady states. For example, in the experiments, we randomly generate announcements of “high” and “low” forecast employment. Letting $A_t \in \{L, H\}$ denote the announcement at time t , where L represents the announcement “Low employment is forecast this period” and H denotes the announcement “High employment is forecast 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.⁶

³Another 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 (under competitive assumptions), but they can affect the behavior of the experimental economy, as will be discussed later.

⁴Lei 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, productivity is low. They find that experimental economies can get into poverty traps with low levels of capital and output.

⁵It would be straightforward to generalize the model to allow for less than fully elastic labor supply.

⁶Equilibria or SSEs can be constructed that depend on any observable; e.g., equilibria can be constructed that switch between steady-state 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 $A_t = L$.

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 subsidies or taxes are assumed to be 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 next determine which steady states are stable under simple adaptive learning rules that have been widely studied in the macro learning literature.

2.3 Temporary equilibrium framework and E-stability

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 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: $n_t^i = \left(\frac{\psi_t^{e,i}}{2w}\right)^2$ where the superscripts e, i denote the expectations of agent i .

Average employment of other firms for firm i is given by

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

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, \overline{N}_t^i and Y_t . The profit of firm i at time t is given by

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

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 \overline{N}_t (equation (3)), it is equivalent to specify expectations in terms of $\overline{N}_t^{e,i}$. That is, given the profile of time t expectations $\{\overline{N}_t^{e,i}\}_{i=1}^K$, employment levels are given by

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

Figure 1 shows a firm’s optimal choice of employment \overline{n} as a function of the firm’s forecast, as given by equation (7). The above equations then determine \overline{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.⁷ For convenience (this is not essential), assume

⁷For details and further discussion of adaptive learning, see Evans and Honkapohja (2001).

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, which they use to determine their employment in period t , according to the adaptive rule

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

where γ_t are the “gain” parameters, which might, for example, be fixed at a number γ such that $0 < \gamma_t = \gamma < 1$.⁸ Then 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.

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 next examine 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 here assume 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 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. Sunspots fluctuating 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 $t \rightarrow \infty$,

⁸We need to assume that $\sum_{t=1}^{\infty} \gamma_t = +\infty$.

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.⁹

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 reporting the results below.

3 Design of experiments

As explained 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 $n = 12.25$. 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 higher than its profit in the low steady state, as shown in Table 1. Therefore, it might seem likely that subjects would coordinate on the high steady state. Payoff domination by the high steady state makes coordination on 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.

The payoff dominance of one of the steady states in our model is the key difference between our experiments and those in Duffy and Fisher (2005). In Duffy and Fisher (2005), there are two possible equilibria, depending on the median price in the market. The marginal valuations of the buyers and the marginal costs of the sellers are high (low) if the median price is above (below) the threshold. However, some of the buyers and some of the sellers are better off in the “high” equilibrium, and some are better off in the “low” equilibrium.

Information about the economy We provide descriptive information about the economy without technical details and equations. For example, the instructions provide the following information. “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 , the 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

⁹For additional discussion and details of adaptive learning of SSEs, see Evans and Honkapohja (1994; 2001, chapters 4.6 and 12).

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). The experiment lasts 50 periods, and the subjects are told this.

Payoffs We conduct two different 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 the forecast squared error:

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

and forecasting payoff is computed as

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

where 8 is the maximum payoff when the 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 forecast in this period” or “High employment is forecast 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 forecast this period” or “High employment is forecast this period.” The announcements are generated randomly. 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 generated randomly by the experimenters before the experiment.

Practice periods Each experiment includes six practice periods, during which subjects can familiarize themselves with the environment. We also use practice periods to “train” (condition) subjects to experience different equilibria 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 three periods of high employment and then three 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 concerns low employment.

Information on the computer screen On the computer screen, subjects can see past data in the table and graph updated with the actual results from the experimental economy as it evolves. In each period, the subjects can see the announcements about the economy. Figure 2 shows a screenshot of the computer screen during the experiment.

The experimental software was programmed using z-Tree (Fischbacher 2007). The experiments were conducted in November 2011 at the Economic Science Institute, Chapman University. We ran two treatments: one with payoff based on the firm’s profits (profits treatment), and the other with payoff based on the forecast squared error (FSE treatment). We ran six sessions of each treatment, with six subjects participating in each session (total of 72 subjects). Each session lasted 50 periods.

4 Results of experiments

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 report the results observed in individual sessions for each treatment, then analyze the data in terms of deviations from the announcements and in terms of efficiency. We also compare the results observed in the two treatments.

4.1 Profits treatment

Figures 3–8 report the results of the profits treatment for each individual session.¹⁰ Each figure consists of two panels. The first panel shows average employment, average forecast and the equilibrium employment corresponding to the announcement.¹¹ (Note that participants were not given the history of equilibrium employment corresponding to the announcement on their screens. We provide these series in our figures for ease of comparison with the actual data.) The second panel shows 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 3 where average employment is the same as the equilibrium corresponding to announcements in most of the periods except for five instances. Figure 3 shows that the percentage deviations from the announced equilibrium are zero for all periods, except for five 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 shown in Figures 4–6, respectively, 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 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, shown in Figure 7, between periods 12 and 31 there are both high and low announcement stretches in which average employment does not correspond to the announcement, while after period 32 average

¹⁰We 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 4.3.

¹¹By the “equilibrium employment corresponding to the announcement” we mean n_H if $A_t = H$ and n_L if $A_t = L$.

employment is 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, shown in Figure 8, 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 is 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. As we will discuss in section 5.1, it is harder to switch from the low to the high steady state than vice versa. This difficulty could explain what otherwise appears as a puzzle.

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.

4.2 FSE treatment

Figures 9–14 show the results of the FSE treatment. Again, the data for each session are shown in a figure that consists of two panels. The first reports average employment, average forecast and equilibrium employment corresponding to the announcement; the second reports percentage deviations of average employment and average forecasts from equilibria corresponding to the announcement.

In sessions shown in Figures 9–11 and 13 and 14, 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 (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 reported in section 4.3.

Figure 12 illustrates the results of session 4 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.

4.3 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.

4.3.1 Employment and forecasts

We collect the data on individual employment in periods with low-employment announcements and in periods with high-employment announcements separately, and pool these data for all experimental sessions for each treatment. Table 2 reports the percentages of observations in the ranges containing two equilibria; this table corresponds to the histograms shown in Figures 15 and 16. The top left panel of Figure 15 shows the histogram of individual employment decisions during periods with high-employment announcements in the FSE treatment: the values of employment are concentrated on the high-employment equilibrium of 16 (83.54 per cent of employment outcomes according to Table 2). The top right panel of Figure 15 shows the histogram of individual employment decisions during periods with low-employment announcements in the FSE treatment: the values of employment are heavily concentrated on the low-employment equilibrium of 6.25 (98.23 per cent of outcomes according to Table 2). The bottom left and right panels of Figure 15 show 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 per cent of the employment outcomes are close to the high-equilibrium employment of 16; and during periods with low-employment announcements, 88.76 per cent of the employment outcomes are close to the low-equilibrium employment of 6.25.

We also collect 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 16 show the histograms of individual forecasts made in periods with high- and low-employment announcements in the FSE treatment, respectively. The forecasts are heavily concentrated on the respective equilibrium values corresponding to the announcements: 61.83 per cent of forecasts are in the range containing the high-equilibrium employment of 16, and 70.20 per cent of forecasts correspond to the low-equilibrium employment of 6.25. The bottom left and right panels of Figure 16 show the histograms of individual forecasts made in periods with high- and low-employment announcements in the profits treatment, respectively. The forecasts are centered around the equilibrium values, but the percentages of forecasts in the ranges containing equilibria are much lower than in the FSE treatment: 28.60 per cent of the forecasts are in the range containing the high-equilibrium employment of 16, and 33.46 per cent 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 section 4.3.2.

Table 3 reports the data on average, median and standard deviations of employment and forecasts in both treatments.

Figures 17 and 18 show the empirical cumulative distribution functions (CDF) of employment outcomes and forecasts in both treatments and the theoretical CDF of equilibrium employment values – low-equilibrium and high-equilibrium – based on the number of periods with low-employment and high-employment announcements during the experiments.

4.3.2 Deviations from the equilibrium employment

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 top left panel of Figure 19 shows the 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 a 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 in 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 top right panel of Figure 19 shows 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 the Kolmogorov-Smirnov test (with a p-value of 0, and a 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.

4.3.3 Efficiency measures

Next, we evaluate how close the subjects' payoffs are to the equilibrium payoffs that would result if the subjects followed the announcements. To facilitate the comparison between the FSE and profits treatments, we compute the forecast squared error of forecasts made in the profits treatment, and we compute profits corresponding to the forecasts made in the FSE treatment. Then we evaluate two efficiency measures: the first measure is efficiency based on the profits, and the second measure is efficiency based on the forecast squared error.¹²

Efficiency based on profits is computed as

$$E_{i,t}^{\Pi} = \frac{\Pi_{i,t}}{\Pi_{t,an}^{eq}} 100\%, \quad (10)$$

where $\Pi_{i,t}$ is the profits of subject i in period t . In the profits treatment, $\Pi_{i,t}$ is the actual profit on which the subjects' performance is evaluated. In the FSE treatment, $\Pi_{i,t}$ is the profit that subject i would have received if the performance were evaluated based on profits, and it is computed based on equation (6). $\Pi_{t,an}^{eq}$ is the profit that can be obtained at the equilibrium employment corresponding to the announcement made in period t .

Efficiency based on the FSE is computed as

$$E_{i,t}^{FSE} = \frac{FP_{i,t}}{FP_{max}} 100\%, \quad (11)$$

¹²Bao et al. (2012) use similar efficiency measures to compare the performance of different treatments in an experimental N-firm cobweb economy.

where $FP_{i,t}$ is the forecasting payoff of subject i in period t . In the FSE treatment, $FP_{i,t}$ is the actual forecasting payoff on which the subjects' performance is evaluated. In the profits treatment, $FP_{i,t}$ is the forecasting payoff that subject i would have received if the performance were evaluated based on the forecasting accuracy, and it is computed based on equation (9). $FP^{max} = 8$ is the maximum forecasting payoff that can be obtained according to the payoff function in equation (9).

We compute these efficiency measures for each subject in each period and then pool the data from all the experimental sessions for each treatment. The bottom left panel of Figure 19 shows the empirical CDF of the efficiency based on the FSE for both treatments. This figure illustrates that the CDF for the profits treatment is larger than the CDF for the FSE treatment, which is statistically significant using the Kolmogorov-Smirnov test (with a p-value of 0 and a test statistic of 0.0981). This implies that the probability mass is larger in the higher values of the efficiency measure in the FSE treatment than in the profits treatment; i.e., there are more accurate forecasts in the FSE treatment. Because subjects are evaluated based on the accuracy of their forecasts in the FSE treatment, their forecasts are indeed more accurate.

The bottom right panel of Figure 19 shows the empirical cumulative distribution function of the efficiency based on the profits for both treatments. This figure illustrates that the CDF for the profits treatment is larger than the CDF for the FSE treatment, which is statistically significant using the Kolmogorov-Smirnov test (with a p-value of 0 and a test statistic of 0.8764). Thus, the probability mass is larger in the higher values of the efficiency measure in the FSE treatment; i.e., profits are higher in the FSE treatment. In the discussion of forecasts and employment outcomes, we have seen that the forecasts are more accurate and the employment outcomes are closer to the equilibrium values in the FSE treatment than in the profits treatment. Therefore, as employment outcomes are closer to the equilibrium values in the FSE treatment, profits are higher in the FSE treatment as well.

5 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 2.3. 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 the reported average employment of others, \bar{N}_t . In many of the experiments this was sufficient to generate convergence or approximate convergence to the SSE throughout the experimental session. 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 the 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, appears 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 $n_M = 12.5$, thereby implying that actual observations of average employment are less than the average forecast, which under adaptive learning pushes agents toward n_L rather than n_H . However, during the extended sequence of high announcements 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, which 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 2.3 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.

5.1 The role of heterogeneity in learning

Continuing with the previous 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. In particular, even if most agents have expectations near, say, the high steady state, if there are several agents that have sufficiently low expectations, this can be enough to destabilize coordination on an SSE.¹³ 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.¹⁴

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 parameterization, choosing high employment is relatively more profitable than choosing low employment when five out of six subjects choose high values (see Table 4). In contrast, if four or fewer 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

¹³More 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.

¹⁴“Direct criterion” versions of the adaptive learning rules, described in section 2.3, 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).

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.

Let us take a closer look at how learning happens during an experimental session and at the role of heterogeneity. For example, in session 4 of the profits treatment (Figure 6) 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 and 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 4 (Figure 12), 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 toward 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 toward the high equilibrium. When fewer than five subjects choose high values, the average employment of others is below their forecasts, which under adaptive learning pushes them toward low equilibrium.

In session 2 of the FSE treatment (Figure 10) 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 the 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 and 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 12), 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.

6 Conclusion

We have conducted experiments in a simple general-equilibrium 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. For the same reason, the FSE treatment demonstrates higher “efficiency” whether measured by forecast squared error or profits.

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 periods 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 for future research.

References

- [1] Adam, Klaus. “Experimental Evidence on the Persistence of Output and Inflation.” *Economic Journal* 117 (2007): 603-636.
- [2] Azariadis, Costas, and Roger Guesnerie. “Sunspots and Cycles.” *Review of Economic Studies* 53 (1986): 725-737.
- [3] Bao, Te, Hommes, Cars, and John Duffy. “Learning, Forecasting and Optimizing: an Experimental Study.” Tinbergen Institute Discussion Papers 12-015/1, Tinbergen Institute, 2012.
- [4] Cass, David, and Karl Shell. “Do Sunspots Matter?” *Journal of Political Economy* 91 (1983): 193-227.
- [5] Clarida, Richard, Gali, Jordi, and Mark Gertler. “Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory.” *Quarterly Journal of Economics* 115 (2000): 147-180.
- [6] Cooper, Russel, De Jong, Douglas V., Forsythe, Robert, and Thomas W. Ross. “Communication in Coordination Games.” *Quarterly Journal of Economics* 107 (1992): 739-771.
- [7] Corbae, Dean, and John Duffy. “Experiments with Network Formation.” *Games and Economic Behavior* 64 (2008): 81-120.

- [8] Duffy, John. “Macroeconomics: A Survey of Laboratory Research.” Working Paper, University of Pittsburgh, 2012.
- [9] Duffy, John, and Eric O.N. Fisher. “Sunspots in the Laboratory.” *American Economic Review* 95(2005): 510-529.
- [10] Evans, George, and Seppo Honkapohja. “Increasing Social Return, Learning and Bifurcation Phenomena.” In *Learning and Rationality in Economics*, edited by Alan Kirman and Mark Salmon, 216–235. Oxford: Blackwell, 1995.
- [11] Evans, George, and Seppo Honkapohja. “On the Local Stability of Sunspot Equilibria under Adaptive Learning.” *Journal of Economic Theory* 64 (1994): 142-161.
- [12] Evans, George, and Seppo Honkapohja. *Learning and Expectations in Macroeconomics*, Princeton: Princeton University Press, 2001.
- [13] Evans, George, Honkapohja, Seppo, and Ramon Marimon. “Convergence in Monetary Inflation Models with Heterogeneous Learning Rules.” *Macroeconomic Dynamics* 5 (2001): 1-31.
- [14] Evans, George, Honkapohja, Seppo, and Ramon Marimon. “Stable Sunspot Equilibria in a Cash-in-Advance Economy.” *The B.E. Journal of Macroeconomics, Advances* 7 (2007): Article 3.
- [15] Evans, George, Honkapohja, Seppo, and Paul Romer. “Growth Cycles.” *American Economic Review* 88 (1998): 495-515.
- [16] Farmer, Roger. *The Macroeconomics of Self-Fulfilling Prophecies*. 2nd edition. Cambridge: MIT Press, 1999.
- [17] Fehr, Dietmar, Heinemann, Frank, and Aniol Llorente-Saguer. “The Power of Sunspots: An Experimental Analysis.” SFB 649 Discussion Paper, Humboldt University, 2011.
- [18] Fischbacher, Urs. “Z-Tree: Zurich Toolbox for Ready-made Economic Experiments.” *Experimental Economics* 10 (2007): 171-178.
- [19] Garratt, Rod, and Todd Keister. “Bank Runs: An Experimental Study.” Working Paper, University of California, Santa Barbara, 2005.
- [20] Heemeijer, Peter, Hommes, Cars, Sonnemans, Joep, and Jan Tuinstra. “Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation.” *Journal of Economic Dynamics and Control* 33 (2009): 1052-1072.
- [21] Heinemann, Frank, Nagel, Rosemarie, and Peter Ockenfels. “The Theory of Global Games on Test: Experimental Analysis of Coordination Games with Public and Private Information.” *Econometrica* 72 (2004): 1583–1599.
- [22] Hommes, Cars, Sonnemans, Joep, Tuinstra, Jan, and Henk van de Velden. “Coordination of Expectations in Asset Pricing Experiments.” *Review of Financial Studies* 18 (2005): 955-980.
- [23] Hommes, Cars, Sonnemans, Joep, Tuinstra, Jan, and Henk van de Velden. “Expectations and Bubbles in Asset Pricing Experiments.” *Journal of Economic Behavior and Organization* 67 (2008): 116-133.

- [24] Lei, Vivian, and Charles N. Noussair. "Equilibrium Selection in an Experimental Macroeconomy." *Southern Economic Journal* 74 (2007): 448-482.
- [25] Lim, Suk, Prescott, Edward C., and Shyam Sunder. "Stationary Solution to the Overlapping Generations Model of Fiat Money: Experimental Evidence." *Empirical Economics* 19 (1994): 255-77.
- [26] Marimon, Ramon, and Shyam Sunder. "Indeterminacy of Equilibria in a Hyperinflationary World: Experimental Evidence." *Econometrica* 61 (1993): 1073-1107.
- [27] Marimon, Ramon, and Shyam Sunder. "Expectations and Learning under Alternative Monetary Regimes: An Experimental Approach." *Economic Theory* 4 (1994): 131-62.
- [28] Marimon, Ramon, and Shyam Sunder. "Does a Constant Money Growth Rule Help Stabilize Inflation?" *Carnegie-Rochester Conference Series on Public Policy* 43 (1995): 111-156.
- [29] Marimon, Ramon, Spear, Stephen E., and Shyam Sunder. "Expectationally Driven Market Volatility: An Experimental Study." *Journal of Economic Theory* 61 (1993): 74-103.
- [30] Ochs, Jack. "Coordination Problems." In *The Handbook of Experimental Economics*, edited by John H. Kagel and Alvin E. Roth, 195-251. Princeton: Princeton University Press, 1995.
- [31] Schotter, Andrew, and Tanju Yorulmazer. "On the Severity of Bank Runs: An Experimental Study." Working paper, New York University, 2003.
- [32] Van Huyck, John B., Battalio, Raymond C., and Richard O. Beil. "Tacit Coordination Games, Strategic Uncertainty, and Coordination Failure." *American Economic Review* 80 (1990): 234-48.
- [33] Woodford, Michael. "Learning to Believe in Sunspots." *Econometrica* 58 (1990): 277-307.

Steady state	Employment, n	Productivity, ψ	Profit, Π
Low	6.25	2.5	$2.5\sqrt{6.25} - 0.5*6.25 = 3.125$
High	16	4	$4\sqrt{16} - 0.5*16 = 8$

Table 1: Employment and profits in the steady states.

bins	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

Table 2: Percentage of observations in each range of values (bin). Explanation of the column titles: E, H = employment during periods with high-employment announcements; E, L = employment during periods with low-employment announcements; F, H = forecasts during periods with high-employment announcements; F, L = forecasts during periods with low-employment announcements.

Employment, periods with high-employment announcements				
Treatment	Average	Median	Std deviation	Skewness
FSE	14.46	16.00	3.51	-1.86
Profits	13.91	16.00	3.83	-1.36
Employment, periods with low-employment announcements				
Treatment	Average	Median	Std deviation	Skewness
FSE	6.42	6.25	1.27	7.34
Profits	7.20	6.25	2.81	2.73
Forecasts, periods with high-employment announcements				
Treatment	Average	Median	Std deviation	Skewness
FSE	14.35	16.00	3.21	-1.84
Profits	14.21	15.00	3.30	-1.06
Forecasts, periods with low-employment announcements				
Treatment	Average	Median	Std deviation	Skewness
FSE	6.66	6.25	1.37	4.78
Profits	7.85	7.00	3.14	1.51

Table 3: Descriptive statistics of the data on individual employment outcomes and forecasts collected in periods with high- and low-employment announcements for the FSE and profits treatments. Low-equilibrium employment is 6.25, and high-equilibrium employment is 16.

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

Table 4: Productivity and profits of subjects forecasting low and high values depending on the number of forecasters of each type.

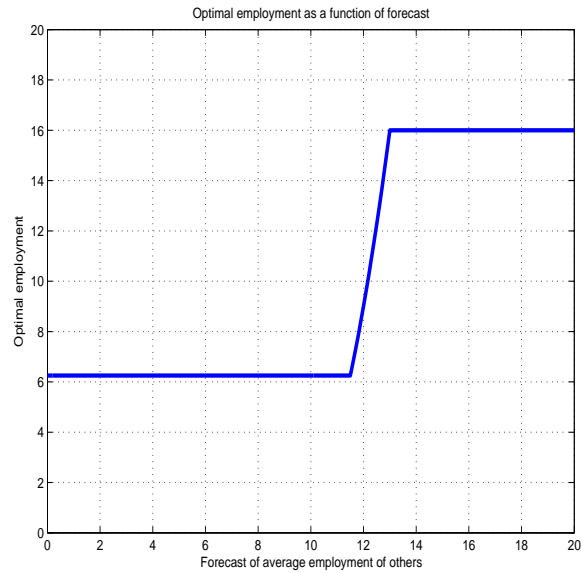


Figure 1: Employment as a function of the forecast of average employment of others according to equation (7).

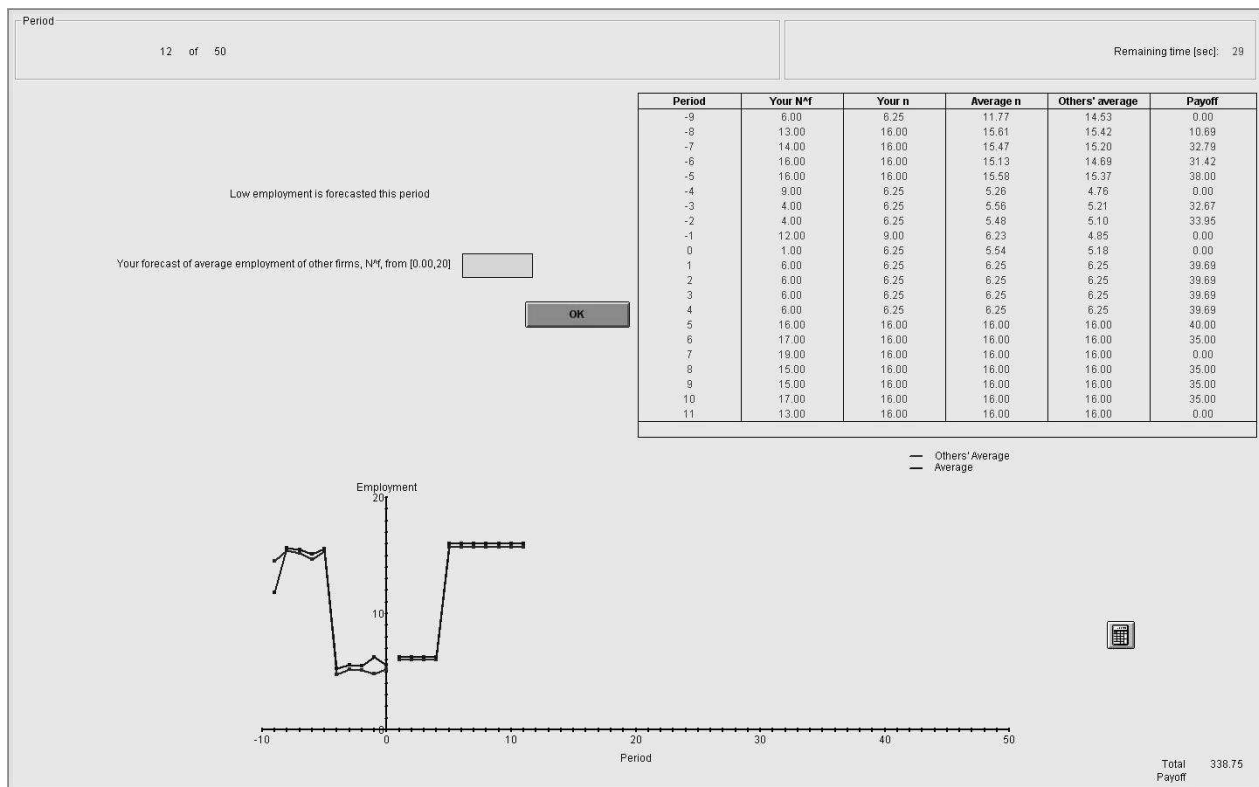


Figure 2: A screenshot.

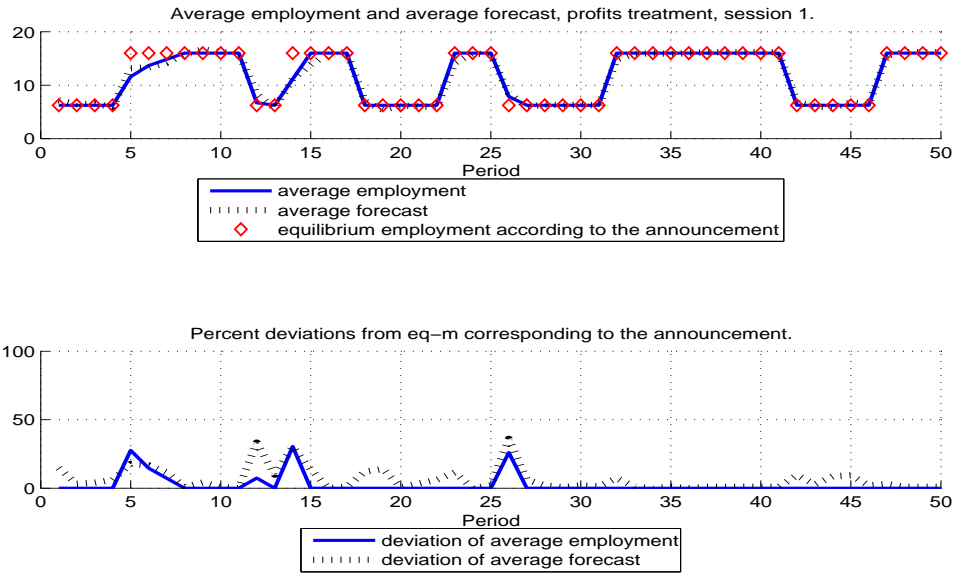


Figure 3: The first panel shows the average employment and average forecast, and the second panel shows the average percentage deviation of employment and average percentage deviation of forecasts from eq-m employment corresponding to the announcement in session 1 of the profits treatment.

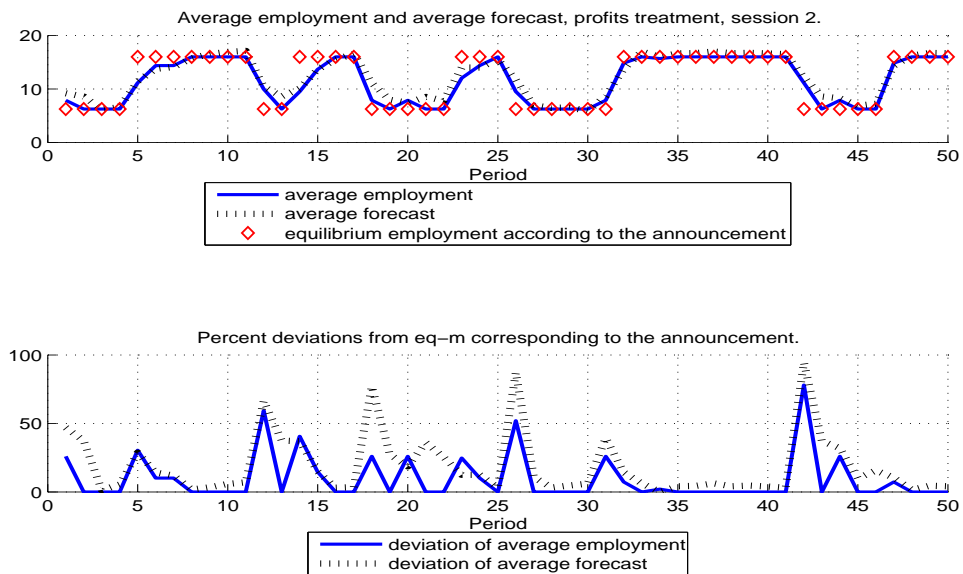


Figure 4: The first panel shows the average employment and average forecast, and the second panel shows the average percentage deviation of employment and average percentage deviation of forecasts from eq-m employment corresponding to the announcement in session 2 of the profits treatment.

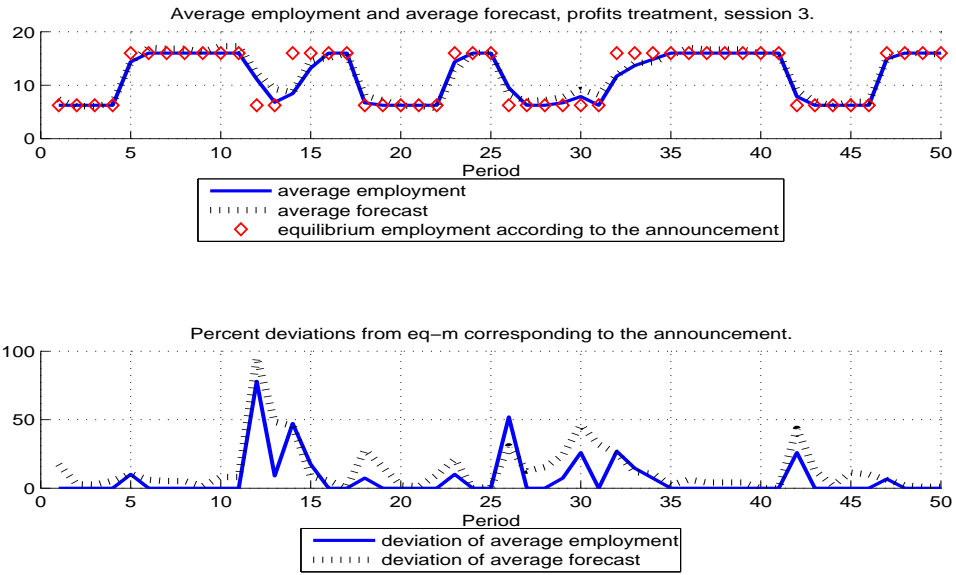


Figure 5: The first panel shows the average employment and average forecast, and the second panel shows the average percentage deviation of employment and average percentage deviation of forecasts from eq-m employment corresponding to the announcement in session 3 of the profits treatment.

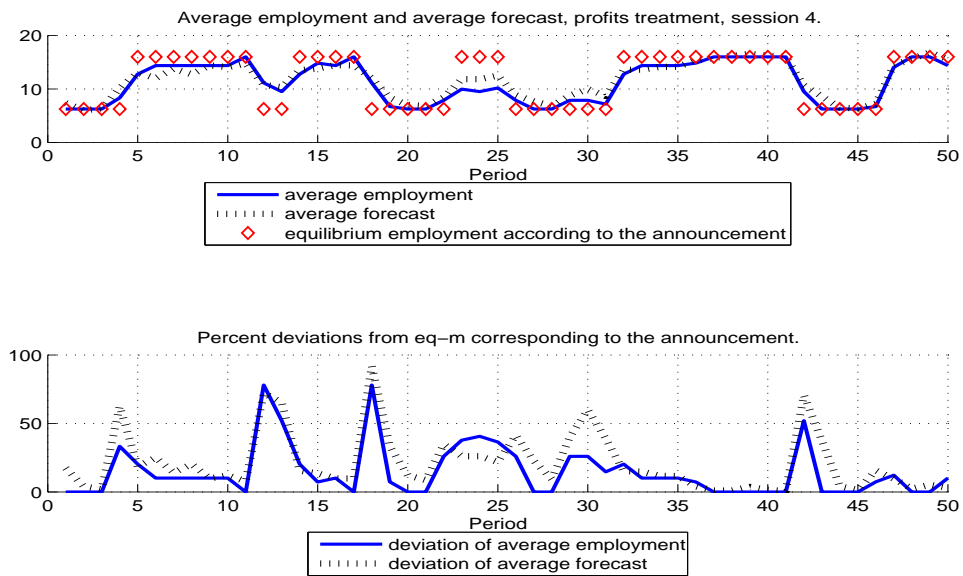


Figure 6: The first panel shows the average employment and average forecast, and the second panel shows the average percentage deviation of employment and average percentage deviation of forecasts from eq-m employment corresponding to the announcement in session 4 of the profits treatment.

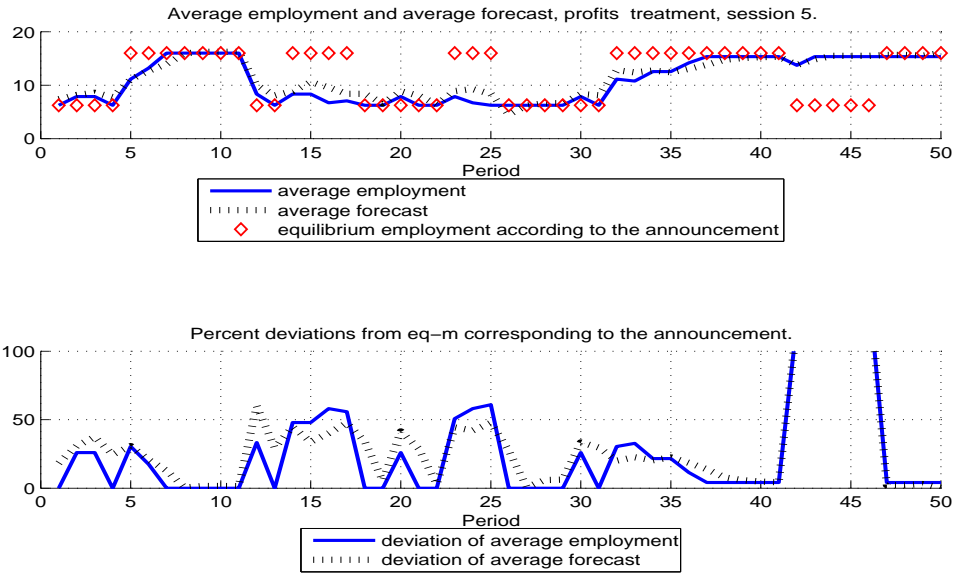


Figure 7: The first panel shows the average employment and average forecast, and the second panel shows the average percentage deviation of employment and average percentage deviation of forecasts from eq-m employment corresponding to the announcement in session 5 of the profits treatment.

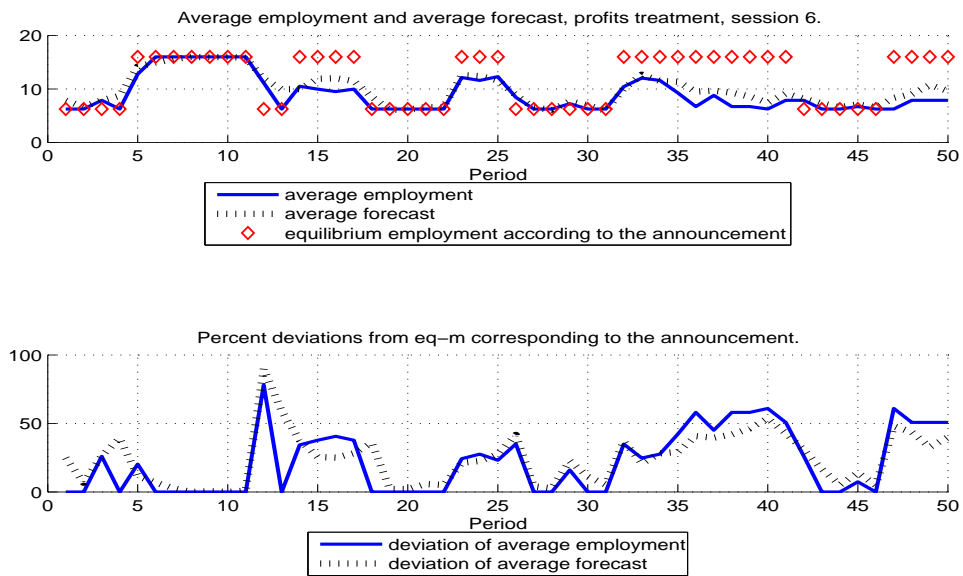


Figure 8: The first panel shows the average employment and average forecast, and the second panel shows the average percentage deviation of employment and average percentage deviation of forecasts from eq-m employment corresponding to the announcement in session 6 of the profits treatment.

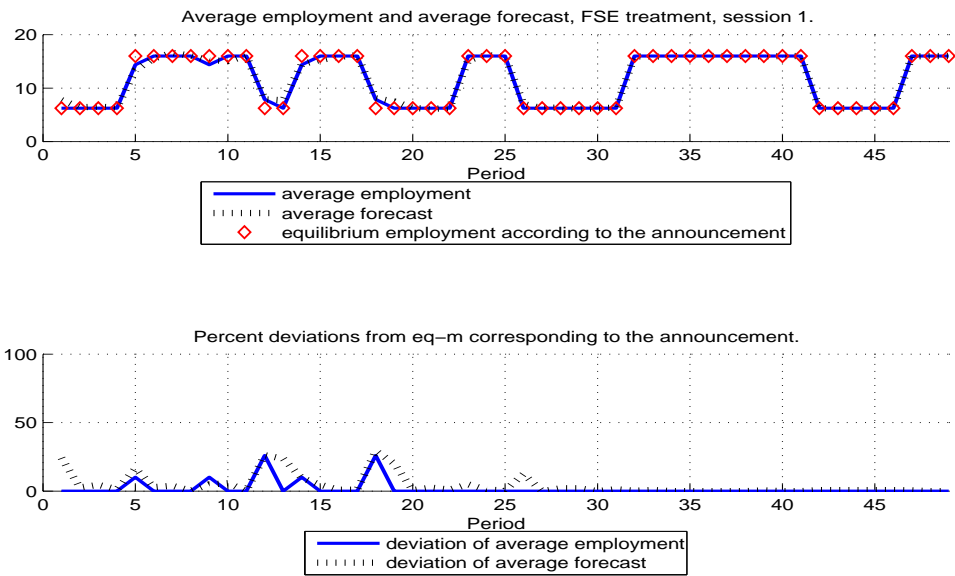


Figure 9: The first panel shows the average employment and average forecast, and the second panel shows the average percentage deviation of employment and average percentage deviation of forecasts from eq-m employment corresponding to the announcement in session 1 of the FSE treatment.

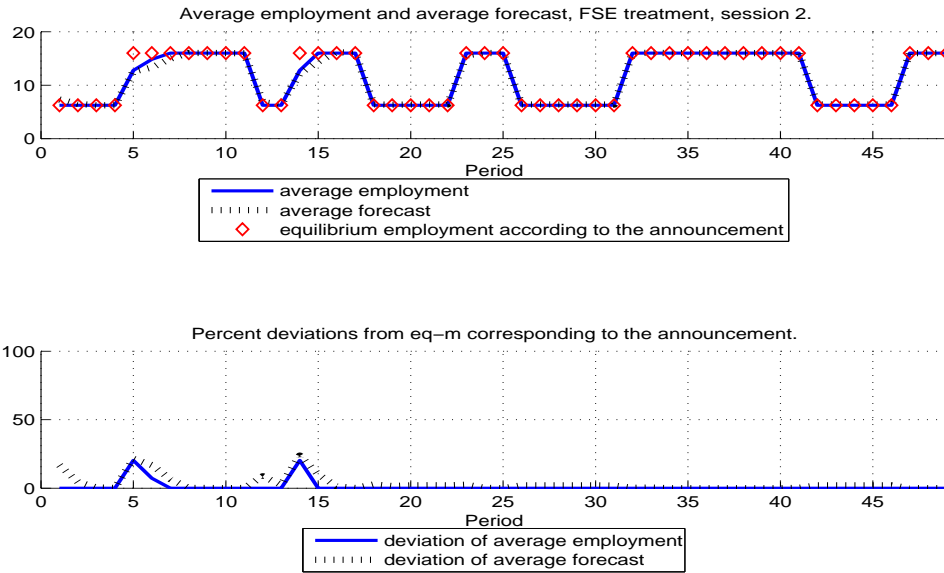


Figure 10: The first panel shows the average employment and average forecast, and the second panel shows the average percentage deviation of employment and average percentage deviation of forecasts from eq-m employment corresponding to the announcement in session 2 of the FSE treatment.

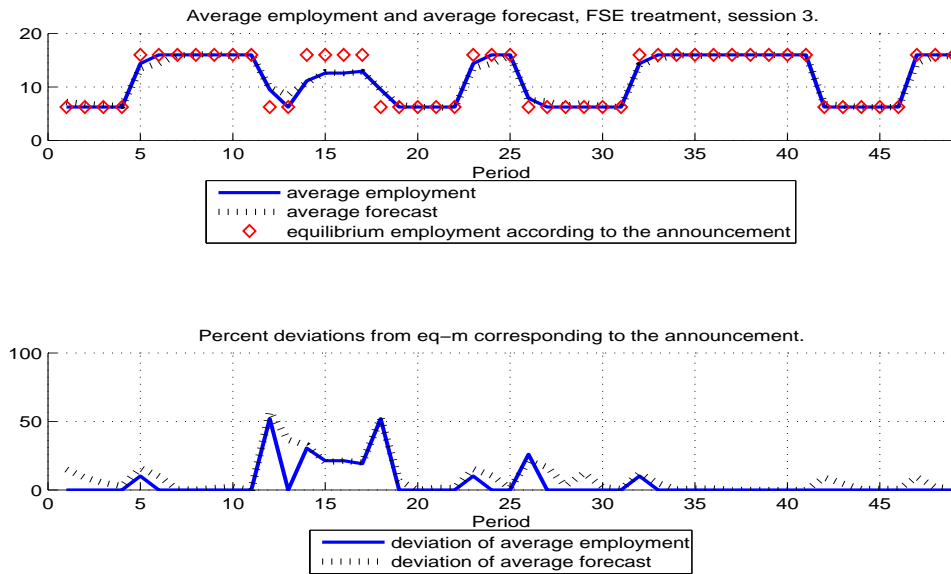


Figure 11: The first panel shows the average employment and average forecast, and the second panel shows the average percentage deviation of employment and average percentage deviation of forecasts from eq-m employment corresponding to the announcement in session 3 of the FSE treatment.

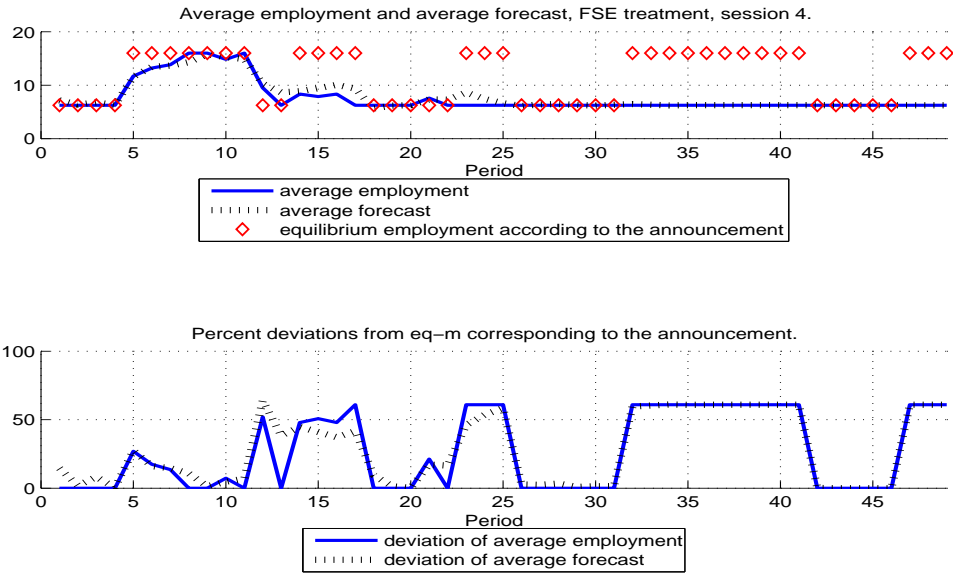


Figure 12: The first panel shows the average employment and average forecast, and the second panel shows the average percentage deviation of employment and average percentage deviation of forecasts from eq-m employment corresponding to the announcement in session 4 of the FSE treatment.

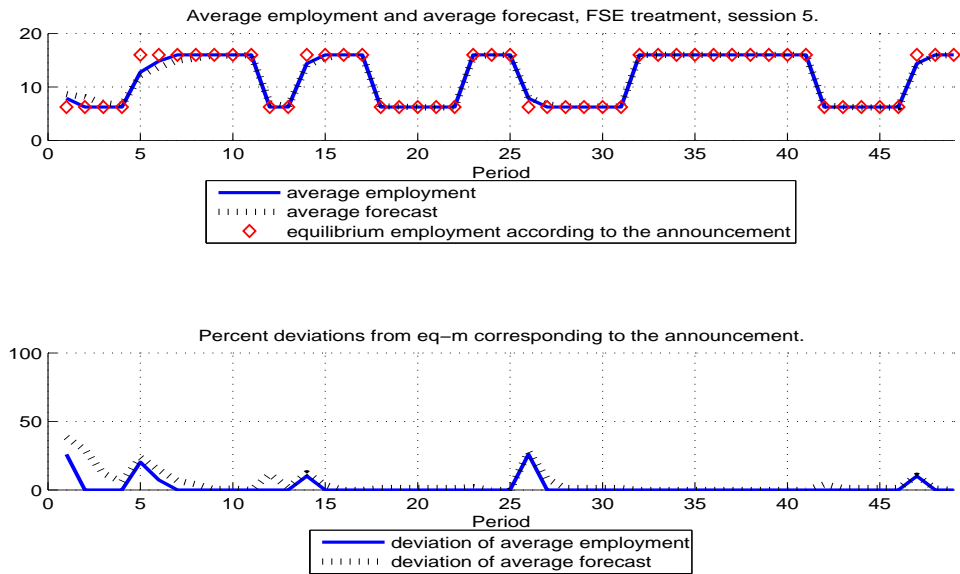


Figure 13: The first panel shows the average employment and average forecast, and the second panel shows the average percentage deviation of employment and average percentage deviation of forecasts from eq-m employment corresponding to the announcement in session 5 of the FSE treatment.

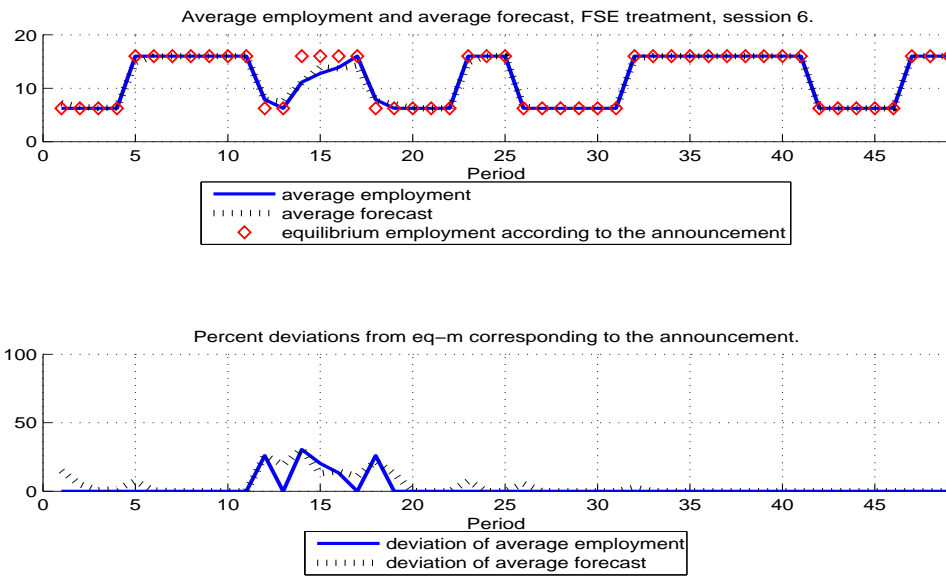


Figure 14: The first panel shows the average employment and average forecast, and the second panel shows the average percentage deviation of employment and average percentage deviation of forecasts from eq-m employment corresponding to the announcement in session 6 of the FSE treatment.

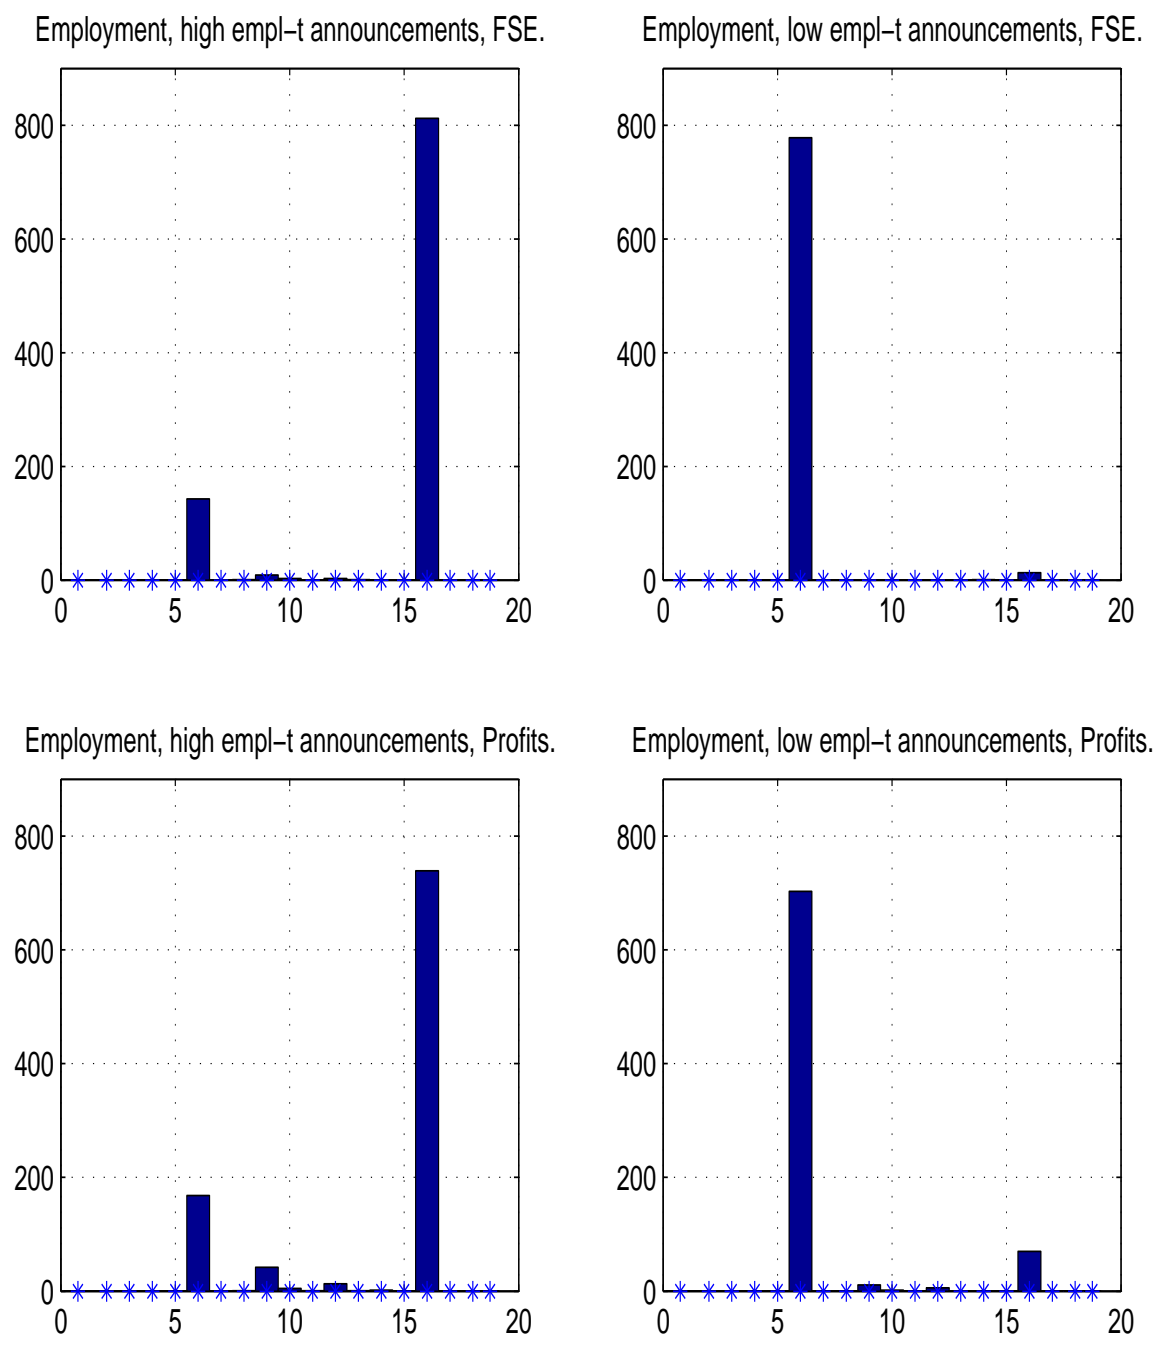


Figure 15: Histogram of individual employment during periods with high-employment and low-employment announcements, collected from all experimental sessions in FSE and profits treatments.

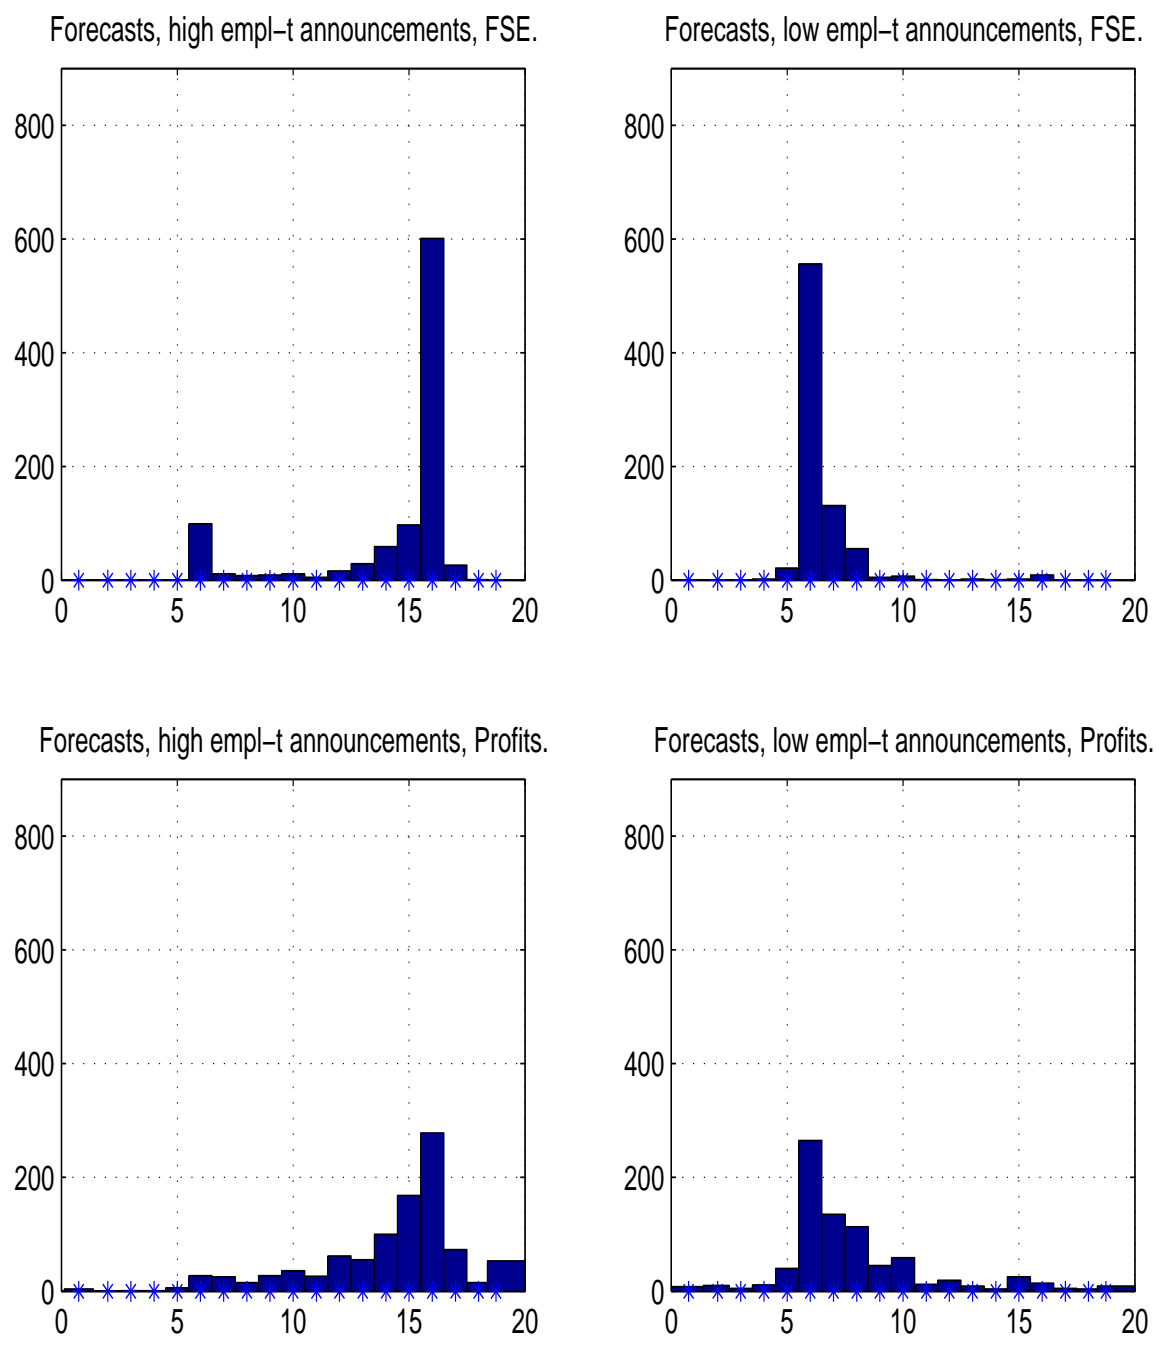


Figure 16: Histogram of individual forecasts during periods with high-employment and low-employment announcements, collected from all experimental sessions in FSE and profits treatments.

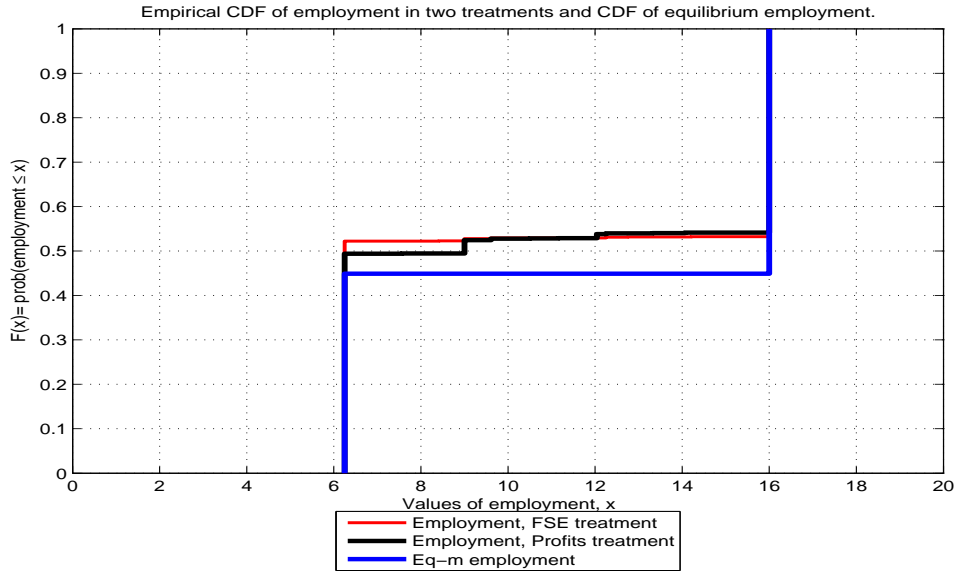


Figure 17: Empirical cumulative distribution function of employment outcomes in two treatments and theoretical CDF of equilibrium employment.

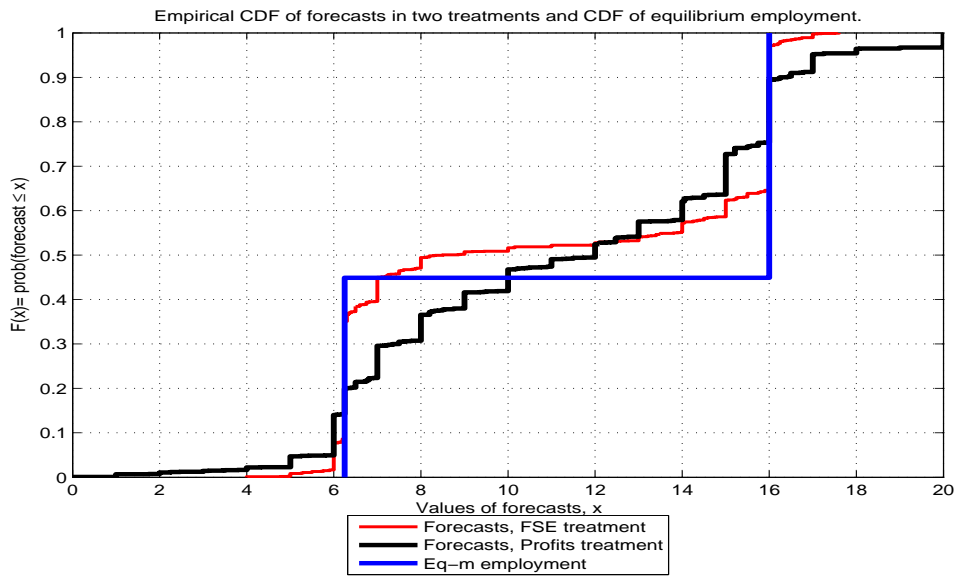


Figure 18: Empirical cumulative distribution function of employment outcomes in two treatments and theoretical CDF of equilibrium employment.

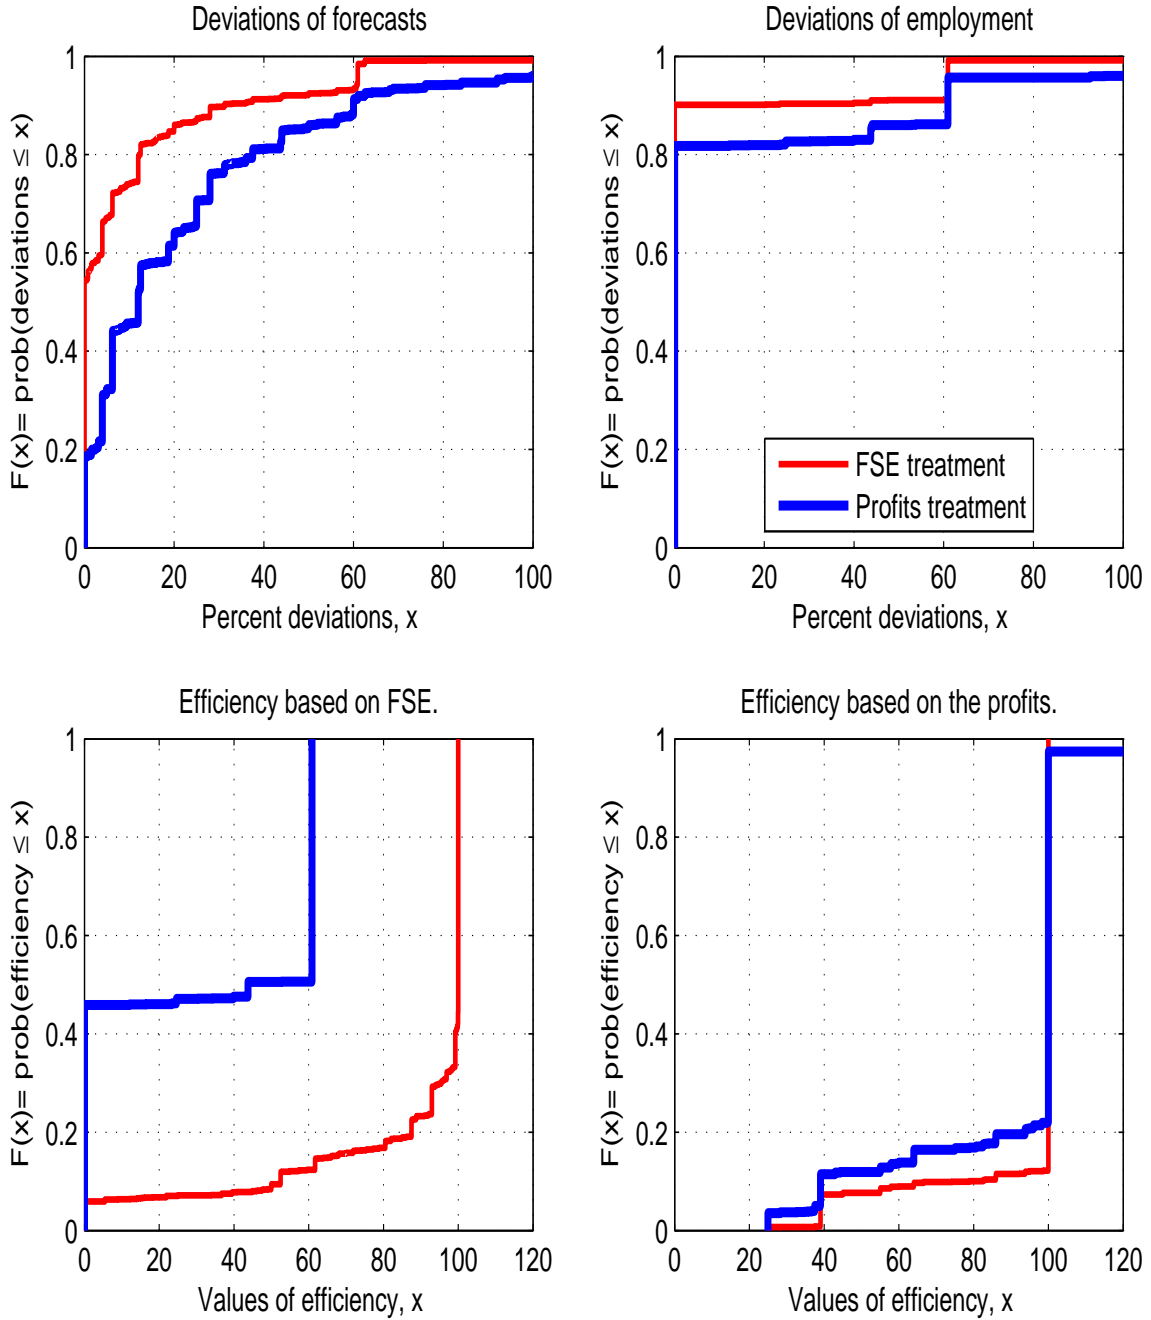


Figure 19: Empirical cumulative density functions of percentage deviations of individual forecasts from equilibrium employment corresponding to the announcements (top left panel), percentage deviations of individual employment from equilibrium employment corresponding to the announcements (top right panel), efficiency based on the FSE payoff (bottom left panel) and efficiency based on the profit (bottom right panel) in both FSE and profits treatments.